WASP WINTER CONFERENCE 2022

Poster Catalouge

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

WASP WINTER CONFERENCE 2022

AI MATH

Agerberg, Jens
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Data, geometry and homology

Homology-based invariants can be used to characterize the geometry of datasets and thereby gain some understanding of the processes generating those datasets. In this work we investigate how the geometry of a dataset changes when it is subsampled in various ways. In our framework the dataset serves as a reference object; we then consider different points in the ambient space and endow them with a geometry defined in relation to the reference object, for instance by subsampling the dataset proportionally to the distance between its elements and the point under consideration. We illustrate how this process can be used to extract rich geometrical information, allowing for example to classify points coming from different data distributions.

AI MATH | Page 1 B

Agerberg, Jens
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Data, geometry and homology

Jens Agerberg, KTH
Math department, Math of data and AI
Joint work with Wojciech Chachólski and Ryan Ramanujam

KTH VETENSKAP OCH KONST

Abstract

Homology-based invariants can be used to characterize the geometry of datasets and thereby gain some understanding of the processes generating those datasets. In this work (under review) we investigate how the geometry of a dataset changes when it is subsampled in various ways. In our framework the dataset serves as a reference object; we then consider different points in the ambient space and endow them with a geometry defined in relation to the reference object, for instance by subsampling the dataset proportionally to the distance between its elements and the point under consideration. We illustrate how this process can be used to extract rich geometrical information, allowing for example to classify points coming from different data distributions.

Methods

Persistent homology: from point clouds to persistence modules
From a point cloud we can construct a Vietoris-Rips complex, a combinatorial object encoding its geometry, parametrized by $\epsilon \in [0, \infty)$.

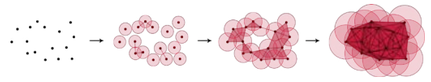

By taking homology we get (for each homological degree) a vector space for each ϵ and a linear map for each $\tau \leq \epsilon \in [0, \infty)$. These linear maps are called *persistence modules* and are the output of persistent homology.

Metrics and machine learning: from persistence modules to stable ranks
Persistence modules can be seen as a summary of geometrical aspects of the point cloud. To be useful we need metrics to compare them and ways to develop machine learning algorithms on them.
For this we use the framework of *stable rank*: persistence modules have a discrete invariant called rank, this invariant can be stabilized by considering instead the minimum rank in a growing neighborhood of the module, leading to a type of homology-based invariant in the form of a non-increasing piecewise constant function. Since this function space is a Hilbert space one can consider a kernel based on stable rank [1] for use in machine learning.

From global to local
Homology-based invariants are often used to characterize global aspects of a dataset. In this work, we instead investigate whether they can be useful in describing a single point in the ambient space, by subsampling a dataset (called reference object) according to the distance of its members to the point:

1. Choose a reference object: a finite subset $\mathcal{R} \subset \mathbb{R}^N$ and a point $p \in \mathbb{R}^N$.
2. Attach a probability distribution to $\mathcal{R}$. We are interested here in distributions that attach high probability to points $r \in \mathcal{R}$ which have low distance to p and low probability to more remote points.
3. Sample s points from the reference object according to the probability distribution. Repeat n times and each time compute persistence modules and stable ranks.
4. Average the stable ranks to get a descriptor characterizing the point p.

Selected Results

We start with data consisting of random points on the plane. For each point, the reference object (the green points, sampled from a circle) is sampled relative to its distance to the point, Persistent homology and stable ranks are computed. The stable ranks clearly group into orange (for points inside the circle) and blue (outside the circle), indicating that interesting geometric properties can be found.

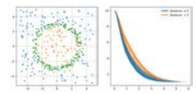

We now use as reference object the MNIST dataset for digit 1. We select two points from the ambient space, $\mathbb{R}^{784}$: the origin and the center of mass of the reference object. Using dimensionality reduction, we can illustrate what it means to sample the reference object relative to those 2 points. Now the stable ranks resulting from the sampling allow to distinguish the 2 points, for all homological degrees.

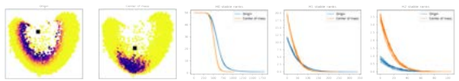

We now use as reference object the union of MNIST training sets for digit 1 and 7. We select out-of-sample 1:s and 7:s and represent them by their stable rank, obtained by sampling the reference object in the same way as before. In many cases, the geometry of the reference object close to the out-of-sample digits allow to distinguish them. This is further quantified by training an SVM classifier based on the stable rank kernel (in a semi-supervised learning setup: the reference object is used in an unlabeled fashion and the SVM is only trained on 10 samples from each class).

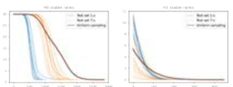

Interestingly, to distinguish a pair of digits, one can also use other digits as reference object. Here 2:s and 3:s are used as reference object to distinguish 1:s and 7:s.

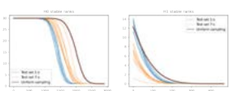

References

[1] Jens Agerberg, Ryan Ramanujam, Martina Scolamiero, and Wojciech Chachólski. Supervised learning using homology stable rank kernels. *Frontiers in Applied Mathematics and Statistics*, 7:39, 2021.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Homogeneous vector bundles and G-equivariant convolutional neural networks

G-equivariant convolutional neural networks (GCNNs) is a geometric deep learning model that uses global symmetry to improve learning. Most GCNNs use convolutional layers to transform data in a translation equivariant manner, like the sliding kernels of CNNs but generalized to other symmetries, e.g. rotation-equivariant transformations of spherical data. We analyze GCNNs and classify those G-equivariant layers that are expressible as convolutional layers. That is, we characterize the expressivity of convolutional layers.

UNIVERSITY OF GOTHENBURG

Equivariant Neural Networks

Jimmy Aronsson
Chalmers University of Technology
Department of Mathematical Sciences

CHALMERS

Homogeneous vector bundles and G-equivariant convolutional neural networks

G-equivariant convolutional neural networks (GCNNs) is a geometric deep learning model that uses global symmetry to improve learning. Most GCNNs use convolutional layers to transform data in a translation equivariant manner, like the sliding kernels of CNNs but generalized to other symmetries, e.g. rotation-equivariant transformations of spherical data. We analyze GCNNs and classify those G-equivariant layers that are expressible as convolutional layers. That is, we characterize the expressivity of convolutional layers.

Suppose that we are given data defined on a homogeneous space $\mathcal{M}$, e.g. meteorological wind vector fields on the rotation symmetric sphere S^2, or digital images defined on the translation symmetric pixel lattice $\mathbb{Z}^2$. The general idea is to transform such data in *translation equivariant* ways, preserving the global symmetry. A special case is translation *invariant* layers, which produce the same output for all translations of the input. Such layers are useful when classifying images, for example, as they make the same class prediction no matter where objects are located within an image.

Globally symmetric spaces are also called homogeneous:

Definition. Let G be a Lie group. A smooth manifold $\mathcal{M}$ is a *homogeneous space* if there exists a smooth, transitive action

$$G \times \mathcal{M} \to \mathcal{M}, \qquad (g, x) \mapsto g \cdot x.$$

Any homogeneous space $\mathcal{M}$ is diffeomorphic to a quotient space G/K for some closed subgroup $K \leq G$. Examples of homogeneous spaces include the Euclidean spaces $\mathbb{R}^n$, lattices $\mathbb{Z}^n$, spheres $S^n \simeq SO(n+1)/SO(n)$, and many others.

Vector bundles over a homogeneous space $\mathcal{M} = G/K$ often inherit its global symmetry, i.e. there often exists a smooth, transitive action $G \times E \to E$. Such vector bundles are also called *homogeneous* and are uniquely characterized by some K-representation ρ. We write $E = E_\rho$.

Data points are viewed as functions that attach a feature vector/scalar at each point of a homogeneous space. They are formalized as sections of homogeneous vector bundles:

Definition. A *data point* is a square-integrable section $s : \mathcal{M} \to E_\rho$. We denote the space of all data points by $L^2(E_\rho)$.

The global symmetry in G induces a representation $\mathrm{Ind}_K^G \rho$ on $L^2(E_\rho)$ that performs translations of data points:

$$\big(\mathrm{Ind}_K^G \rho(g) s\big)(x) = g \cdot s(g^{-1} x)$$

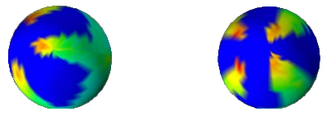

$s \in L^2(E_\rho)$ $\qquad$ $\mathrm{Ind}_K^G \rho(g) s \in L^2(E_\rho)$

Definition. Let E_ρ and E_σ be homogeneous vector bundles over $\mathcal{M}$. A linear transformation $\Phi : L^2(E_\rho) \to L^2(E_\sigma)$ is called a *G-equivariant layer* if it intertwines the induced representations:

$$\mathrm{Ind}_K^G \sigma \circ \Phi = \Phi \circ \mathrm{Ind}_K^G \rho.$$

For example, if Φ produces bounding boxes around objects in an image, then translating the image will also translate the bounding boxes.

Implementations of GCNNs primarily use *convolutional layers*[1]

$$(\Phi s)(g) = \int_G \kappa(g^{-1} g') s(g')\, dg,$$

where κ is a matrix-valued kernel. All convolutional layers are G-equivariant layers but the latter notion is much more general, so implementations that only use convolutional layers could be unnecessarily restrictive. It is thus interesting to analyze the relation between G-equivariant and convolutional layers.

Our main theorem characterizes those G-equivariant layers that are expressible as convolutional layers. It does so for extremely general homogeneous spaces $\mathcal{M} \simeq G/K$, including the Euclidean spaces, grids, spheres, and even Minkowski spacetime (which is homogeneous with respect to the Poincaré group).

Theorem. Consider a homogeneous space $\mathcal{M} = G/K$ with G a unimodular Lie group of type I and $K \leq G$ a compact subgroup. Let E_ρ, E_σ be homogeneous vector bundles over $\mathcal{M}$ and let

$$\Phi : L^2(E_\rho) \to L^2(E_\sigma),$$

be a G-equivariant layer. If Φ maps into a space of bandlimited functions, then Φ is a convolutional layer.

The following corollary is especially useful, since implementations of GCNNs primarily use finite or discrete groups.

Corollary. If G is either discrete abelian or finite, then any G-equivariant layer is a convolutional layer.

Using only convolutional layers is therefore not a restriction, when G is either discrete abelian or finite; Implementations based on these layers are maximally expressive. For instance, convolutional layers

$$[\kappa \star s](x) = \sum_{y \in \mathbb{Z}^2} \kappa(y - x) s(y), \qquad (\mathcal{M} = G = \mathbb{Z}^2)$$

are the only possible equivariant transformations of digital images s.

[1] Convolutional layers actually transform *feature maps* rather than data points, but these objects are equivalent. See [1] for details.

References

[1] Aronsson, J.. "Homogeneous vector bundles and G-equivariant convolutional neural networks." *arXiv preprint arXiv:2105.05400* (2021).

[2] Gerken, J.E., Aronsson, J., et al. "Geometric deep learning and equivariant neural networks." *arXiv preprint arXiv:2105.13926* (2021).

[3] Cohen, T. and Welling, M. "Group equivariant convolutional networks." *International conference on machine learning*. PMLR, 2016.

[4] Kondor, R., and Trivedi, S. "On the generalization of equivariance and convolution in neural networks to the action of compact groups." *International Conference on Machine Learning*. PMLR, 2018.

Bengtsson Bernander, Karl
Uppsala University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Robust learning of geometric equivariances

We extend convolutional neural networks (CNNs) to provide rotation equivariance. We evaluate on the oral cancer dataset to diagnose malignant cancer, using the VGG16 classifier architecture. We also evaluate on the BBB038 dataset of highly varied cell nuclei, this time using the U-net architecture combined with a discriminative loss function for semantic instance segmentation. We expect that incorporating rotation equivariance into CNNs will increase the expressive capacity without increasing the number of parameters, reducing overfitting. Also, since data augmentation can be reduced, misclassification due to interpolation artifacts should decrease. The results indicate that this holds for the classifier network, but more experiments are needed to verify this for the semantic instance segmentation network.

Bengtsson Bernander, Karl
Uppsala University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Centrum för bildanalys, Uppsala universitet

UPPSALA UNIVERSITET

Robust learning of geometric equivariances

Karl Bengtsson Bernander, Joakim Lindblad, Nataša Sladoje, Robin Strand, Ingela Nyström
Centre for Image Analysis, Department of Information Technology, Uppsala University, Sweden

Abstract

We extend convolutional neural networks (CNNs) to provide rotation equivariance. We evaluate on the oral cancer dataset to diagnose malignant cancer, using the VGG16 classifier architecture. We also evaluate on the BBB038 dataset of highly varied cell nuclei, this time using the U-net architecture combined with a discriminative loss function for semantic instance segmentation. We expect that incorporating rotation equivariance into CNNs will increase the expressive capacity without increasing the number of parameters, reducing overfitting. Also, since data augmentation can be reduced, misclassification due to interpolation artifacts should decrease. The results indicate that this holds for the classifier network, but more experiments are needed to verify this for the semantic instance segmentation network.

Classification on the oral cancer dataset

One feature of standard convolutional neural networks (CNNs) is translational invariance: the result of convolving an input with a filter and then shifting the output is identical to shifting the input and then applying the convolution. We are interested in other equivariances, such as rotations and scaling. Recent works on rotation equivariance in CNNs include:

- **Group-equivariant convolutional networks (G-CNNs) [1], using group-convolutions.**
- **General E(2)-Equivariant Steerable CNNs [2], available as a library in Pytorch.**

$$T_g\Phi(f)=\Phi(T_g f)$$

Rotational equivariance: filtering (Φ) an input , then rotating (T_g), gives the same result as filtering on the rotated input.

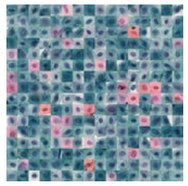

Microscopy image of cells from the oral cavity

The oral cancer dataset. We modified the VGG16 classifier to use group-equivariant convolutions on the p4 group, consisting of translations and rotations of multiples of 90 degrees.

The baseline CNN version combined with data augmentation of rotations of multiples of 90 degrees yield an accuracy score of around 56 %. The equivariant version , without data augmentation, yields 60 %. The latter architecture is less sensitive to overrfitting.

Semantic Instance Segmentation

We further modify the U-net architecture with a discriminative loss function [3] to be equivariant to rotations of multiples of 90 degrees. The methods yield a DICE score of about 0.7 for both versions of the U-net.

Instance segmentation on the modified BBBC038 dataset. The left image shows the input image, the middle one the results from the baseline U-net architecture, and the right one the results from the U-net architecture modified to be equivariant to rotations of multiples of 90 degrees.

Further directions

We plan to train the instance segmentation network for hundreds of more epochs to verify the hypothesis. That is, that ordinary CNN architectures, combined with data augmentation of multiples of rotations of 90 degrees, can be replaced with networks that are equivariant by design to those same transformations. We also plan to use another clustering method than K-means, preferably one without a predetermined number of clusters.

Both the instance segmentation network and the classifier networks can be tested on other datasets, and with other symmetry groups.

For larger datasets, moving to distributed training over multiple GPUs show promise for speeding up the training phase. Moving to a cloud computational environment could also allow for more flexibility, with the drawback that you lose some control over your own development environment.

References

1. T.S. Cohen, M. Welling. Group Equivariant Convolutional Networks. Proceedings of the International Conference on Machine Learning (ICML), 2016
2. Maurice Weiler and Gabriele Cesa. General E(2)-equivariant´steerable CNNs. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019.
3. Bert De Brabandere, Davy Neven, and Luc Van Gool.Semantic instance segmentation with a discriminative loss´function. In IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW 2017), 2017.

Bozorgpanah, Aso
Doctoral student, Department of Computing Science

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

The Interpretable Protected Machine Learning Model with Privacy

Machine learning (ML) models have the potential to enhance products. It is a type of Artificial Intelligence (AI) that allows software applications to predict outcomes.
Data-driven models built using ML have proven their usefulness. Nevertheless, ML algorithms do not explain their predictions, which is a barrier to ML adoption. To address this issue, the researcher uses eXplainable Artificial Intelligence (XAI). XAI explains why a ML model yields a predicted output for a certain input.
Understanding why a model makes a prediction is important, but it is not enough. So, other principles need to be addressed for ML deployment in the real world. In the current work, privacy is one of the challenges that is discussed.
We studied the effect of data privacy techniques^{[1]} on SHapley Additive exPlanations (SHAP)^{[2]}.
By applying SHAP the output of any ML model can be explained. The output model is interpretable.
Our aim is to understand how data protection affects the measures related to explainability. Hence, we performed a series of experiments comparing the effects of data masking procedures on the explainability of models according to SHAP on the data set.

Bozorgpanah, Aso
Doctoral student, Department of Computing Science

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

The Interpretable Protected Machine Learning Model with Privacy

Aso Bozorgpanah, Ph.D., Umeå University
Dept. Computing Science, NAUSICA: PrivAcy-AWare traNSparent deCIsions group
Supervisors: Prof. Vicenç Torra (Umu), Associate professor. Lili Jiang (Umu)

UMEÅ UNIVERSITET

WASP WALLENBERG AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Motivation & Research Goals

Machine learning (ML) models have the potential to enhance products. It is a type of Artificial Intelligence (AI) that allows software applications to predict outcomes. Data-driven models built using ML have proven their usefulness. Nevertheless, ML algorithms do not explain their predictions, which is a barrier to ML adoption. To address this issue, the researcher uses eXplainable Artificial Intelligence (XAI). XAI explains why a ML model yields a predicted output for a certain input. Understanding why a model makes a prediction is important, but it is not enough. So, other principles need to be addressed for ML deployment in the real world. In the current work, privacy is one of the challenges that is discussed. We studied the effect of data privacy techniques[1] on SHapley Additive exPlanations (SHAP)[2]. By applying SHAP the output of any ML model can be explained. The output model is interpretable. Our aim is to understand how data protection affects the measures related to explainability. Hence, we performed a series of experiments comparing the effects of data masking procedures on the explainability of models according to SHAP on the data set.

Methods

An implications' analysis of applying data privacy techniques to explainability was performed. It is claimed[3] privacy and explainability are incompatible. While we designed an explainable model along with privacy. In this regard, Maximum Distance to Average Vector (MDAV) was applied for achieving microaggregation. The MDAV is a masking method that provides k-anonymity to protect data[4]. Microaggregation is one of the most efficient approaches in relation to the trade-off risk-utility. It consists of building small clusters with the original data and then replacing each of the data with a cluster center that is representative of the whole cluster. Microaggregation is flexible and permits implementing k-anonymity for any kind of data. We supposed $k = [1, 15]$. Although the range of k is different for various datasets, the k value should be selected in a reasonable range to have high accuracy. After masking the dataset by MDAV. SHapley Additive exPlanations (SHAP) was done on the masked dataset. SHAP[2] is a method to explain individual predictions. It is based on the Shapley Value of game theory. TreeSHAP is an estimation approach of SHAP that was used. TreeSHAP defines the value function in terms of the conditional expectation to estimate effects instead of the marginal expectation.

As the above progress is shown, we present a privacy-preserving explainable ML model. The explainable machine learning algorithms were applied to the protected data to train machine learning models and explain the result of their predictions. They were compared with the one obtained without masking.

Selected Results

A baseline model was trained on the original dataset, then, additional models were trained on the masked datasets. The explainable models were not changed even after protecting data for $k = [1, 12]$. The results showed that explainability for the protected model by MDAV was similar to the one obtained with the original data. Therefore, decisions on the amount of distortion to achieve protection through microaggregation and k-anonymity should be led by the desired trade-off between disclosure risk and model accuracy.
We presented an approach, what kind of data privacy methods are more feasible to explainability after applying SHAP to make an explainable ML model. The ML models trained on masked data were evaluated by their results explainability. We considered feature importance analysis of the final models (based on decision trees) using SHAP. Our approach were applied on 'USA Housing' dataset, and the results were compared between the results for the original and the masked data. The results for $k = [3, 6, 11]$ are shown in (b), (c), and (d) respectively in the below figure. It is clear that the extracted explainability are similar among all four models.

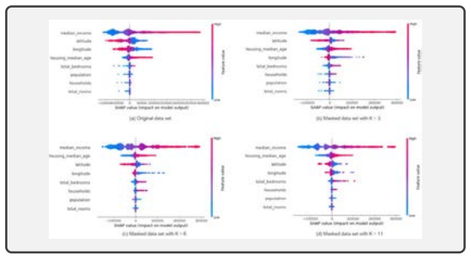

We found that interpretability using SHAP is studied for k-anonymous data. The results showed that qualitative properties of attributes were maintained for masked data. Then, the decision on which level of privacy and the amount of distortion was appropriate needs to focus on the risk-utility trade-off. For instance, a user needs to take into account both the value of k and the utility of the masked data set.

The explainable ML models can be considered along with privacy. We found that how explainability can be affected by data privacy methods and masking methods keep utility. Future research should seek to address other XAI requirements within a privacy-preserving framework to assess to what extent these tools apply in privacy-by-design MI.

References

[1] Data privacy: Foundations, new developments and the big data challenge
Torra, Vicenç
Springer, 2017

[2] A unified approach to interpreting model predictions
Lundberg, Scott and Lee, Su-In
arXiv preprint arXiv:1705.07874

[3] Show Us the Data: Privacy, Explainability, and Why the Law Can't Have Both
Grant, Thomas D and Wischik, Damon J
George Washington Law Review Arguendo, 2020

[4] Efficient k-anonymous microaggregation of multivariate numerical data via principal component analysis
Monedero, David Rebollo and Mezher, Ahmad Mohamad and Colomé, Xavier Casanova and Forné, Jordi and Soriano, Miguel
Information Sciences, Elsevier, 2019

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Breitholtz, Adam
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Data dependent bounds for domain adaptation

The study of generalization is one of the cornerstones of machine learning theory. Tight generalization bounds are potential tools for guaranteeing adequate performance and the PAC-Bayes framework has proven useful in deriving such bounds when good model priors are known and test cases match training cases in distribution.
However, in real world tasks, where deep neural networks are the models of choice and training and test cases come from different domains, deriving tight and estimable bounds remains an unresolved challenge.
In our work, we combine recent advances in PAC-Bayes domain adaptation with data-dependent priors to give estimable and informative bounds for problems where classical bounds are vacuous. We apply this method to a domain adaptation image classification task and find that it produces tighter bounds. We study which terms dominate the bounds and identify possible directions for further improvement.

Breitholtz, Adam
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Data Dependent Priors for Domain Adaptation Bounds

Adam Breitholtz, PhD, Chalmers University of Technology
Dept. Computer Science and Engineering, Data Science and AI division
Supervisors: Ass.Prof. Fredrik D. Johansson (CTH) and Prof. Devdatt Dubhashi (CTH)

UNIVERSITY OF TECHNOLOGY

Motivation & Research Goals

The study of generalization is a cornerstone of machine learning theory. Our understanding of how generalization functions is crucial to confidently engineer and deploy models in high stakes, real world domains, such as healthcare. Tight generalization bounds are potential tools for guaranteeing adequate performance and the PAC-Bayes framework has proven useful in deriving such bounds when good model priors are known and test cases match training cases in distribution.

However, in real world tasks, where deep neural networks are the models of choice and training and test cases come from different domains, deriving tight and estimable bounds remains an unresolved challenge. Recent work has shown that using data dependent priors is a promising way to achieve tighter bounds for deep neural networks in stationary domains. In this work, we combine recent advances in PAC-Bayes domain adaptation with data-dependent priors to give estimable and informative bounds for problems where classical bounds are vacuous. We apply this method to a domain adaptation image classification task and find that it produces tighter bounds. We study which terms dominate the bounds and identify possible directions for further improvement.

Methods

We apply data dependent priors[1] on two bounds from the literature[2] for a domain adaptation image classification task. We seek to understand how the addition of data dependent priors affects the sample generalization part of the bound. Further, it is of interest to find if any specific part of the bounds dominates and in which range it does so. Moreover, we also want to investigate if the dominating terms change as the training of the model progresses. I.e., we evaluate the bound at several different points during the training of the model, as the KL term is expected to increase increase as the posterior drifts away from the prior.

Theorem 1 (Additive bound). *For any real numbers $\omega, \alpha > 0$ we have with probability at least $1-\delta$ over the random choice of $S \times T_x \sim (\mathcal{S} \times \mathcal{T}_{\mathbf{X}})^m$; for every posterior ρ on $\mathcal{H}$*

$$\mathop{\mathbb{E}}_{h\sim\rho} R_{\mathcal{T}}(h) \leq \mathop{\mathbb{E}}_{h\sim\rho} \omega' \hat{R}_{\mathcal{S}}(h) + \alpha' \frac{1}{2} \hat{Dis}_{\rho}(S, T_x) + \Big(\frac{\omega'}{\omega} + \frac{\alpha'}{\alpha}\Big) \frac{KL(\rho\|\pi) + \log\frac{3}{\delta}}{m} + \lambda_{\rho} + \frac{1}{2}(\alpha' - 1),$$

where $\hat{Dis}_{\rho}(S, T_x) = |\hat{d}_{T_x} - \hat{d}_{\mathcal{S}_x}|$ is the empirical domain disagreement, $\lambda_{\rho} = |e_{\mathcal{T}}(\rho) - e_{\mathcal{S}}(\rho)|$ and $\omega' = \frac{\omega}{1-e^{-\omega}}$ and $\alpha' = \frac{2\alpha}{1-e^{-2\alpha}}$.

Theorem 2 (Multiplicative bound). *For any real numbers $a, b > 0$ we have with probability at least $1-\delta$ over the choices $S \sim (\mathcal{S})^m$ and $T_x \sim (\mathcal{T}_x)^n$*

$$\mathop{\mathbb{E}}_{h\sim\rho} R_{\mathcal{T}}(h) \leq a' \frac{1}{2} \hat{d}_{T_x} + b' \beta_{\infty}(\mathcal{T}\|\mathcal{S}) \hat{e}_{\mathcal{S}} + \eta_{\mathcal{T}\setminus\mathcal{S}} + \Big(\frac{a'}{na} + \frac{b'\beta_{\infty}(\mathcal{T}\|\mathcal{S})}{mb}\Big)\Big(2KL(\rho\|\pi) + \ln\frac{2}{\delta}\Big)$$

where $a' = \frac{a}{1-e^{-a}}$, $b' = \frac{b}{1-e^{-b}}$,

$$\beta_{\infty}(\mathcal{T}\|\mathcal{S}) = \sup_{(x,y)\sim supp(\mathcal{S})} \frac{\mathcal{T}(x,y)}{\mathcal{S}(x,y)}$$

and

$$\eta_{\mathcal{T}\setminus\mathcal{S}} = \Pr_{(x,y)\sim\mathcal{T}}\Big((x,y) \notin supp(\mathcal{S})\Big) \sup_{h\in\mathcal{H}} R_{\mathcal{T}\setminus\mathcal{S}}(h).$$

Selected Results

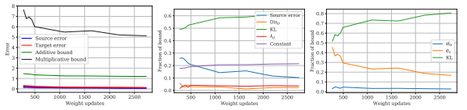

The two bounds are evaluated on the learning task described earlier. In the center and rightmost figures, corresponding to the additive and multiplicative bound respectively, we show the contribution of different terms in the bounds. The labels refer to the term in the bound including any multiplicative constants.

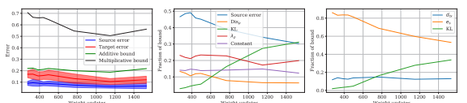

The same type of figure as the one above, however, here we have used 30% of the source data to inform the prior. We see that the bounds are no longer vacuous and that when the KL divergence is small the unobservable λ_{ρ} term is a significant part of the additive bound. The shaded area around source and target error represents one standard deviation.

References

[1] On the role of data in PAC-Bayes bounds
Dziugaite, G. K.; Hsu, K.; Gharbieh, W.; Arpino, G.; andRoy, D. M.
In Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS)

[2] PAC-Bayes and Domain Adaptation
Germain, P.; Habrard, A.; Laviolette, F.; and Morvant, E
Neurocomputing, 2020

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Bågmark, Kasper
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Energy-based approach for the nonlinear filtering problem using a deep splitting method

In this work the main goal is to approximate the optimal nonlinear filter of an underlying high dimensional process through deep learning. This work utilise the deep splitting method, developed for the approximation of solutions to (stochastic) partial differential equations. We solve the Zakai equation, which in turn solves the filtering problem, with an energy-based model. Taking the observations as input, a computationally fast filter is obtained. The model is employed on a nonlinear bistable problem and shows promising performance. The bootstrap particle filter is used for comparison.

Bågmark, Kasper
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Energy-based approach for the nonlinear filtering problem using a deep splitting method

Kasper Bågmark, PhD student
Department of Mathematical Sciences
Supervisors: Adam Andersson (Chalmers and SAAB), Stig Larsson (Chalmers)

CHALMERS UNIVERSITY OF TECHNOLOGY
UNIVERSITY OF GOTHENBURG

1. The optimal filtering problem

Consider a system of stochastic differential equations (SDE) (X,Y) given by

$$X_t = X_0 + \int_0^t \mu(X_s)\,\mathrm{d}s + \int_0^t \sigma(X_s)\,\mathrm{d}W_s, \tag{1}$$

$$Y_t = \int_0^t h(X_s)\,\mathrm{d}s + V_t, \tag{2}$$

where X is called the **underlying (unobserved) state process** in $L^2(\Omega;\mathbb{R}^d)$ and Y is the **observation process** in $L^2(\Omega;\mathbb{R}^d)$. W and V are two independent $\mathbb{R}^d$-valued Brownian motions. The optimal filtering problem consists of finding the probability density of the state given the observation, $p\left(X_t|\,(Y_s)_{0\le s\le t}\right)$. This is called the **filtering density**.

2. The Zakai equation

The unnormalized version of the filtering density $p_t := p\left(X_t|\,(Y_s)_{0\le s\le t}\right)$ can be shown to satisfy the stochastic partial differential equation (SPDE) known as the **Zakai equation**. The strong form of the Zakai equation reads

$$p_t = p_0 + \int_0^t \mathcal{A}p\,\mathrm{d}s + \int_0^t p_s h^\top\,\mathrm{d}Y_s, \tag{3}$$

where $\mathcal{A}$ is the second order operator from the Kolmogorov forward equation related to X, and Y is the observed process. By substitution, the second integral contains an Ito integral.

3. Methods

Deep splitting method: In [1] a splitting method for SPDE, including (3), is introduced. The splitting method is based on solving the linear part of the equation analytically via a Feynman–Kac formula, and adding the nonlinearity in a second step. The scheme is formulated as a recursive (in time) nonlinear least squares problem. The recursion reads

$$p_t = \arg\min_{u\in C(\mathbb{R}^d,\mathbb{R})} \mathbb{E}\left|u(X_{T-t_{n+1}}) - (p_{t_n}(X_{T-t_n}) + f(X_{T-t_n}, p_{t_n}(X_{T-t_n}), (\nabla p_{t_n})(Y_{T-t_n}), \Delta t, \Delta Y))\right|^2. \tag{4}$$

Here f is short-notation for the Euler–Maruyama or Milstein schemes for (3). In [1] u is approximated with a deep neural network for every realization of Y. We consider a more general framework where we let the model take the observation sequence as input.

Energy-based approach: In probabilistic model learning, one successful technique in density estimation and maximum likelihood estimation is the use of energy-based methods (EBM) [2]. The idea is to approximate $p(x|y)$ by associating an scalar energy f^θ to each pair of (x,y) where in our setting $x := X_{t_n}$ and $y := Y_{t_1:t_n}$. The model is trained to associate high energies to pairs that are unlikely and low energy to values that are likely. In our setting we use the unnormalized parametric model

$$\widehat{p}_t(x|y) := e^{-f^\theta(x,y)}, \tag{5}$$

where θ denotes the parameters of our energy-based model.

References

[1] T. Beck, C., Becker, S., Cheridito, P., Jentzen, A., & Neufeld, A. Deep learning based numerical approximation algorithms for stochastic partial differential equations and high-dimensional nonlinear filtering problems. arXiv preprint arXiv:2012.01194.

[2] Gustafsson, F. K., Danelljan, M., Bhat, G., & Schön, T. B. Energy-based models for deep probabilistic regression. In European Conference on Computer Vision (pp. 325-343). Springer, Cham.

4. Numerical Results

Our proposed method is to combine the energy-based approach with the deep splitting method. We demonstrate the method on a nonlinear bistable problem.

Example. *Consider a process (X,Y) satisfying (1) and (2) with nonlinear drift $\mu(x) = 5x - x^3$, constant diffusion $\sigma(x) = 1$ and linear observation $h(x) = x$ with initial density $p_0 = \mathcal{N}(0,1)$. Below we see the underlying density of the state at time $T = 0.5$.*

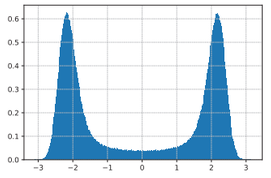

We compare the result from our approximation to the bootstrap particle filter by measuring the distance from the true state X_t to the mean of our method and the bootstrap particle filter (PF), respectively. Formally this is the $L^1(\Omega;\mathbb{R}^d)$-norm

$$\mathbb{E}\|X_{t_n} - \mathbb{E}\left[X_{t_n}|Y_{t_1:t_n}\right]\| \quad \text{and} \quad \mathbb{E}\|X_{t_n} - \widehat{\mu}_{t_n}\|$$

for $n = 1,\dots,25$, where $\mathbb{E}\left[X_t|Y_{t_1:t_n}\right]$ is approximated by the particle filter and $\widehat{\mu}_{t_n}$ is the estimated mean from our method.

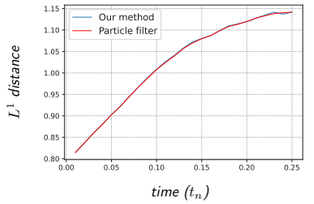

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Geometric Deep Learning and Equivariant Neural Networks

When constructing a convolutional network for image analysis one cannot truly escape the risk that real world data will not respect the the orientation of data the model was trained on. For example, satellite images have, by their nature, no preferred orientation. How can one deal with this problem in an easy way? One solution is to use explicitly equivariant convolutions. This poster discusses some points for why the equivariant convolutions are needed and discusses an implementation made by Cohen et al. in 2016 as well as presents a visualisation of its effect on a network. It also discusses some parts of the mathematical structure as well as current and future work.

AI MATH | Page 7 B
Carlsson, Oscar
Chalmers
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Use the **symmetries** in your problems to your advantage.

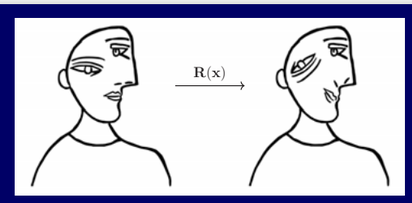

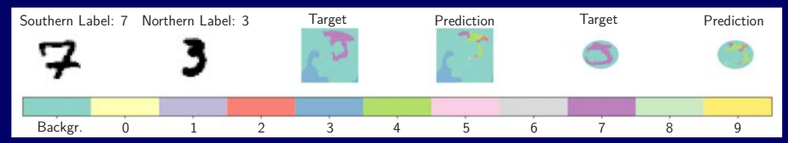

Geometric Deep Learning and Equivariant Neural Networks

Oscar Carlsson[†⋆] Daniel Persson[†] Jimmy Aronsson[†⋆] Fredrik Ohlsson[•]
Jan Gerken[†] Christoffer Petersson[◇] Hampus Linander[◇]

(**Affiliations** †: Chalmers Univerity of Technology, Department of Mathematical Sciences. ⋆: WASP. •: Umeå University, Department of Mathematics and Mathematical Statistics. ◇: Zenseact.)

Intro

How does one deal with rotations? Options are:

- You don't, you assume all your data will have your prefered orientation. (Dangerous: real life throws curveballs at your models)
- Augement training data so that everything is represented. (Every orientation becomes a **lot** of data to deal with)
- Make sure that all your data has the right orientation. (A lot of work, either manually or making an algorithm to unrotate data)
- Modify your architecture and layers to deal with the rotation automatically. (The easy way.)

(Bullet image source: "rotation" by Adrien Coquet from the Noun Project)

One way: transform kernels [Cohen and Welling 2016]

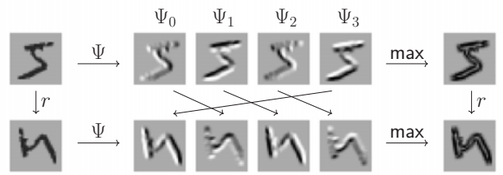

Example: Magnitude of classification invariance for four fold rotation symmetry applied to single MNIST digit

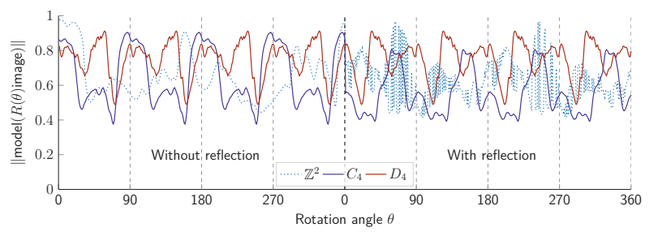

Some mathematics

A map Φ is equivariant with respect to a transformation T of the data if it doesn't matter if one transforms the data before or after one applies the map Φ:

$$\Phi \circ T = T \circ \Phi. \tag{1}$$

An example is that normal convolutions are equivariant to translation. One can extend this to a larger equivariance if we allow convolution kernels and data to be functions on a group:

$$[\Psi \star f](g) = \int_G \Psi(g^{-1}g')f(g')\,dg'. \tag{2}$$

This is equivariant if the kernel transforms in a special way under the group action:

$$\Psi(hgh') = \rho_2(h)\Psi(g)\rho_1(h'). \tag{3}$$

[Cohen and Welling 2016] discretise this to allow for easy computation, see figure on the left.

This can be generalised to local transformations by taking a viewpoint of local coordinates changes. Equivariant convolutions in this general context was introduced by [Cheng et al. 2019] without much mathematical details. We expand on the details of this and provide some generalisation in our recent article [Gerken et al. 2021].

Current work

We're currently examining the details of how imposed equivariance affects semantic segmentation and how it compares to data augmentation.

References

Cheng, Miranda C. N. et al. (June 6, 2019). *Covariance in Physics and Convolutional Neural Networks*. arXiv: 1906.02481 [hep-th, stat]. URL: http://arxiv.org/abs/1906.02481 (visited on 01/27/2020).

Cohen, Taco S. and Max Welling (June 3, 2016). *Group Equivariant Convolutional Networks*. arXiv: 1602.07576 [cs, stat]. URL: http://arxiv.org/abs/1602.07576 (visited on 11/07/2019).

Gerken, Jan E. et al. (May 28, 2021). *Geometric Deep Learning and Equivariant Neural Networks*. arXiv: 2105.13926 [hep-th]. URL: http://arxiv.org/abs/2105.13926 (visited on 05/31/2021).

← Download [Gerken, Jan E. et al. 2021]

WASP
GÖTEBORGS UNIVERSITET

Dadras, Ali
Umeå university

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Solving stochastic/deterministic constrained optimization problems in statistical learning.

Modern learning machines, such as deep neural networks, are often over-parametrized and tuned to perfectly interpolate the training data. Recent works have shown that first-order methods could converge fast in non-convex optimization problems such as overparameterized neural networks, satisfying certain interpolation conditions (e.g., zero training loss). We seek to investigate and understand the convergence of first-order methods in non-convex optimization problems with deterministic or stochastic constraints.

Dadras, Ali
Umeå university

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Solving stochastic/deterministic constrained optimization problems in statistical learning

UMEÅ UNIVERSITY

Ali Dadras, Umeå University
Department of Mathematics and Mathematical Statistics

Abstract

Modern learning machines, such as deep neural networks, are often over-parametrized and tuned to perfectly interpolate the training data. Recent works have shown that first-order methods could converge fast in non-convex optimization problems such as overparameterized neural networks, satisfying certain interpolation conditions (e.g., zero training loss). We seek to investigate and understand the convergence of first-order methods in non-convex optimization problems with deterministic or stochastic constraints.

Problem Statement

Let $\{x_i\}_{i=1}^n$ be a given training set, and let $\{y_i\}_{i=1}^n$ be training labels. We would like to minimize

$$\min_{\theta \in D} f(\theta) = \sum_{i=1}^{n} f_i(\theta;\, x_i, y_i)$$

where θ is the model parameter and D is a set of stochastic or deterministic constraints .

Objectives

- Investigating the potential of first-order methods in optimizing non-convex optimization problems with deterministic or stochastic constraints.
- Investigating the existence of first order methods for solving constrained optimization problems motivated by learning problems.
- Investigating and understanding the convergence of desired first order methods.

Methods

Considering deterministic constraints, a vast number of studies have been done to solve the above optimization problem. There are different solving strategies, many of them rely on gradient descent and its variants. To improve these gradient-based methods, different strategies are proposed.

- Preconditioning (e.g.,data normalization, layer and batch normalization)
- Mmomentum (e.g., Polyak and Nestrov)
- Variance reduction (e.g., SAG, SDCA, SVRG)
- Adaptive stepsizes (e.g., Adagrad, ADAM)
- Importance sampling

References

1. Meng, Si Yi, et al. "Fast and furious convergence: Stochastic second order methods under interpolation." International Conference on Artificial Intelligence and Statistics. PMLR, 2020.
2. Loizou, Nicolas, et al. "Stochastic polyak step-size for SGD: An adaptive learning rate for fast convergence." arXiv preprint arXiv:2002.10542 (2020).
3. Vaswani, Sharan, et al. "Adaptive Gradient Methods Converge Faster with Over-Parameterization (and you can do a line-search)." arXiv preprint arXiv:2006.06835 (2020).

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Gorenstein discrete decomposable models

Discrete hierarchical models are statistical models that are widely used throughout statistics and data science. An advantage of these models is that there are established methods that can be used to make inference.

The goal of this project is to explore at a deeper level the combinatorial objects arising from discrete decomposable models beyond their graph. Specifically, we aim to answer when enumerative properties, such as the Gorenstein property, hold for the polytope associated to a discrete decomposable model.

AI MATH | Page 9 B
Deligeorgaki, Danai
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Gorenstein discrete decomposable models

Danai Deligeorgaki, Department of Mathematics (KTH)
Supervisor: Liam Solus (KTH)

Introduction

We study discrete decomposable models, a family of statistical models that lie in the class of **hierarchical models**. Decomposable models and their corresponding graphs are of wide use throughout **statistics** and data science. For instance, directed acyclic graphs (**DAGs**) can be approximated by decomposable graphs. The complexity of this approximation determines the complexity of probabilistic inference algorithms for DAG models such as **variable elimination**. Therefore, the combinatorics of the graphs defining the decomposable models carry important information in regard to **probabilistic inference**. The goal of this project is to explore at a deeper level the information encoded in combinatorial objects associated to **decomposable models**.

Definition
A **decomposable simplicial complex** Γ is a collection of simplices, i.e. nodes, edges, triangles, tetrahedra, etc., that are glued together (in a certain way). The simplices in Γ are called **faces** and the (non-trivial) inclusion-maximal faces are called **facets**.
For example, the graph on the right denotes a decomposable simplicial complex on 9 nodes. The edge {1,2}, the triangle {2,3,4} and the node {7} are some of its facets.

decomposable simplicial complex

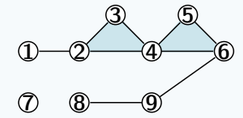

Discrete decomposable models

Let $r_1, ..., r_m \in \mathbb{N}$ be the number outcomes of the discrete variables $X_1, X_2, ..., X_m$, respectively, and let $\mathcal{R} = r_1 \times \cdots \times r_m$ be the set of all possible outcomes. The joint distribution of $X_1, ..., X_m$ lies in the $(\#\mathcal{R} - 1)$-dimensional **probability simplex**

$$\Delta_{\#\mathcal{R}-1} = \{p \in \mathbb{R}^{\#\mathcal{R}} : p_i \geq 0, \text{ for all } i \in \mathcal{R} \text{ and } \sum_{i \in \mathcal{R}} p_i = 1\}.$$

The **decomposable model** associated with a decomposable simplicial complex Γ is

$$M_\Gamma = \{p \in \Delta_{\#\mathcal{R}-1} : p_i = \frac{1}{Z(\theta)} \prod_{F \in \text{facet}(\Gamma)} \theta_{i_F}^{(F)} \text{ for all } i \in \mathcal{R}\},$$

for $\theta_{i_F}^{(F)}$ positive parameters and $Z(\theta)$ normalizing constant.

From the model to the polytope

Apart from the graph Γ, there are other combinatorial objects linked to a decomposable model M_Γ. In fact, M_Γ can be written as the intersection of a toric variety V_{M_Γ} with the probability simplex $\Delta_{\#\mathcal{R}-1}$.

For example, for $\#\mathcal{R} = 3$,

$$M_\Gamma = V_{M_\Gamma} \cap \Delta_2$$

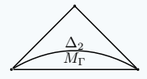

From the toric variety, which is an algebro-geometric object, we can pass to a polytope P_{M_Γ}, a geometric object. It is a property of toric varieties that the geometric properties of V_{M_Γ} are encoded in the **polytope** P_{M_Γ}.

In this project, we are investigating the structure of this polytope to see if it carries useful information in relation to probabilistic inference.

References

[1] Markov bases of binary graph models
M. Develin, S. Sullivant
Annals of Combinatorics 7, 2003

[2] Gröbner bases and polyhedral geometry of reducible and cyclic models
S. Hosten, S. Sullivant
Journal of Combinatorial Theory, Series A 100.2, 2002

Getting to know the polytope

When investigating a polytope's combinatorics, there are several questions to be explored, such as

- What are the facets of the polytope P_{M_Γ}? Answered in [1].

  The facets are given by $x_{i_F}^F \geq 0$ for $F \in$ facets(Γ) and $i_F \in \mathcal{R}_F$.

- Does P_{M_Γ} admit a regular unimodular triangulation? Answered in [2].

  Yes!

What combinatorial information does this triangulation carry? Open.

- What are the enumerative properties of P_{M_Γ}? Our results.

To this end, we study the structure of an integer polynomial associated to P_{M_Γ}, called the ***h****-polynomial**,

$$h^*(x) = h_0^* + h_1^* x + \cdots + h_{\#\mathcal{R}-1}^* x^{\#\mathcal{R}-1}.$$

This polynomial captures important information about the polytope, including its volume and whether or not the polytope P_{M_Γ}, and hence the model M_Γ, has the Gorenstein property. In fact, the decomposable model M_Γ is **Gorenstein** if and only if $h^*(x)$ is palindromic.
We characterize all Gorenstein discrete decomposable binary models for forests Γ.

Theorem 1
Let Γ be a forest on m nodes and $X_1, ..., X_m$ be binary variables.
Then M_Γ is Gorenstein if and only if all connected components of Γ have

- exactly one vertex,
- strictly more than one vertex.

Gorenstein examples

Γ_1

Γ_2

Future work

We will continue exploring the combinatorial properties of discrete decomposable models, and their interpretation in terms of statistics. Our current goals are to

1. Interpret the observations in Theorem 1 statistically.
2. Generalize Theorem 1 to characterize all discrete decomposable models. We already have a conjecture in this direction.
3. Analyze the information that the triangulation constructed in [2] carries.

Deducing function from structure

Neuroscientists are working hard to map and understand the intricate connections of neurons in a brain. What will the knowledge of that structure give us? Using which mathematical framework can we in a useful yet practical way describe the (supposed) link between the structure of a network and the tasks it can perform? To find a rigorous answer, we study the impact that different structures and dynamics can have on networks. The aim is to combine pure mathematics and neuroscience, using methods from statistical physics, combinatorics, geometry, percolation and probability theory.

Even endowing a simple structure with simple dynamics can yield surprisingly intricate results. We now study emerging structures in the Hopfield model as well as cellular automata containing inhibitory and excitatory 'neurons'. The latter can be thought of as a generalisation of bootstrap percolation with highly non-monotone behaviour!

Polyhedral geometry of Wasserstein distances

Every discrete probability distribution corresponds to a point in the standard simplex. Given a model consisting of probability distributions and sample data, we want to find a candidate in the model that best explains the data. Studying the Wasserstein distance between probability distributions is one route to this. The approach we take is to use polyhedral geometry - in particular bisectors and bisection fans - to better understand the Wasserstein distance.

AI MATH | Page 11 B
Jal, Aryaman
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Polyhedral geometry of Wasserstein distances

Aryaman Jal Katharina Jochemko
Department of Mathematics, Royal Institute of Technology (KTH)

KTH VETENSKAP OCH KONST

Motivation

Every probability distribution on $[n] = \{1, 2, \ldots, n\}$ corresponds to a point $\mu \in \Delta_{n-1} = \{x \in \mathbb{R}^n : \sum_{i=1}^{n} x_i = 1, x_i \geq 0 \ \forall i = 1, \ldots, n\}$.

Given a finite model $\mathcal{M} = \{\mu_1, \ldots, \mu_k\} \subseteq \Delta_{n-1}$ of discrete probability distributions and samples $x_1, \ldots, x_N \in [n]$, we want to find $\mu_i \in \mathcal{M}$ that best explains the samples. We follow the approach in [1] and use polyhedral geometry to minimise the Wasserstein distance between the empirical distribution and the model.

Wasserstein distance

Given a metric d on $[n]$, the Wasserstein distance is defined by

$$W_d(\mu, \nu) = \operatorname{dist}(\mu, \nu) = \min\{\alpha \in \mathbb{R}_{\geq 0} : \nu - \mu \in \alpha B_d\}$$

where

$$B_d = \operatorname{conv}\left\{\frac{1}{d_{ij}}(e_i - e_j) : 1 \leq i \neq j \leq n\right\}$$

is called the Wasserstein unit ball with respect to W_d. Given $x_1, \ldots, x_N$, the Wasserstein distance estimator equals

$$\min_{i=1,\ldots,k} W_d\left(\frac{1}{N}\sum_{j=1}^{N} \delta_{x_j}, \mu_i\right)$$

Bisectors

If $N = 2$, the decision boundary is given by the bisector

$$\operatorname{bis}(\mu, \nu) = \{x \in \mathbb{R}^n : \operatorname{dist}(\mu, x) = \operatorname{dist}(\nu, x)\}.$$

Proposition. [2] The bisector $\operatorname{bis}(\mu, \nu)$ is a polyhedral complex.

Remark.

1. Polyhedrality of the bisectors is true for all norms.
2. The bisector can have non-empty interior (see Figure 3). This happens if and only if μ and ν lie on a hyperplane parallel to a facet of B_d.

Euclidean and Wasserstein bisectors

Figure 1 shows the bisector of two points with respect to the Euclidean distance. In Figures 2 and 3, the bisector for W_d for $n = 3$ (projected down to $\mathbf{1}^{\perp}$) and $d_{ij} = 1$, for all $i \neq j$, are depicted.

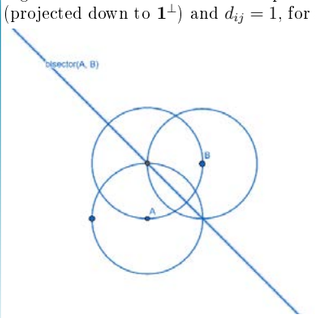

Figure 1: Bisector w.r.t. Euclidean distance

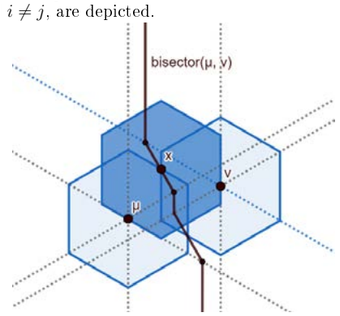

Figure 2: One-dimensional bisector w.r.t. W_d

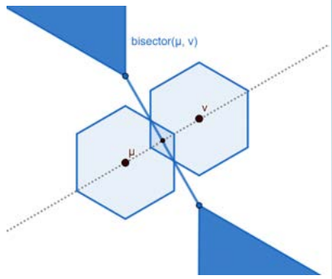

Figure 3: Two-dimensional bisector w.r.t. W_d

Goals

1. Investigating the combinatorics of $\operatorname{bis}(\mu, \nu)$ when moving ν. This leads to the concept of a bisection fan (see [2]).
2. Determining the number of maximal cells in $\operatorname{bis}(\mu, \nu)$ as a measure of complexity of the decision boundary.
3. Determining Voronoi diagrams with respect to Wasserstein distances.

Partial Results

For every graph $G = ([n], E)$, define $d_{ij} = 1$ if $ij \in E$ and ∞ otherwise. In this case, B_d is a symmetric edge polytope ([3]). If G is a tree, then W_d equals the 1-norm up to an affine transformation and B_d is affinely isomorphic to the cross-polytope

$$\diamondsuit_{n-1} = \operatorname{conv}\{\pm e_i : i \in [n-1]\}.$$

Proposition (Jal, Jochemko 2021+)**.** The bisection fan of W_d corresponding to the case of d being a graphical metric on a tree is, up to affine transformation, induced by the hyperplane arrangement

$$\mathcal{H} = \bigcup_{i=1}^{n-1}\{x_i = 0\} \bigcup_{I \subseteq [n-1]} \left\{\sum_{i \in I} x_i = \sum_{i \in I^c} x_i\right\}$$

A similar, more involved result was obtained in the case of G equal to the complete graph K_n.

References

[1] Çelik, T.Ö., Jamneshan, A., Montúfar, G., Sturmfels, B. and Venturello, L., 2021. Wasserstein distance to independence models. Journal of Symbolic Computation, 104, pp.855-873.

[2] Criado, F., Joswig, M. and Santos, F., 2019. Tropical bisectors and Voronoi diagrams. arXiv preprint arXiv:1906.10950.

[3] Higashitani, A., Jochemko, K. and Michaek, M., 2019. Arithmetic aspects of symmetric edge polytopes. Mathematika, 65(3), pp.763-784.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Jansson, Erik
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

ResNets Understood as Sub-Riemannian Landmark Matching

Residual neural networks can be interpreted as time discretizations of optimal control problems. This observation means that it is possible to use sub-Riemannian landmark matching, a method from the field of shape analysis, to study and understand ResNets. For instance, as demonstrated in the poster, the impact of regularization on the smoothness of transformations can be studied from a diffeomorphic point of view. The connection between the ResNets and sub-Riemannian landmark matching demonstrates that it is possible to study and understand neural networks using shape analysis methods.

AI MATH | Page 12 B

Jansson, Erik
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

ResNets understood as Sub-Riemannian Landmark matching

Erik Jansson, Chalmers and University of Gothenburg
Supervisors: Klas Modin (Chalmers and GU)

UNIVERSITY OF GOTHENBURG
CHALMERS UNIVERSITY OF TECHNOLOGY

sub-Riemannian landmark matching

1. Problem: find a diffeomorphism, determined by a vector field, warping *initial landmarks* $x_1, ..., x_n$ on a *manifold* M to targets $c_1, ..., c_n$ on a metric space N.
2. ***Idea: Consider only vector fields parametrized by variables in some Hilbert space*** $\mathcal{U}$.

This is an optimal control problem.

$$\min_{u:[0,1]\to\mathcal{U}} \sum_{i=1}^{m} d_N^2(h(y_i(1), c_i)) + \int_0^1 \ell(F(u))\mathrm{d}t,$$

$$\text{s.t.} \quad \dot{y}_i = F(u)(y_i),\ t \in [0,1],\ y_i(0) = x_i.$$

Note that F maps from controls to vector fields and formally parametrizes the vector fields!

We derive an *equaton of motion* describing the ***evolution of the controls.***

$$A(u)\dot{u} = (D_u^T m \circ L^{-1}) \cdot \left(\mathrm{ad}_v^T m + \mathrm{div}(v)m\right)$$

$$Lv = (1 - \alpha\Delta)^k v = m$$

Connection with residual neural networks

Residual neural networks can be considered as *time discretizations* of time-continuous control problems.

In some settings: ResNets discretize a high dimensional sub-Riemannian landmark matching problem!

Interpretation 1: If we know the dynamics of the control, increasing the number of layers does not increase the number of parameters we need to optimize.
Interpretation 2: Shape analysis, a mathematically well understood subject, can be used to interpret neural networks, *and vice-versa*.

Numerical example

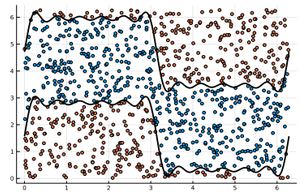

We let $\mathcal{U} = \mathbb{R}^{16}$ and set $F(u) = X^0 + \sum_{i=1}^m u_i X^i$ for vector fields on $\mathbb{R}^2/(2\pi\mathbb{Z})^2$ that diagonalizes the Laplace-Beltrami operator. We generate 1000 points and classify those in a variable band of width π as a 1 and those outside as a 0. We want to determine an initial value $u(0)$ so that the points are moved towards a uniform band.

We first try without regularization. Note that the points are moved very erratically. This is not computationally stable.

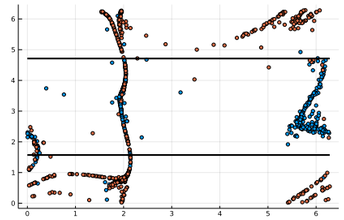

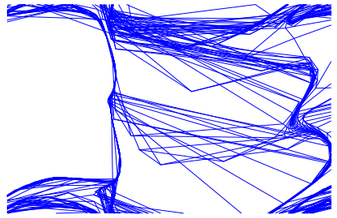

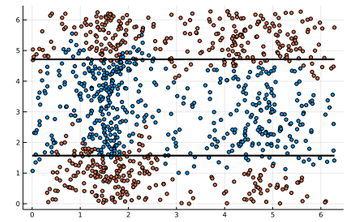

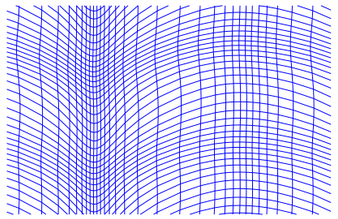

We increase the regularity sharply. Note that the resulting warp is much smoother. This transformation is more robust.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Accelerated Deterministic Methods in Optimization

Optimization problems play an important role in the process of learning a machine using previously available data. This process, can be time consuming and therefore many researchers have tried to reduce it through various techniques. One method to attack this problem is to reduce the optimization time of the learning process. As a result, accelerated methods in optimization gain a remarkable attention. Among the first order algorithms for smooth convex functions, Nesterov's accelerated gradient descent(NAG) is proven to be the fastest. For decades, various studies tried to enlighten the essence of acceleration through Nesterov updates. Recently, using Ordinary Differential Equations(ODE), it is shown that for a fixed convergence rate, accelerated algorithms may not be unique. This research, proposes a general algorithm which can achieve various convergence rates for different choices of parameters.

AI MATH | Page 13 B
Maskan, Hoomaan
Umeå University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Accelerated Deterministic Methods in Optimization

Hoomaan Maskan, Umeå University
Department of Mathematics and Mathematical Statistics
Main Advisor: Armin Eftekhari
UMEÅ UNIVERSITY

Abstract

Optimization problems play an important role in the process of learning a machine using previously available data. This process, can be time consuming and therefore many researchers have tried to reduce it through various techniques. One method to attack this problem is to reduce the optimization time of the learning process. As a result, accelerated methods in optimization gain a remarkable attention. Among the first order algorithms for **smooth convex functions, Nesterov's accelerated gradient descent(NAG)** is proven to be the fastest. For decades, various studies tried to enlighten the essence of acceleration through Nesterov updates. Recently, using **Ordinary Differential Equations(ODE)**, it is shown that for a fixed convergence rate, accelerated algorithms may not be unique. This research, proposes a general algorithm which can achieve various convergence rates for different choices of parameters.

Motivation and Methods

The learning problem can be formulized as

$$\min_{\theta\in\mathbb{R}^{n\times n}} \frac{1}{N}\sum_i \mathcal{L}(\boldsymbol{y}_i, f(\boldsymbol{x}_i,\boldsymbol{\theta}))$$

which is known as **Empirical Risk Minimization** (ERM) problem. Depending on the features of f and $\mathcal{L}$, this problem can be non-linear and non-convex. Therefore, if not impossible, it would be so hard to find the global minimizer(s) of this problem. For simplicity, from now on we consider the objective function to be smooth and μ-strongly convex and denote it as $F_L(\boldsymbol{x},\boldsymbol{y},\boldsymbol{\theta})$.

NAG updates for $\min_{\theta\in\mathbb{R}^{n\times n}} F_L(\boldsymbol{x},\boldsymbol{y},\boldsymbol{\theta})$ are[]

$$\begin{cases}\boldsymbol{v}_{k+1} = \gamma\boldsymbol{v}_k + h\nabla F_L(\boldsymbol{x},\boldsymbol{y},\boldsymbol{\theta_k}-\gamma\boldsymbol{v}_k)\\ \boldsymbol{\theta}_{k+1} = \boldsymbol{\theta_k} - \boldsymbol{v}_{k+1}\end{cases},$$

which gets the best of the momentum term $\boldsymbol{v}_k$ and calculates the gradient near the future point $\boldsymbol{\theta}_{k+1}$(see Figure).

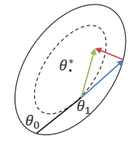

Shi, et. al. proposed high-resolution ODEs for modeling acceleration methods [1]. Specifically, if $\boldsymbol{\vartheta}$ denotes the continuous trajectory of the NAG, then

$$\begin{cases}\ddot{\boldsymbol{\vartheta}} = -\sqrt{h}\nabla F_L(\boldsymbol{\vartheta}) - \sqrt{\mu}(\dot{\boldsymbol{\vartheta}} - V)\\ \dot{V} = -\sqrt{\mu}(V-\boldsymbol{\vartheta}) - \left(\frac{1}{\sqrt{\mu}}\right)\nabla F_L(\boldsymbol{\vartheta})\end{cases} \quad (1)$$

with $\boldsymbol{\vartheta}(0) = \boldsymbol{\vartheta_0}, \dot{\boldsymbol{\vartheta}}(0) = \frac{-2\sqrt{h}\nabla F_L(\boldsymbol{\vartheta}_0)}{1+\sqrt{\mu h}}$ will converge to the global minimizer with rate

$$F_L(\boldsymbol{\vartheta}) - F_L(\boldsymbol{\vartheta}^*) \le \frac{2||\boldsymbol{\vartheta_0}-\boldsymbol{\vartheta}^*||^2}{h} e^{\frac{-\sqrt{\mu}t}{4}}.$$

Interestingly, if one discretizes the above ODE with semi-implicit Euler scheme, then with small step size NAG is approximately a symplectic method. Also, implicit Euler scheme leads to acceleration, but it is not easy to use in practice [2].

References

1. Shi, B., Du, S.S., Jordan, M.I. et al. Understanding the acceleration phenomenon via high-resolution differential equations. Math. Program. (2021). https://doi.org/10.1007/s10107-021-01681-8
2. Shi, Bin. "Acceleration via Symplectic Discretization of High-Resolution Differential Equations." (2019).
3. Zhang, Peiyuan, et al. "Revisiting the Role of Euler Numerical Integration on Acceleration and Stability in Convex Optimization." International Conference on Artificial Intelligence and Statistics. PMLR, 2021.

Main Theorem

The ODE (1) can be generalized by replacing the coefficients with positive parameters m, n, p, q

$$\begin{cases}\ddot{\boldsymbol{\vartheta}} = -m\nabla F_{\mathcal{L}}(\boldsymbol{\vartheta}) - n(\dot{\boldsymbol{\vartheta}} - V)\\ \dot{V} = -q(V-\boldsymbol{\vartheta}) - p\nabla F_{\mathcal{L}}(\boldsymbol{\vartheta})\end{cases},$$

which can be rephrased as the ODE

$$\ddot{\boldsymbol{\vartheta}} + ((n+q) + m\nabla^2 F_{\mathcal{L}}(\boldsymbol{\vartheta}))\dot{\boldsymbol{\vartheta}} + (mq+np)\nabla F_{\mathcal{L}}(\boldsymbol{\vartheta}) = 0 \quad (2)$$

Preliminary Result: The following theorem shows that for a fixed rate of convergence, one can find many accelerated ODEs.

Theorem 1: Assume $F_L(\boldsymbol{\vartheta})$is L-smooth and μ-strongly convex. Then if $\boldsymbol{\vartheta}(t)$ and $V(t)$ are such that (2) holds, the **Lyapunov function**

$$\varepsilon(t) = F_L(\boldsymbol{\vartheta}(t)) - F_L(\boldsymbol{\vartheta}^*) + A\|V(t) - \boldsymbol{\vartheta}^*\|^2$$

will decrease as

$$\varepsilon(t) \le e^{-\min\{n,\frac{q}{4}\}t}\varepsilon(0)$$

with $\max\left\{\frac{m}{q}, \frac{n\mu}{q}\right\} \le A \le \min\left\{\frac{n}{2(q+p)}, \frac{4n\mu}{3q}\right\}$ and $m, n, p, q \ge 0$.

We can apply **semi-implicit Euler integrator** to achieve the corresponding algorithm

$$\begin{cases}\boldsymbol{v}_{k+1} - \boldsymbol{v}_k = -p\sqrt{h}\nabla F_L(\boldsymbol{x},\boldsymbol{y},\boldsymbol{\theta_{k+1}}) - q\sqrt{h}(\boldsymbol{v}_k - \boldsymbol{\theta_k})\\ \boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k = -m\sqrt{h}\nabla F_L(\boldsymbol{x},\boldsymbol{y},\boldsymbol{\theta_k}) - n\sqrt{h}(\boldsymbol{\theta_k} - \boldsymbol{v}_k)\end{cases} \quad (3)$$

The following theorem shows the convergence rate of this algorithm.

Theorem 2: Assume $F_L(\boldsymbol{\vartheta})$is L-smooth and μ-strongly convex. Then if $\boldsymbol{\theta}_k$ and $\boldsymbol{v}_k$ follow the updates (3), the Lyapunov function

$$\varepsilon(k) = F_L(\boldsymbol{\vartheta}_k) - F_L(\boldsymbol{\vartheta}^*) + A'\|\boldsymbol{v}_k - \boldsymbol{\vartheta}^*\|^2$$

will decrease as

$$\varepsilon(k+1) \le (1-\lambda)^k\varepsilon(0)$$

with

$$\frac{n}{2q(1-q\sqrt{h}-p\sqrt{h})} \le A' \le \min\left\{\frac{1}{4L(q\sqrt{h}-q^2h)}, \frac{\mu^2}{4L(q\sqrt{h}-p\sqrt{h\mu^2}-p^2h\mu^2)}\right\},$$

$$p\sqrt{h} \le q\sqrt{h} \le \frac{1}{2}, \frac{1}{2L\sqrt{h}} \le m \le \frac{1}{L\sqrt{h}}, n \le \frac{1-q\sqrt{h}-p\sqrt{h}}{2L(1-q\sqrt{h})\sqrt{h}}, m, n, p, q \ge 0,$$

$$\lambda = \min\{q\sqrt{h} - p\sqrt{h}, \frac{2}{L}(\frac{\mu^2}{4L} + A'p\sqrt{h}\mu^2 + A'p^2h\mu^2 - Aq\sqrt{h})\}.$$

For comparison, the trajectory of Lyapunov function above (3) is compared with the Lyapunov function in [3]. Of course, there are conditions under which the behaviours are different. Deeper analysis is left for future work.

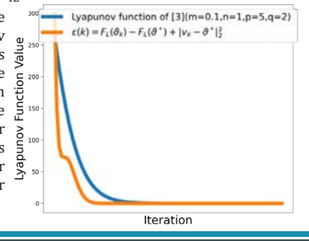

Normative reasoning for Social AI

Norms are both a regular occurrence in human reasoning, as well as a useful tool for governing the behaviour of agent populations. However, what exactly norms are, and how we can effectively use them in computer science is poorly understood. For example, in CS, norms are often purely used as constraints of behaviour, while they can also have a strong motivating component.
In our research, we try to address this issue by formalizing sociological and psychological theories of norms. This gives us a framework for studying norms and their interactions. Besides this, we also study how norms influence human reasoning, and how agents can use them in their reasoning.

AI MATH | Page 14 B
Mellema, René
Umeå University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Normative reasoning for Social AI

René Mellema, Umeå University
UMEÅ UNIVERSITY
Computing Science
Supervisors: Frank Dignum
Juan Carlos Nieves Sanchez

Motivation & Research goals

Norms are both a regular occurrence in human reasoning, as well as a useful tool for governing the behaviour of agent populations. However, what exactly norms are, and how we can effectively use them in computer science is poorly understood. For example, in CS, norms are often purely used as constraints of behaviour, while they can also have a strong motivating component.
In our research, we try to address this issue by formalizing sociological and psychological theories of norms. This gives us a framework for studying norms and their interactions. Besides this, we also study how norms influence human reasoning, and how agents can use them in their reasoning.

Why norms?

In human societies, norms are used to build **accountability** and **cooperation**. This means they play an integral part in all our day-to-day interactions, they can have internal effects, and norm following/breaking behaviour carries meaning as well. This means that they have a strong **motivational** component.
In **social simulation** we want to model human societies to e.g. study them or make predictions about reactions to changes. Since norms play such a pivotal role in human societies, they can have a large effect in these simulations.

Why new formalizations?

However, in CS norms get used in multi-agent systems to control the behaviour of heterogenous populations of agents. This means the focus is often on norms as **constraints**. This means current formalizations ignore the motivational aspects, as well as **sanctioning** behaviours. Both of these are important for **norm change**, which is currently also not well understood.

The formalizations

We are interested in **representing** norms in our simulations such that agents can **reason** with them. This requires that various aspects of the norm are incorporated in the design:
- Activation/deactivation conditions
- Violation condition
- Sanctions for breaking the norm

Similarly, the agents need to be able to react to other agents norm breaking behaviour, which means the representation also needs to take violations into account. Current research is ongoing on how to best represent violations, and which aspects are necessary to differentiate between them. Using this, a framework for normative reasoning in CTL is being developed.

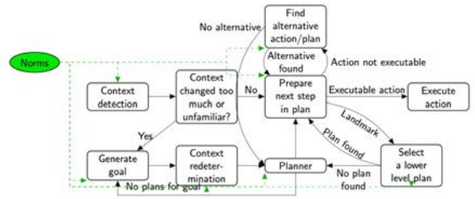

Agents with normative deliberation

If we want the agents to use norms, the norms do not only need to be represented, but the agents also need to know when to take them into account in their reasoning. This needs to happen in:
- **Goal generation**: Norms can have obligations attached, which can motivate an agent to take actions it otherwise would not have
- **Context detection**: Part of a context is what norms are active, which can influence what actions the agent considers
- **Planning**: Not every plan an agent can follow has to follow all norms, but it should make the decision whether or not to follow a plan/take an action based on the norms it might follow/break

In all of these, whether or not the agent should follow a norm is an important consideration.

Norm change

While norms can stabilize a society, they are not static. Since we are interested in peoples possible responses to new policy, how this influences old norms and shapes new ones is vital. Currently, our work here focusses on how agents can break norms, and why they might do so.

References

1. Mellema R., Jensen M., Dignum F. (2021) Social Rules for Agent Systems.
2. Dignum, F. (2021) Social Simulation for a Crisis
3. Brennan G., Eriksson L., Goodin R., Southwood N. (2013) Explaining norms
4. Vázquez-Salceda, J., Aldewereld, H., Grossi, D., & Dignum, F. (2008). From human regulations to regulated software agents' behavior

Further questions

- Where in the decision making process should norms exert an influence?
- How do we combine normative reasoning with other types of motivation?
- When should an agent break a norm?
  - What factors influence this decision?
- How do identities and norms interact?
- When should an agent sanction another agent for breaking a norm?

Interacting Particle Dynamics for Deep Learning

Neural networks (MLPs) and GANs can be interpreted/represented as systems of interacting particles. This may enable using techniques from statistical physics, probability theory and partial differential equations in the understanding of neural networks. Future work includes establishing laws of large deviations (LDP) to help make these connections.
This poster shows two different frameworks in which single hidden layer neural networks, an GANs are treated from this perspective.

AI MATH | Page 15 B
Nilsson, Viktor
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Interacting Particle Dynamics for Deep Learning

Viktor Nilsson, Pierre Nyquist - KTH Mathematics Dept.

Introduction

Several frameworks have been proposed that establish a particle dynamic view of neural networks. In two different fashions, one can see the training and inference of a network as the behavior of a many-particle system, consisting of say N particles. Further, such systems with N particles have a 'mean-field' behavior when letting $N \to \infty$, i.e., having the characteristic of a 'smooth' distribution. This lends itself to so called mean-field approximation, where for large N, the system is approximated by the limit behavior instead. Thus, the discrete probability distribution of the N-particle system is replaced by a continuous distribution instead. This distribution and its evolution under the training dynamics can then be described by a PDE, or a so-called gradient flow.
Several questions remain about the convergence to the mean-field limit.
Current literature establishes some convergence results of law of large numbers (LLN) central limit theorem (CLT) type, while not giving any convergence rates. The current goal is to go beyond these results and use the *theory of large deviations* to develop a *large deviations principle* (LDP), which gives convergence rate guarantees based on a *rate function*.

One hidden layer neural network

Consider a one hidden layer neural network $f_{\boldsymbol{\theta}} : \mathbb{R}^d \to \mathbb{R}$. Its prediction can be seen as an average of the N hidden neurons, i.e.,

$$f_{\boldsymbol{\theta}}(x) = \frac{1}{N}\sum_{i=1}^{N} \varphi(x, \boldsymbol{\theta}_i).$$

How is the behavior when $N \to \infty$? With an L2-loss function it turns out that the loss can be written as

$$l(\boldsymbol{\theta}_1, ..., \boldsymbol{\theta}_N) = \sum_{i=1}^{N} F(\boldsymbol{\theta}_i) + \frac{1}{2N}\sum_{i,j=1}^{N} K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j). \quad (1)$$

Defining the *empirical measure* of the weights, $\mu_t = \sum_{i=1}^{N} \delta_{\boldsymbol{\theta}_{i,t}}$, we have that a standard gradient descent (with infinitessimal timestep) follows the PDE

$$\partial_t \mu_t = \nabla \cdot (\mu_t \nabla V), \quad (2)$$

in the many-particle limit.

GANs

Generative adversarial networks consist of a pair of networks, $G : \mathcal{Z} \to \mathcal{X}$ and $D : \mathcal{X} \to [0, 1]$, that compete in some two-player game, for instance the following zero-sum game.

$$\min_G \max_D \mathbb{E}_{\mathbf{x}}[\log(D(\mathbf{x}))] + \mathbb{E}_{\mathbf{z}}[\log(1 - D(G(\mathbf{z})))] \quad (3)$$

Existence of pure Nash equilibria are not guaranteed in continuous games. However, the existence of *mixed Nash equilibria* is guaranteed. A mixed Nash equilibrium is a Nash equilibrium for the relaxed game

$$\mathcal{L}(\mu_x, \mu_y) := \int\int l(x, y)\, \mu_x(dx)\mu_y(dy). \quad (4)$$

Thus, we consider *mixed strategies* μ_x, μ_y instead of pure strategies x, y. In practice, this is done by having multiple "particles" $\{x_t^i\}_{i=1}^n, \{y_t^i\}_{i=1}^n$ and letting their empirical measures approximate μ_x, μ_y.

$$\mu_{x,t}^n := \frac{1}{n}\sum_{i=1}^{n} \delta_{x_t^i}, \quad \mu_{y,t}^n := \frac{1}{n}\sum_{i=1}^{n} \delta_{y_t^i} \quad (5)$$

How should we optimize $\delta_{x_t^i}$, $\delta_{y_t^i}$? Gradient descent-ascent (DA) dynamics correspond to

$$\begin{aligned} dX_t^i &= -\frac{1}{n}\sum_{j=1}^{n} \nabla_x l(X_t^i, Y_t^j)dt, \\ dY_t^i &= \frac{1}{n}\sum_{j=1}^{n} \nabla_y l(X_t^j, Y_t^i)dt. \end{aligned} \quad (6)$$

Future work

The mean-field behavior is described in [2] and [1]. Currently, we want to strengthen those results by providing a LDP.
The dynamics of equation (2) and equation (6) can be modified by including a diffusion term, e.g. adding the term $\sqrt{2\beta^{-1}}dW_t^i$ to equation (6). We would further like to see how the size of the inverse temperature β affects the convergence.

References

[1] Carles Domingo-Enrich et al. "A mean-field analysis of two-player zero-sum games". In: *arXiv preprint arXiv:2002.06277* (2020).
Grant Rotskoff et al. "Global convergence of neuron birth-death dynamics". In: *arXiv preprint arXiv:1902.01843* (2019).

Osipov, George
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Resolving Inconsistencies in Simple Temporal Problems: A Parameterized Approach

Constraint satisfaction problems (CSPs) have many applications in AI including planning, knowledge representation, and reasoning. Given a set of variables and constraints, the goal in a CSP is to find an assignment that satisfies all constraints. The computational complexity of a CSP depends on the set of allowed constraints. For all sets of constraints over a finite domain, the dichotomy theorem of Bulatov and Zhuk distinguishes between problems that are in P and NP-complete. Almost CSP is an optimization version, where the goal is to find an assignment that violates as few constraints as possible. Applications include handling noise, dealing with faulty measurements, and repairing merge conflicts in databases. With an additional assumption that the number of violated constraints is small, Almost CSP becomes interesting from the point of view of parameterized complexity. Here one needs to distinguish between problems in P, in FPT, W[1]-hard and NP-hard. In our work, we give a full classification for Almost Simple Temporal Problem (STP), an influential reasoning formalism for temporal information.

Osipov, George
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Almost Constraint Satisfaction Problems

George Osipov
Linköping University

Toy Problem

Setting:
You teach a class and n students signed up.
You need to assign them into 2 groups.
Some students are friends and asked to be in the same group.
Some students are foes and asked to be in different groups.

Question 1: Can you divide students and respect all preferences?

Mathematical model:
Introduce a variable x_i for every student $i \in \{1, \dots, n\}$.
Let's set $x_i = 1$ ($x_i = 2$) if student i is assigned to group 1 (2).
If i and j are friends, the assignment should satisfy the constraint $x_i = x_j$.
If i and j are foes, the assignment should satisfy the constraint $x_i \neq x_j$.
Find an assignment $f : \{x_1, \dots, x_n\} \to \{1, 2\}$ that satisfies all constraints.

Follow-up:
What if no such assignment exists?
One can try to find an assignment that violates as few constraints as possible (upsets the students as little as possible).

Question 2: Can you divide students and disregard as few preferences as possible?

General Problem: Constraint Satisfaction

Let D be a domain (i.e., a set of values) and $\mathcal{A}$ be a set of relations over D.

Constraint Satisfaction Problem (CSP($\mathcal{A}$))
Instance: A set of variables V and a set of constraints C of the form $R(v_1, \dots, v_r)$, where R is a relation from $\mathcal{A}$ and $v_1, \dots, v_r \in V$.
Question: Is there an assignment $f : V \to D$ that satisfies all constraints in C?

Almost Constraint Satisfaction

Almost Constraint Satisfaction Problem (AlmostCSP($\mathcal{A}$))
Instance: An instance (V, C) of CSP($\mathcal{A}$) and an integer k.
Question: Is there an assignment $f : V \to D$ that satisfies all but k constraints in C?

Toy Problem as a CSP

Variables: $V = \{x_1, \dots, x_n\}$ (one variable per student).
Domain: $\{1, 2\}$ (groups).
Relations: $\{=, \neq\}$.
Call the problem CSP($\mathcal{T}$).

Complexity:
CSP($\mathcal{T}$) is in P. AlmostCSP($\mathcal{T}$) is NP-hard and *fixed-parameter tractable*.

Parameterized Complexity of AlmostCSP

Parameter k – number of violated constraints.
PC of AlmostCSP($\mathcal{A}$) is only interesting if it is NP-hard and CSP($\mathcal{A}$) is in P.
Can be solved in $n^{O(k)}$ time by considering every subset of $n - k$ constraints.
Impractical even for small values of k.
If it can be solved in FPT time, i.e. $f(k) \cdot n^C$ for some function f of k and a constant C independent of k, then the algorithm is efficient for small values of k.

AlmostCSP Classification Project
For which $\mathcal{A}$ does AlmostCSP($\mathcal{A}$) admit an FPT algorithm?

Parameterized Complexity Theory

k-Vertex Cover

Goal: Cover all edges with k vertices.
Solvable in $f(k) \cdot n^C$ time - in FPT.

k-Independent Set

Goal: Find k non-adjacent vertices.
Solvable in $n^{O(k)}$ time - W[1]-hard.

Main Conjecture: FPT $\neq$ W[1].

Simple Temporal Problem (STP)

Introduced by Dechter, Meiri, and Pearl in 1989.
Variables $x_1, x_2, \dots, x_n$ represent points in time.
Constraints: $a \leq x_i - x_j \leq b$, where $a, b \in \mathbb{Q} \cup \{-\infty, \infty\}$.
Denote the problem by CSP($\mathcal{S}$).

$$
\begin{aligned}
-1 &\leq x_4 - x_1 \leq 4 \\
2 &\leq x_2 - x_1 \leq 4 \\
2 &\leq x_3 - x_2 \leq 3 \\
1 &\leq x_5 - x_4 \leq 2 \\
2 &\leq x_2 - x_4 \leq \infty \\
-\infty &\leq x_3 - x_5 \leq 1 \\
-\infty &\leq x_2 - x_5 \leq 3.
\end{aligned}
$$

An instance is satisfiable iff its distance graph has no negative cycle. In P.

AlmostSTP Classification

Theorem: Let $\mathcal{A} \subseteq \mathcal{S}$. Then AlmostCSP($\mathcal{A}$) is
- in constant time if $\mathcal{A}$ is trivial (all relations satisfied by setting $x_i = x_j$).
- in FPT if $\mathcal{A}$ only contains left/right-sided relations ($x_i - x_j \lesseqgtr c$ for $c \in \mathbb{Q}_{\geq 0}$).
- in FPT if $\mathcal{A}$ only contains equation relations ($x_i - x_j = c$ for $c \in \mathbb{Q}$).
- W[1]-hard otherwise.

Equation STP

An interesting special case. Satisfiable if and only if has no non-zero cycles.

$$
\begin{aligned}
x_1 - x_2 &= 1 \\
x_3 - x_2 &= 2 \\
x_4 - x_3 &= 3 \\
x_4 - x_5 &= 2 \\
x_6 - x_5 &= 1 \\
x_6 - x_1 &= 3 \\
x_3 - x_6 &= 2.
\end{aligned}
$$

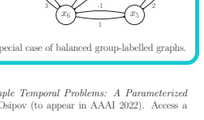

Generalize balanced signed graphs. Special case of balanced group-labelled graphs.

LINKÖPING UNIVERSITY

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

All results for Almost STP can be found in "*Resolving Inconsistencies in Simple Temporal Problems: A Parameterized Approach*" by Konrad K. Dabrowski, Peter Jonsson, Sebastian Ordyniak, George Osipov (to appear in AAAI 2022). Access a preliminary version here: https://tinyurl.com/almostp.

LaTeX TikZposter

High-Dimensional Bayesian Optimization with Gaussian Processes

Standard Bayesian Optimization (BO) is known to perform only well for up to 20-30 input dimensions. Optimizing higher-dimensional functions requires changes to the model or further assumptions on the problem itself. One line of current research focuses on sparse problems, where one assumes that the problem lies in a low-dimensional subspace of the higher-dimensional (ambient) space. Such methods perform BO in a lower-dimensional subspace that ideally captures the true effective subspace. Most algorithms for such problems, however, require an appropriate guess for the effective dimensionality as they rely on fixed embeddings. We present an algorithm that softens this requirement by introducing adaptive embeddings that increase the lower-dimensional subspace over time. Our algorithm outperforms the state-of-the-art on many benchmarks while being more computationally efficient than many contemporary approaches.

High-dimensional Bayesian Optimization with Adaptive Embeddings

Leonard Papenmeier, PhD Student, Lund University
Department of Computer Science
Supervisor: Dr. Luigi Nardi, Lund University, Coauthor: Matthias Poloczek, Amazon
Lund University

Motivation & Research Goals

Standard Bayesian Optimization (BO) is known to perform only well for up to 20 input dimensions [2]. Optimizing higher-dimensional functions requires changes to the model or further assumptions on the problem itself. Current research considers sparse problems where the problem lies in a low-dimensional subspace of a higher-dimensional (ambient) space. Such methods perform BO in a lower-dimensional subspace that aims at capturing the true effective subspace. Yet, most algorithms for such problems require an appropriate guess on the effective dimension. We present an algorithm (AdaTheSBO) that softens this requirement by using adaptive embeddings that increase the subspace dimension over time. [*Unpublished, preliminary state.*]

Problem and Algorithm

Minimization of expensive-to-evaluate black-box function $f : \mathbb{R}^D \mapsto \mathbb{R}$: $x^* \in \arg\min_{x \in \mathcal{X}} f(x)$, where $\mathcal{X}$ is D-dimensional ($D \gg 20$). We assume that there exists a low-(d-)dimensional ($d \leq 20$) subspace $\mathcal{Y}$ that can be mapped to by a linear embedding, and f is axis-aligned.

Use HeSBO embedding [3] to train an information-preserving Gaussian Process (GP) in trust region [1] of a subspace of increasing dimension.

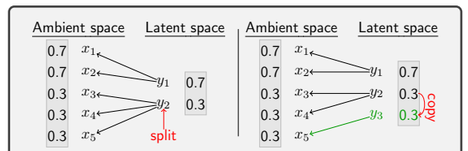

Figure 1: Increasing the dimension of the embedding from 2 to 3 (y_2 is *split*). Information can be preserved with the HeSBO embedding when increasing the dimension.

Algorithm 1 Ada-TheSBO Algorithm Outline

Require: initial latent dimension d
  sample random HeSBO embedding, defining up-projection S^T
  while not converged or budget available **do**
    while trust region sufficiently large **do**
      find candidate $\boldsymbol{x}^{(t)}$ by maximizing Thompson sample
      evaluate $f(S^T \boldsymbol{x}^{(t)})$; update GP; update TR
    end while
    if no progress in inner `while`-loop **then**
      re-start with new embedding and new GP
    else
      split latent dimension(s) with smallest GP length scale
    end if
    $d \leftarrow d + 1$
  end while
  Return Overall best $\boldsymbol{x}$ so far

Contributions

- First algorithm with an embedding of increasing dimension
- Outperforms state-of-the-art on a variety of problems
- Works in arbitrarily high-dimensions as long as d bounded

Selected Results

AdaTheSBO (with TuRBO and CMA-ES for comparison) on different input dimensions of the Ackley2 function. Empirically, AdaTheSBO is agnostic with respect to the ambient dimension for a fixed effective dimension.

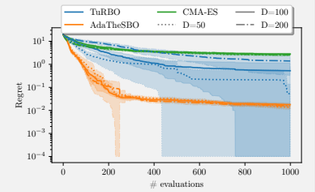

Strong performance on sparse-axis aligned, competitive on sparse, non-axis aligned problems. Poor performance on truly high-dimensional problems.

AdaTheSBO and TuRBO for comparison on different input dimensions of the Mopta08 function. The benchmark is assumed to be dense.

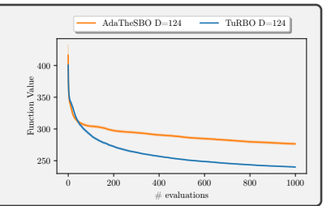

References

References

[1] David Eriksson et al. "Scalable global optimization via local bayesian optimization". In: Advances in Neural Information Processing Systems 32 (2019), pp. 5496–5507.

[2] Peter I Frazier. "A tutorial on Bayesian optimization". In: arXiv preprint arXiv:1807.02811 (2018).

[3] Amin Nayebi, Alexander Munteanu, and Matthias Poloczek. "A Framework for Bayesian Optimization in Embedded Subspaces". In: Proceedings of the 36th International Conference on Machine Learning. 2019.

Over the Air Computation For Machine Learning Over Wireless

With the increasing popularity of mobile devices and the development of the internet of things (IoT), accessibility to vast amounts of data has been grown. Further, the global number of connected IoT devices will reach more than 4 billion by 2024. On the flip side, taking advantage of large data sets can aid us in solving many complex problems in Machine Learning (ML). The primary challenges are communication latency, bandwidth consumption, energy limitations, privacy, and security. With limited communication resources, it is challenging to achieve efficient data aggregation over a large volume of IoT devices, as a critical point for exploiting the potential of the distributed ML. Unlike the standard "transmit-then-compute" approach, the over-the-air computation approach integrates communication and computation steps and provides ultra-fast wireless data aggregation in IoT networks.

AI MATH | Page 18 B
Razavikia, Saeed
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Over the Air Computation For Machine Learning Over Wireless

Saeed Razavikia, KTH Royal Institute of Technology
Electrical Engineering and Computer science
Main Advisor: Carlo Fischione

Motivation

With the increasing popularity of mobile devices and the development of the internet of things (IoT), accessibility to vast amounts of data has been grown. Further, the global number of connected IoT devices will reach more than 4 billion by 2024. On the flip side, taking advantage of large data sets can aid us in solving many complex problems in Machine Learning (ML). The primary challenges are communication latency, bandwidth consumption, energy limitations, privacy, and security. With limited communication resources, it is challenging to achieve efficient data aggregation over a large volume of IoT devices, as a critical point for exploiting the potential of the distributed ML. Unlike the standard "transmit-then-compute" approach, the over-the-air computation approach integrates communication and computation steps and provides ultra-fast wireless data aggregation in IoT networks.

Problem Setup

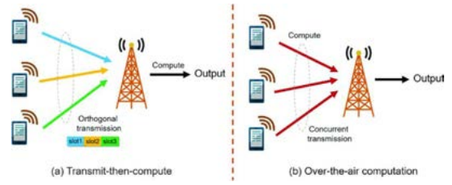

Chen et al. [1]

- Computing a function of input data at edge server by leveraging superposition property of electromagnetic waves
- Wireless communications become wireless computations

Benefits:

1. Integrate the communication and computation
2. Bandwidth is shared among all users in the time, frequency and code domain
3. More spectral efficient than transmit-then-compute
4. Transmission relies on analog communication
5. Fast wireless data aggregation

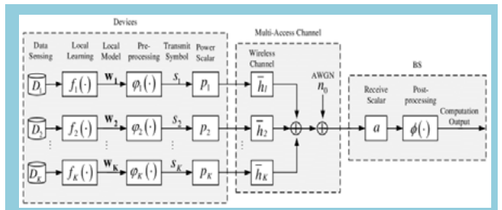

Ni et al. [2]

- Nomographic functions can be calculated
  - Mean $r = \frac{1}{K}\sum_{k=1}^{K} w_k$
  - Multiplication $r = \prod_{k=1}^{K} w_k$
  - Maximum $r = \max_k w_k$

Digital Communication

Analog Modulation

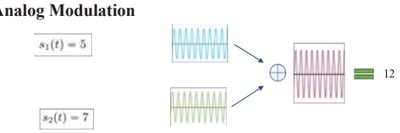

Digital Modulation

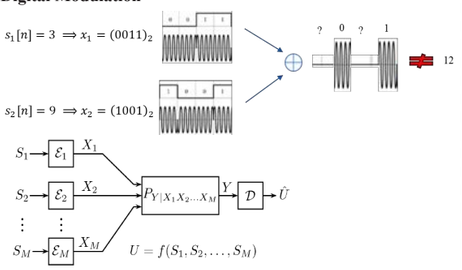

- Digital Communications can be seen as a probabilistic digital map/function between input code-words and output code-words
- In so doing, we can allow digital communication to perform probabilistic function computations
- The probability map i.e. $P_{(Y|X_1,X_2,\ldots,X_M)}$ can be designed by artificially enforcing the probability correspondence input output (changing the modulation formats)

References

1. X. Chen and Q. Qi, Convergence of Energy, Communication and Computation in B5G Cellular Internet of Things. Berlin, Germany: Springer, 2020
2. N. Wanli, L. Yuanwei, Y. Zhaohui, T. Hui, and S. Xuemin. "Federated learning in multi-RIS aided systems". IEEE Internet of Things Journal, 2021.

Greedy Causal Discovery is Geometric

Finding a directed acyclic graph (DAG) that best encodes the conditional independence statements observable from data is a central question within causality. Algorithms that greedily transform one candidate DAG into another given a fixed set of moves have been particularly successful, for example the GES, GIES, and MMHC algorithms. In 2010, Studený, Hemmecke and Lindner introduced the characteristic imset polytope, CIM_p , whose vertices correspond to Markov equivalence classes, as a way of transforming causal discovery into a linear optimization problem. We show that the moves of the aforementioned algorithms are included within classes of edges of CIM_p and that restrictions placed on the skeleton of the candidate DAGs correspond to faces of CIM_p . Thus, we observe that GES, GIES, and MMHC all have geometric realizations as greedy edge-walks along CIM_p.

Greedy Causal Discovery is Geometric

S. Linusson, P. Restadh and L. Solus
Email: linusson@kth.se, petter@kth.se, and solus@kth.se
KTH Royal Institute of Technology

Graphical Models

Directed Acyclic graphs (DAGs) are commonly used for complex models and causal interference [2]. DAGs encode conditional independence statements in the following way:

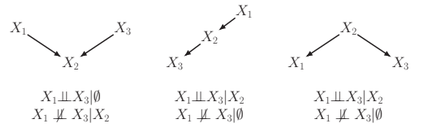

Goal 1 *Given data D on $X_1, \ldots, X_p$, find the DAG that "best" fits our data.*

Several algorithms have been proposed, for example PC, greedy SP, GES, GIES, and MMHC.

Characteristic Imset Polytope

The **characteristic imset** [4] of a DAG $\mathcal{G}$ is a vector in $\mathbb{R}^{\Upsilon}$ where $\Upsilon := \{S \subseteq [p] : |S| \geq 2\}$. It is defined as

$$c_{\mathcal{G}}(S) = \begin{cases} 1 & \text{if } \exists i \in S \text{ s.t. } S \subseteq \mathrm{pa}_G(i) \cup \{i\} \\ 0 & \text{otherwise.} \end{cases}$$

It was shown [4] that any (additive) decomposable score equivalent function can be written as a linear function in $\mathbb{R}^{\Upsilon}$. Thus we consider the polytopes

$$\mathrm{CIM}_p := \mathrm{conv}(c_{\mathcal{G}} : \mathcal{G} \text{ is a DAG with } p \text{ nodes})$$

and if we have two graphs $G \subseteq H$ we define

$$\mathrm{CIM}_{G,H} := \mathrm{conv}(c_{\mathcal{G}} : \mathcal{G} \text{ a DAG with skeleton } D \text{ such that } G \subseteq D \subseteq H).$$

Example 1 *Let us consider* $\mathrm{CIM}_{\smile}$.

Proposition 1 $\mathrm{CIM}_{G,H}$ *is a face of* CIM_p.

Hence Goal 1 can be formulated as the following.

Goal 2 *Given a (additive) decomposable score equivalent function s_D. Maximize s_D over* CIM_p.

GES, GIES, and MMHC all use the Bayesian Information Criterion (BIC) as a score function.

Edges of CIM_p and CIM_G

If $\mathcal{G}$ is a directed graph with $i \to j \in \mathcal{G}$ we denote by $\mathcal{G}_{i \leftarrow j}$ the directed graph identical to $\mathcal{G}$ except that the edge $i \to j$ is replaced with $i \leftarrow j$.

Proposition 2 *If $\mathcal{G}_{i \leftarrow j}$ is a DAG, then either $\mathcal{G}$ and $\mathcal{G}_{i \leftarrow j}$ are Markov equivalent, or* $\mathrm{conv}\left(c_{\mathcal{G}}, c_{\mathcal{G}_{i \leftarrow j}}\right)$ *is an edge of* CIM_G.

Let $\mathcal{G}$ be a DAG and assume i and j are not adjacent in the skeleton of $\mathcal{G}$. We denote by $\mathcal{G}_{i \leftarrow\!\!\rightarrow j}$ the directed graph identical to $\mathcal{G}$ with the edge $i \leftarrow j \in \mathcal{G}_{i \leftarrow\!\!\rightarrow j}$.

Proposition 3 *If $\mathcal{G}_{i \leftarrow\!\!\rightarrow j}$ is a DAG, then* $\mathrm{conv}\left(c_{\mathcal{G}}, c_{\mathcal{G}_{i \leftarrow\!\!\rightarrow j}}\right)$ *is an edge of* CIM_p.

We also show several more classes of edges.

Theorem 4 *The following causal discovery algorithms are greedy edge-walks along a face of* CIM_p*:*

1. *GES,*
2. *GIES with purely observational data,*
3. *MMHC, and*
4. *Greedy SP [1].*

Thus we can define two algorithms **Greedy CIM** and **Skeletal Greedy CIM** as greedy depth-first edge-walks along CIM_p and CIM_G respectively. By the above propositions **Greedy CIM** generalize GES and GIES with observational data. The graph G in **Skeletal Greedy CIM** was determined using conditional independence test similar to PC. A recurrent phased breadth-first version of **Greedy CIM** was as well implemented, as an easier comparison to GES and GIES.

Computational Results

The algorithms were implemented on simulated data using linear structural equation models with Gaussian noise. The code is available at [3].

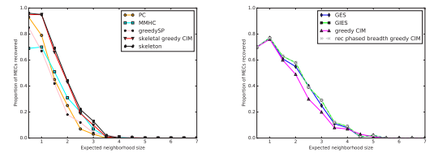

We see that among algorithms using conditional independence test **Skeletal Greedy CIM** performs better than previous algorithms. In the purely score based methods the breadth first algorithms have a higher recovery rate, but that can change if more edges of CIM_p were found and classified.

[1] F. Mohammadi, C. Uhler, C. Wang, and J. Yu, *Generalized permutohedra from probabilistic graphical models*, SIAM Journal on Discrete Mathematics, 32 (2018), pp. 64–93.

[2] J. Pearl, *Causality : Models, Reasoning, and Inference*, Cambridge University Press, Cambridge, U.K. New York, 2000.

[3] P. Restadh and L. Solus, *causalCIM. GitHub Repository*, 2021.

[4] M. Studený, R. Hemmecke, and S. Lindner, *Characteristic imset: A simple algebraic representative of a bayesian network structure*, Proceedings of the 5th European Workshop on Probabilistic Graphical Models, PGM 2010, (2010), pp. 257–265.

Algebraic Vision

The applications of reconstructing 3D models from 2D images include modelling of cities and objects for movies and video games, modelling clouds to predict the wheather, and helping robots and vehicles to orient themselves in new environments. Algebraic vision, which describes the algebraic component, is a prominent connection between methods from algebraic geometry and artificial intelligence. I investigate the geometry of points and lines projected onto the images of a set of cameras, and the stability of different approaches in the algebraic part of the reconstruction. This can help engineers in building new algorithms.

Algebraic Vision

Felix Rydell, KTH
Mathematics for Data and AI

Motivation & Research goals

The applications of reconstructing 3D models from 2D images include modelling of cities and objects for movies and video games, modelling clouds to predict the wheather, and helping robots and vehicles to orient themselves in new environments. *Algebraic vision*, which describes the algebraic component, is a prominent connection between methods from algebraic geometry and artificial intelligence. I investigate the geometry of points and lines projected onto the images of a set of cameras, and the stability of different approaches in the algebraic part of the reconstruction. This can help engineers in building new algorithms.

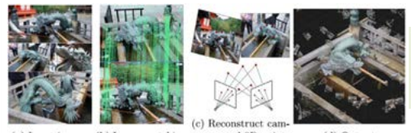
(a) Input images (b) Image matching (c) Reconstruct cameras and 3D points (d) Output

The reconstruction pipeline. Algebraic vision concerns itself with part c).

This table gives alternatives to part c). The first rows stand for number of cameras and the last row the number of solutions of camera positions **[1]**. The fewer number of solutions, the faster the algorithms.

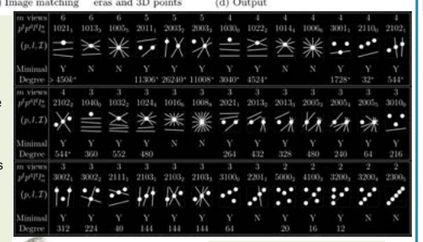

Two examples of reconstructions **[2]**. Once we have found the camera positions using an approach from the table above, we find the point/line cloud in 3D.

A systematic understanding of reconstruction is provided via 1) *projective* and 2) *algebraic* geometry.

1) *Projective space* $\mathbb{P}^n$ is the set of lines in $\mathbb{R}^{n+1}$ passing through the origin. It is a compactification of $\mathbb{R}^n$, meaning it is equal to $\mathbb{R}^n$ and some additional points at infinity.

2) *Algebraic geometry* is the study of systems of polynomials equations. A set of solutions to such a system is an *algebraic variety*.

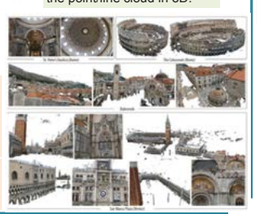

The Multiview Variety

Let $\mathtt{C} = \{C_1, \dots, C_m\}$ be a collection of cameras, meaning a set of m 3×4 matrices that project points and lines in $\mathbb{P}^3$ to their images in the cameras. Note that this map is not linear affinely.

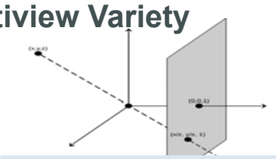

$$\Upsilon_C : \mathrm{Gr}(1,3) \dashrightarrow \mathrm{Gr}(1,2)^m,$$
$$L \mapsto (\ell_1, \dots, \ell_1) = (C_1 L, \dots, C_m L).$$

The *Grassmannian* $Gr(k,n)$ is the set of linear k-dimensional subspaces of $\mathbb{P}^n$ **[3]**.

The image of Υ_C describes the geometry of projected lines in the camera planes. Its Zariski closure is the smallest variety containing it, and is called the *multi-view variety*.

The point case is well-understood **[4]**.

A computer can understand an algebraic variety as a finite set of generators of a *defining ideal* in $\mathbb{R}[x_1, \dots, x_n]$.

There is a one-to-one correspondence sending an ideal $\mathcal{I}$ to the variety $\mathcal{V}(\mathcal{I}) = \{x \in \mathbb{R}^n : f(x) = 0 \ \forall f \in \mathcal{I}\}$

Algebraic vision can help practitioners in
- 1: Finding camera specifications,
- 2: Choosing point/line correspondences,
- 3: Improving speed of algorithms,
- 4: Determining robustness under noise.

Results

In joint work with E. Shehu, P. Breiding and A. Torres, we describe the multi-view variety geometrically and its defining ideal.

Theorem. The multi-view variety is the set of image lines whose back-projected planes meet in a line.

$$\begin{aligned}
&\text{if } m = 3, \text{ then } \mathrm{EDD} \geq 74;\\
&\text{if } m = 4, \text{ then } \mathrm{EDD} \geq 934;\\
&\text{if } m = 5, \text{ then } \mathrm{EDD} \geq 3651;\\
&\text{if } m = 6, \text{ then } \mathrm{EDD} \geq 9887;\\
&\text{if } m = 7, \text{ then } \mathrm{EDD} \geq 21807;\\
&\text{if } m = 8, \text{ then } \mathrm{EDD} \geq 42073;
\end{aligned}$$

There is numerical evidence that lines are more numerically robust than points. In this direction we consider the 1) *ED degree* **[5]**, the 2) *multi-degree* and numerical experiments.

1) The ED degree tells us how quickly we can ``denoise'' the data. The smaller the better.
2) The multi-degree tells us how ``curvy'' the variety is.

Future directions include explicit computations of ED degrees, theoretical understanding of numerical stability, and investigation of *rolling shutter cameras*.

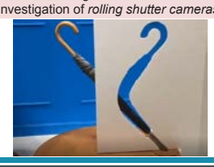

References

1. Duff, T., Kohn, K., Leykin, A., & Pajdla, T. (2019). *Plmp-point-line minimal problems in complete multi-view visibility.*
2. Agarwal, S., Furukawa, Y., Snavely, N., Curless, B., Seitz, S. M., & Szeliski, R. (2010). *Reconstructing rome*.
3. Gathmann, A. (2002). *Algebraic geometry*. Notes for a class taught at the University of Kaiserslauten *(2002/2003)*
4. Trager, M., Hebert, M., & Ponce, J. (2015). *The joint image handbook*.
5. Draisma, J., Horobeţ, E., Ottaviani, G., Sturmfels, B., & Thomas, R. R. (2016). *The Euclidean distance degree of an algebraic variety*.

LassoBench:
A High-Dimensional Hyperparameter Optimization Benchmark Suite for Lasso

Even though Weighted Lasso regression has appealing statistical guarantees, it is typically avoided due to its complex hyperparameter space described with thousands of hyperparameters. On the other hand, the latest progress with high-dimensional HPO methods for black-box functions demonstrates that high-dimensional applications can indeed be efficiently optimized. Despite this initial success, the high-dimensional hyperparameter optimization (HPO) approaches are typically applied to synthetic problems with a moderate number of dimensions which limits its impact in scientific and engineering applications. To address this limitation, we propose LassoBench, a new benchmark suite tailored for an important open research topic in the Lasso community that is Weighted Lasso regression. LassoBench consists of benchmarks on both well-controlled synthetic setups (number of samples, SNR, ambient and effective dimensionalities, and multiple fidelities) and real-world datasets, which enable the use of many flavors of HPO algorithms to be studied and extended to the high-dimensional Lasso setting. We evaluate 6 state-of-the-art HPO methods and 3 Lasso baselines, and demonstrate that Bayesian optimization and evolutionary strategies can improve over the methods commonly used for sparse regression while highlighting limitations of these frameworks in very high-dimension and noisy settings. Remarkably, TuRBO and CMA-ES improve the Lasso baselines on 60, 100, 300, and 1000 dimensional synthetic benchmarks, and the real-world benchmark based on the leukemia dataset by 42.3%, 23%, 22.3%, 12.6% and 75%, respectively.

LassoBench: A High-Dimensional HPO Benchmark Suite for Lasso

Kenan Šehić, Lund University (Computer Science)
Join work with Alexandre Gramfort[1], Joseph Salmon[2] and Luigi Nardi[3,4]
1-Université Paris-Saclay, Inria, CEA, France 2-IMAG, Université de Montpellier, CNRS,Montpellier, France 3-Lund University, Sweden 4-Stanford University, USA)

Summary

Even though Weighted Lasso regression has appealing statistical guarantees, it is typically avoided due to its complex hyperparameter space described with thousands of hyperparameters. On the other hand, the latest progress with high-dimensional HPO methods for black-box functions demonstrates that high-dimensional applications can indeed be efficiently optimized. Despite this initial success, the high-dimensional hyperparameter optimization (HPO) approaches are typically applied to synthetic problems with a moderate number of dimensions which limits its impact in scientific and engineering applications. To address this limitation, we propose **LassoBench**, a new benchmark suite tailored for an important open research topic in the Lasso community that is Weighted Lasso regression. LassoBench consists of benchmarks on both well-controlled synthetic setups (number of samples, SNR, ambient and effective dimensionalities, and multiple fidelities) and real-world datasets, which enable the use of many flavors of HPO algorithms to be studied and extended to the high-dimensional Lasso setting. We evaluate 6 state-of-the-art HPO methods and 3 Lasso baselines, and demonstrate that Bayesian optimization and evolutionary strategies can improve over the methods commonly used for sparse regression while highlighting limitations of these frameworks in very high-dimension and noisy settings. Remarkably, **TuRBO** [Eriksson, 2019] and **CMA-ES** [Hansen, 2016] improve the Lasso baselines on 60, 100, 300, and 1000 dimensional synthetic benchmarks, and the real-world benchmark based on the *leukemia* dataset by 42.3%, 23%, 22.3%, 12.6% and 75%, respectively.

LassoBench

We introduce a benchmark suite called **LassoBench** that addresses the limitations of current high-dimensional optimization benchmarks found in the literature while providing an opportunity for AutoML researchers to help advance Lasso research. New insights from the AutoML community will reflect directly on Lasso applications, whose seminal paper has so far been cited more than 40,000 times [Tibshirani, 1996].

LassoBench revolves around the non-convex optimization problem named Weighted Lasso regression, where the objective is to improve a linear model by optimizing the hyperparameters λ_j of the penalty term that promotes the sparsity in regression coefficients β^* [Bertrand, 2020]. The challenge is that the penalty term is defined typically in a high-dimensional setting (e.g., d=10^6) .

$$\boldsymbol{\beta}^*(\boldsymbol{\lambda}) \in \arg\min_{\boldsymbol{\beta}\in\mathbb{R}^d} \frac{1}{2n} \|\boldsymbol{y} - \mathbf{X}\boldsymbol{\beta}\|_2^2 + \sum_{j=1}^{d} e^{\lambda_j} |\beta_j|$$

LassoBench exposes a number of features, such as both noisy and noise-free benchmarks, well-defined effective dimensionality subspaces, and multiple fidelities, which enable the use of many flavors of Bayesian optimization algorithms to be improved and extended to the high-dimensional setting.

LassoBench includes the baselines commonly used in the Lasso community such as LassoCV [Massias, 2018], AdaptiveLassoCV [Massias, 2018] and Sparse-HO [Bertrand, 2020] for the comparison.

Benchmark Name	# Samples n	Ambient Dimensions d	Effective Dimensions d_e
synt_simple	30	60	3
synt_medium	50	100	5
synt_high	150	300	15
synt_hard	500	1000	50

Table 1 Predefined synthetic benchmarks in LassoBench when the true regression coefficients β_{true} are known.

Benchmark Name	# Samples n	Ambient Dimensions d	Approx. Effective Dimensions d_e
breast_cancer	683	10	3
diabetes	768	8	5
leukemia	72	7,129	22
dna	2,000	180	43
rcv1	20,242	19,959	75

Table 2 Real-world benchmarks in LassoBench. d_e is derived with Sparse-HO as $\widehat{d}_e = \|\widehat{\beta}\|_0$.

For a simple 4-line tutorial on how to run LassoBench follow github.com/ksehic/LassoBench

Results

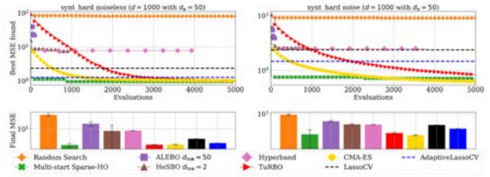

Figure 1 Baselines and HPO algorithms comparisons on synt_hard, left and right are noiseless and noisy, respectively. The bottom subplots show the best found estimations from each method, with confidence interval (for random methods) defined by one standard deviation out of 30 replications.

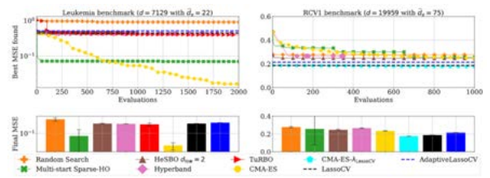

Figure 2 Comparison between the Lasso baselines and the HD-HPO methods for the Leukemia benchmark (left) and the RCV1 benchmark (right). The bottom subplot includes the best found MSE from each method and confidence intervals for random methods defined by one standard deviation out of 30 replications.

Method	Noise	synt_simple (d=60) (N=1000)	synt_medium (d=100) (N=1000)	synt_high (d=300) (N=5000)	synt_hard (d=1000) (N=5000)	Leukemia (d=7,129) (N=2000)	RCV1 (d=19,959) (N=1000)
LassoCV	False True	4.73 4.58	1.67 1.65	2.48 2.48	2.37 2.38	0.44 NA	0.18 NA
Adaptive LassoCV	False True	2.06 7.98	1.52 2.48	1.18 1.32	1.27 1.46	0.51 NA	0.21 NA
Multi-start Sparse-HO	False True	0.697 ± 0.34 0.59 ± 0.31	1.23 0.73 ± 0.49	1.11 ± 0.92 0.76 ± 0.37	**0.96** ± 0.27 0.71 ± 0.58	0.06 ± 0.1 NA	0.25 ± 0.17 NA
Random Search	False True	67.22 ± 58.9 8.31 ± 6.9	60.68 ± 35.5 7.93 ± 3.6	69.41 ± 21.3 8.83 ± 2.0	78.45 ± 13.6 8.96 ± 1.1	0.85 ± 0.21 NA	$0.27 \pm 8e\text{-}3$ NA
CMA-ES	False True	**0.695** ± 0.08 0.34 ± 0.1	1.07 ± 0.06 **0.48** ± 0.08	0.96 ± 0.03 0.64 ± 0.06	1.00 ± 0.02 **0.62** ± 0.03	**0.015** $\pm 7e\text{-}3$ NA	$0.23 \pm 3e\text{-}3$ NA
CMA-ES $\lambda_{LassoCV}$	False True	NA NA	NA NA	NA NA	NA NA	NA NA	**0.17** $\pm 2e\text{-}4$ NA
ALEBO	False True	14.59 ± 26.1 4.95 ± 3.5	18.16 ± 14.1 4.48 ± 2.6	21.68 ± 18.4 4.75 ± 1.7	21.84 ± 7.1 3.89 ± 0.5	NA NA	NA NA
HeSBO	False True	3.20 ± 0.2 3.56 ± 0.6	1.74 ± 0.2 1.74 ± 0.1	2.66 ± 0.2 2.82 ± 0.4	7.57 ± 10.0 2.56 ± 0.3	$0.45 \pm 2e\text{-}2$ NA	$0.24 \pm 7e\text{-}3$ NA
Hyperband	False True	1.52 ± 0.3 1.44 ± 0.3	4.53 ± 3.2 1.94	5.38 2.49	7.87 2.51	0.43 NA	$0.26 \pm 2e\text{-}3$ NA
TuRBO	False True	0.78 ± 0.7 **0.30** ± 0.07	**0.95** ± 0.1 0.55 ± 0.1	**0.90** ± 0.03 **0.59** ± 0.1	1.01 ± 0.02 0.84 ± 0.09	$0.39 \pm 9e\text{-}2$ NA	NA NA

Table 3 Best-found MSE obtained for all optimizers, the synthetic benchmarks with both conditions (noiseless and noisy) and the real-world benchmarks based on the leukemia dataset and RCV1. We report means and standard deviation across 30 runs of each optimizer with N as the number of evaluations. For each benchmark, bold face indicates the best MSE.

For more details, follow our preprint via arxiv.org/abs/2111.02790

References

1. R. Tibshirani. Regression Shrinkage and Selection via the Lasso. Journal of the Royal Statistical Society. Series B (Methodological), 58(1), 267–288. 1996
2. Q. Bertrand, Q. Klopfenstein, M. Blondel, S. Vaiter, A. Gramfort, and J. Salmon. Implicit differentiation of Lasso-type models for hyperparameter optimization. In Proceedings of the 37th International Conference on Machine Learning, pages 119:810–821, 2020
3. D. Eriksson, M. Pearce, J. Gardner, R. D. Turner, and M. Poloczek. Scalable global optimizationvia local Bayesian optimization. In Advances in Neural Information Processing Systems, pages 5496–5507, 2019
4. M. Massias, A. Gramfort, and J. Salmon. Celer: a Fast Solver for the Lasso with Dual Extrapolation. In ICML, volume 80, pages 3315–3324, 2018.
5. Hansen N. The CMA Evolution Strategy: A Tutorial. ArXiv e-prints, arXiv:1604.00772 [cs.LG], 2016.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Sparsification of Infinite-domain CSPs

Many problems encountered in computer science and mathematics can be viewed as CSPs: for example in spatio-temporal reasoning, computer vision, machine learning, scheduling, and bioinformatics, and this makes the CSP a problem of central importance.
The goal of this research is to study the complexity of constraint satisfaction problems (CSPs) over infinite domain.
There has been a lot of results for sparsification of finite domain CSPs, but not as many in the case of infinite-domain CSPs.
It is an important direction of research since a domain of infinite size can capture many problems encountered in AI and logical reasoning, that cannot be formalised in finite-domain.
We aim to construct faster algorithms and methods that may be useful to analyse the kernelisation of parameterized versions of infinite-domain CSPs.

AI MATH | Page 22 B
Sharma, Abhijat
Linköping University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Sparsification of Infinite-domain CSPs

Abhijat Sharma, PhD Student, Linköping University
TCSLab, IDA
Supervisors: Prof. Peter Johnsson (LiU) and Dr. Victor Lagerkvist (LiU)
LINKÖPINGS UNIVERSITET

Motivation & Research Goals

Many problems encountered in computer science and mathematics can be viewed as CSPs: for example in spatio-temporal reasoning, computer vision, machine learning, scheduling, and bioinformatics, and this makes the CSP a problem of central importance. The goal of this research is to study the complexity of constraint satisfaction problems (CSPs) over infinite domain. There has been a lot of results for sparsification of finite domain CSPs, but not as many in the case of infinite-domain CSPs. It is an important direction of research since a domain of infinite size can capture many problems encountered in AI and logical reasoning, that cannot be formalised in finite-domain. We aim to construct faster algorithms and methods that may be useful to analyse the kernelisation of parameterized versions of infinite-domain CSPs.

Background

What is a CSP?
An instance of the constraint satisfaction problem consists of the following as input:

- A set of variables V and a domain D of allowed values for the variables.
- A set of constraints C imposing certain restrictions on the value assignments to the variables.

The solution to a CSP is a value-assignment $f : V \mapsto D$ such that all constraints are satisfied. The most prominent complexity-theoretic questions concerning CSPs is the following: given a set of relations Γ (the constraint language), what is the complexity of CSP(Γ), i.e. the CSP where the constraints contain only the relations from Γ.

> **Finite Domain Dichotomy:** Bulatov([4]) and Zhuk([1]) independently proved the long-standing conjecture: a finite-domain CSP(Γ) is either polynomial-time solvable or NP-complete.

Infinite-domain CSPs are undecidable in general, however there exist many well-understood fragments which admit complexity dichotomies. There seems to be significant variance in the time-complexity of several CSPs, even if most of them are NP-complete. Even though there is vast research on fine-grained complexity of finite-domain and infinite-domain CSPs([2]), it is interesting to study the fine-grained complexity of infinite-domain CSPs restricted to certain kinds of constraints.

Sparsification and Kernelisation
Kernelisation is a pre-processing algorithm that takes an input instance and reduces it to an smaller equivalent instance, called a "kernel". Our goal is to compute efficient kernels for CSPs parameterized by the number of input variables.
In the context of CSPs, we achieve kernelisation by efficiently reducing the number of constraints in terms of the number of variables, while preserving the solution. This is also known as sparsification of CSP instances.

References

[1] A Proof of the CSP Dichotomy Conjecture
Zhuk, Dmitriy
Foundations of Computer Science (FOCS), 2017

[2] A Survey on the Fine-grained Complexity of Constraint Satisfaction Problems Based on Partial Polymorphisms
Couceiro, Miguel and Haddad, Lucien and Lagerkvist, Victor
Journal of Multiple-Valued Logic and Soft Computing, 2020

[3] The Complexity of Equality Constraint Languages
Bodirsky, Manuel and Kára, Jan
Theory of Computing Systems, July 2008

[4] A Dichotomy Theorem for Nonuniform CSPs
Bulatov, Andrei A.
Foundations of Computer Science (FOCS), 2017

[5] Optimal Sparsification for Some Binary CSPs Using Low-degree Polynomials
Jansen, Bart M. P. and Pieterse, Astrid
Mathematical Foundations of Computer Science (MFCS), 2016

Preliminary Results

Equality CSPs
For initial results, we focus on a specific class of CSPs, where the constraint language Γ only consists of equality relations.([3])

> A relation $R \subseteq N^k$ is an equality relation of arity k if it can be defined as $R = \{(x_1, x_2, ..., x_k) : \phi(x_1, x_2, ..., x_k)\}$ where ϕ is a first-order formula over the structure $(N; =)$.

When a constraint language Γ contains only relations of arity at-most k, there exists a trivial sparsification of any CSP(Γ) instance to $O(n^k)$ constraints. Our goal is to either achieve sparsification that is better than this bound, or prove that such sparsification does not exist.

Kernel Lower Bounds
One of the most powerful tools used to analyse complexity of finite-domain CSPs is the standard *algebraic approach*. This involves constructing a framework that allows polynomial-time solution-preserving reductions between constraint languages, and their associated CSPs.
We have introduced algebraic methods that extend the above framework for obtaining stronger lower bounds on kernel size. The following result makes use of QFPP-definitions and additionally some novel reduction techniques inspired from the existing framework.

> **Lower Bound:** Let Γ be an equality language such that CSP(Γ) is NP-hard. Then CSP(Γ) admits no kernel of size $O(n^{2-\epsilon})$ where n is the number of variables and $\epsilon > 0$, unless NP $\subseteq$ co-NP/poly.

Sparsification Techniques
We have introduced new sparsification methods based on the *linear-algebraic framework* of viewing constraints as low-degree polynomials([5]). We have applied these ideas to equality constraints and obtained optimal results in certain cases. A natural research direction now is to better understand sparsification of equality constraints, and preferably obtain optimal bounds for all equality languages. This will need the development of even more powerful methods.

Beyond Equality Relations
Apart from equality relations, our aim is to generalise the linear-algebraic techniques to more interesting cases such as temporal constraints over the domain of rational numbers. For example, consider the following well-studied relation, used for gene mapping in bioinformatics.

> **Betweenness:** $B = \{(x, y, z) \in Q^3 \mid x < y < z \text{ or } z < x < y\}$

Our techniques allow us to sparsify both the above relations to a kernel of $O(n^2)$ constraints. This is encouraging as it illustrates that the our methods are applicable to CSPs far beyond equality relations. Our goal is to utilize these sparsification techniques to more complex spatio-temporal constraints used in AI, for instance the Allen's Interval Algebra and RCC-8 Calculus.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Toft, Albin
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Scalable Causal Inference in Mass Media

This project centers around the theory of causality and its applications to the complex data sets arising from social media platforms, mass media and the financial market. The primary industrial objective of the project is to gain a systematic understanding of the cause-effect network underlying (i) events reported in mass media, (ii) individual interactions in social media and (iii) measurable financial and economic indicators in the globally-coupled markets. As a first step in this direction, we are exploring a different approach to causal inference in time dependent data using Hawkes processes in contrast to the more classical time series approach.

Toft, Albin
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Scalable Causal Inference in Mass Media

Albin Toft, Ind. PhD, Combient Mix and KTH
Dept. Statistics, Causal Inference in Financial Time Series
Supervisors: Liam Solus (KTH), Raazesh Sainudiin (Combient Mix) and Tatjana Pavlenko (UU)

KTH VETENSKAP OCH KONST

mix

Motivation & Research Goals

This project centers around the theory of causality and its applications to the complex data sets arising from social media platforms, mass media and the financial market. The primary industrial objective of the project is to gain a systematic understanding of the cause-effect network underlying (i) events reported in mass media, (ii) individual interactions in social media and (iii) measurable financial and economic indicators in the globally-coupled markets. As a first step in this direction, we are exploring a different approach to causal inference in time dependent data using Hawkes processes in contrast to the more classical time series approach.

Granger Causality & Time Series Analysis

[1] When considering causality in time series, Granger causality and its variations are popular approaches. Consider a multivariate time series $(\mathbf{X}_t)_{t\in Z}$, such that the induced joint distribution is faithful with respect to the corresponding full time graph. Then the summary graph has an arrow $X^j \to X^k$ if and only if there exist a $t \in Z$ such that

$$X_t^k \not\perp X_{past(t)}^j | X_{past(t)}^{-j}.$$

Typically, one way of determining whether one time series Granger causes another, is by modelling the multivariate time series as vectorized auto regressive (VAR) processes

$$\mathbf{X}_t = \mathbf{c} + \sum_{i=1}^{p} \mathbf{A}_i \mathbf{X}_{t-i} + \epsilon_t.$$

The task of determining the Granger causal relationships among the time series, boils down to assessing which elements of the matrices $\mathbf{A}_i, i = 1...p$ are non-zero.

Hawkes Processes

[2] A multi-dimensional Hawkes process is a counting process, where the intensity of each separate counting process at time t can be written as

$$\lambda_i(t) = \mu_i + \sum_{j=1}^{D} \sum_{t_k^j < t} \phi_{ij}(t - t_k^j).$$

The $\phi_{ij}(t)$ functions are called kernel functions, and typically one uses exponential kernels where

$$\phi_{ij}(t) = \alpha_{ij}\beta_{ij} exp\{-\beta_{ij} t\}.$$

These kernels can be used to describe how events of type j might increase the intensity of events of type i occurring.

References

[1] Elements of Causal Inference
J.Peters, D.Janzing, B.Schölkopf
The MIT Press
2017

[2] Learning Granger Causality for Hawkes Processes
H.Xu, M. Farajtabar, H.Zha
2018

Selected Results

VAR-Results

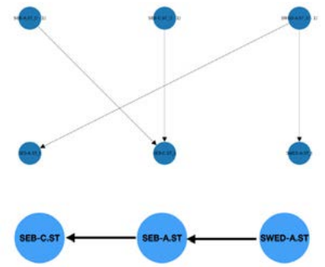

Modelling the log-returns of the stocks SEB-A.ST, SEB-C.ST and SWED-A.ST, as VAR processes the current results show that some Granger causal connections between the time series can be found.

Hawkes-Results

By instead estimating trend reversal points in the time series, it was possible to extract event sequences for the stock prices. The interactions between these event sequences were then modelled using multi dimensional Hawkes processes with exponential kernels.

Planned Work

Most of the research is still to be done in this project, and some ideas are:

- Model expansion: Extend the models to analyse more stocks and over a longer time period.
- Mass media action as intervention: Include news and mass-media data into the analysis.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Homotopical decompositions of simplicial and Vietoris-Rips complexes

When we decompose a simplicial complex and reassemble it, it might happen that the resulting complex has a different homotopy type from the initial one. However, it is sometimes possible to understand this change by looking at subcomplexes living in the intersection of the two decomposing pieces, the so called obstruction complexes. In this poster it is outlined how the homotopy type of a simplicial complex is related to the one of its decompositions. It is also explained with an example how to use these ideas to find out the homotopy type of given Vietoris-Rips complexes. This is a joint work with Wojciech Chachólski, Martina Scolamiero and Alvin Jin.

AI MATH | Page 25 B
Tombari, Francesca
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Homotopical decompositions of Vietoris-Rips complexes

Wojciech Chachólski, Alvin Jin, Martina Scolamiero, Francesca Tombari
Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden

① Introduction

The use of algebraic topology is rapidly growing in understanding data. The general pipeline of TDA can be summarized by the following:

Dataset - Points with relations between them

⇓

Geometrical structure - Simplicial complex

⇓

Algebraic invariants - Homology

⇓

Features for statistics and machine learning

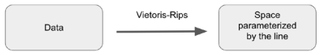

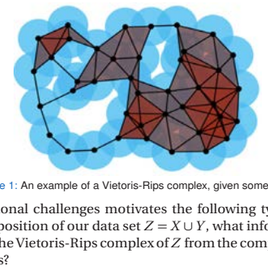

Figure 1: An example of a Vietoris-Rips complex, given some $r > 0$.

The computational challenges motivates the following type of question: given a decomposition of our data set $Z = X \cup Y$, what information can we recover about the Vietoris-Rips complex of Z from the component Vietoris-Rips complexes?

② A general approach

Let K be a simplicial complex, $K_0 = X \cup Y$ the set of its vertices and $A = X \cap Y$. Let $K_X = K \cap X$ and $K_Y = K \cap Y$. One can easily notice that the union of K_X and K_Y does not give the initial simplicial complex K. A natural question that might arise is whether the following inclusion is a homotopy equivalence or not

$$K_X \cup K_Y \hookrightarrow K.$$

Figure 2: Example of a simplicial complex with high complexity. (Image courtesy of the authors of arXiv:1608.03520)

A special case of this problem occurs when a pseudo-metric space (Z, d) is considered. Fixing $r > 0$ and a covering of Z consisting in two subspaces X and Y, we get the inclusion

$$VR_r(X) \cup VR_r(Y) \hookrightarrow VR_r(Z).$$

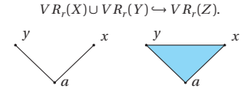

Figure 3: The two figures show a simplicial complex K (on the right) and $K_X \cup K_Y$ (on the left), where $X = \{x, a\}$ and $Y = \{y, a\}$.

③ Main result

We define the obstruction complex:

$$F(\sigma, A) := \{\mu \subset A \,|\, \mu \cup \sigma \in K\}.$$

Theorem. Let $\mathscr{C}$ be a closed collection of simplicial sets. If, for every σ in $\{\sigma \in K \,|\, \sigma \cap X \neq \emptyset \text{ and } \sigma \cap Y \neq \emptyset \text{ and } \sigma \cap A = \emptyset\}$, the simplicial complex $F(\sigma, A)$ satisfies $\mathscr{C}$, then the homotopy fibers of $K_X \cup K_Y \subset K$ also satisfy $\mathscr{C}$.

Corollary. If, for every σ as above, the simplicial complex $F(\sigma, A)$ is contractible, then $K_X \cup K_Y \subset K$ is a weak equivalence.

We get a long exact sequence in the case when adding one vertex:

$$H_n(F(x, A)) \to H_n(K_A) \to H_n(K) \to H_{n-1}(F(x, A)) \to H_{n-1}(K_A)$$

and another when adding two vertices:

$$H_n(\Sigma F(x, y, A) \to H_n(K_X \cup K_Y) \to H_n(K) \to H_{n-1}(\Sigma F(x, y, A) \to H_{n-1}(K_X \cup K_Y).$$

These sequences give information about the global homology of K with respect to local information.

④ Examples

Consider the metric space $Z = \{x_1, x_2, a_1, a_2, y\}$, with the metric such that every two points of Z has distance 1 except for x_1, a_2 and x_2, a_1 having distance 1.1. Let $X = \{x_1, x_2, a_1, a_2\}$, $Y = \{y, a_1, a_2\}$ be a cover for Z. We can easily see that $VR_1(X) \cup VR_1(Y)$ has the homotopy type of S^1, while $VR_1(Z)$ is contractible. This is due to the fact that $F(\sigma, A)$ is empty, hence non-contractible, when σ is the 2-simplex with vertices x_1, x_2 and y.

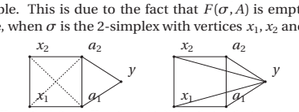

Figure 4: $K_X \cup K_Y$ on the left and K on the right. Notice that all the triangles in this example are filled, because K is a clique complex.

The following picture shows an example of a decomposition of $Z = \{x_1, x_2, y_1, y_2, a_{11}, a_{12}, a_{21}, a_{22}\}$ that has the same homology as the total simplicial complex up to degree 2, but different H_3.

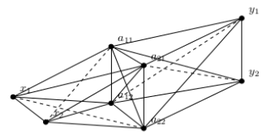

Figure 5: The figure represents a 2-dimensional visualization of the Vietoris-Rips complex $VR_r(X) \cup VR_r(Y)$. $VR_r(X \cup Y)$ is obtained by the above simplicial complex adding the simplex $\{x_1, x_2, y_1, y_2\}$

The metric is given by:

$$\begin{aligned}
&d(a_{11}, a_{21}) = d(a_{11}, a_{12}) = d(a_{21}, a_{22}) = d(a_{12}, a_{22}) = 4,\\
&d(a_{11}, a_{22}) = d(a_{12}, a_{21}) = 6,\\
&d(x_1, a_{11}) = d(y_1, a_{21}) = d(x_2, a_{22}) = d(y_2, a_{12}) = 3,\\
&d(x_1, a_{12}) = d(y_1, a_{11}) = d(x_2, a_{21}) = d(y_2, a_{22}) = 5,\\
&d(x_1, a_{21}) = d(y_1, a_{22}) = d(x_2, a_{12}) = d(y_2, a_{11}) = 7,\\
&d(x_1, a_{22}) = d(y_1, a_{12}) = d(x_2, a_{11}) = d(y_2, a_{21}) = 9,\\
&d(x_1, x_2) = d(y_1, y_2) = 6,\\
&d(x_1, y_1) = d(x_1, y_2) = d(x_2, y_1) = d(x_2, y_2) = 8.
\end{aligned}$$

As we have already noticed, the study of this problem for Vietoris-Rips complexes is actually a consequence of the same problem stated for generic simplicial complexes. Analogously, the conditions that we put on a metric space are just a translation of hypothesis on simplicial complexes.

⑤ References

[1] W. Chacholski, A. Jin, M. Scolamiero, and F. Tombari. Homotopical decompositions of simplicial and Vietoris–Rips complexes. J Appl. and Comput. Topology 5, 215–248 (2021). https://doi.org/10.1007/s41468-021-00066-2

[2] Adamaszek et. al. On Homotopy Types of Vietoris–Rips Complexes of Metric Gluings. Proceedings of the 34th Symposium on Computational Geometry (2018), 3:1-3:15.

Upadhyaya, Manu
Lund University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Performance estimation of iterative algorithms and closed-loop systems

The aim of this research is to combine ideas from the performance estimation problem (PEP) framework in the optimization literature and integral quadratic constraints (IQC) framework from the control theory literature into a novel computer-aided automated Lyapunov analysis framework. Applications that we are considering are 1) the analysis of the worst-case performance of optimization algorithms and iterative algorithms in general and 2) the stability verification of neural network controlled systems and model predictive control (MPC) schemes. Moreover, besides the analysis in 1) and 2), the framework allows for a systematic approach to optimize algorithm or system performance with respect to design parameters.

AI MATH | Page 26 B

Upadhyaya, Manu
Lund University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Performance estimation of iterative algorithms and closed-loop systems

Manu Upadhyaya, Lund University
Department of Automatic Control
Main advisor: Pontus Giselsson

LUND UNIVERSITY

www.control.lth.se

Motivation & research goals

The aim of this research is to combine ideas from the performance estimation problem (PEP) framework in the optimization literature and integral quadratic constraints (IQC) framework from the control theory literature into a novel computer-aided automated Lyapunov analysis framework. Applications that we are considering are 1) the analysis of the worst-case performance of optimization algorithms and iterative algorithms in general and 2) the stability verification of neural network-controlled systems and model predictive control (MPC) schemes. Moreover, besides the analysis in 1) and 2), the framework allows for a systematic approach to optimize algorithm or system performance with respect to design parameters.

Methods

Performance estimation problem (PEP)
The PEP framework, first introduced in [1], provides a systematic method to analyze the worst-case performance of optimization algorithms, and iterative algorithms in general. Roughly speaking, in the optimization algorithm case and in the basic set-up, we have the following components:

- An appropriate class of functions $\mathcal{F}$ with members $f:\mathcal{H}\to\mathbb{R}\cup\{\pm\infty\}$ for some real Hilbert space $\mathcal{H}$.
- Each function $f\in\mathcal{F}$ has a minimizer x^*.
- Some fixed oracle $\mathcal{O}_f(x)$ that provides information about f at x. This could include the function value, and/or the gradient, etc.
- Some initial iterate $x_0\in\mathcal{H}$.
- A fixed algorithm $\mathcal{A}$ that is allotted N iterations such that it generates a sequence

$$x_1=\mathcal{A}_1(x_0,\mathcal{O}_f)$$
$$x_2=\mathcal{A}_2(x_0,x_1,\mathcal{O}_f)$$
$$\vdots$$
$$x_N=\mathcal{A}_N(x_0,x_1,\cdots,x_{N-1},\mathcal{O}_f).$$

- An appropriate performance metric $\mathcal{P}(x^*,x_0,x_1,\cdots,x_N,\mathcal{O}_f)$. Some simple examples include function value suboptimality $f(x_N)-f(x^*)$, norm of gradient $\|\nabla f(x_N)\|$ and distance to an optimal solution $\|x^*-x_N\|$.

The performance estimation problem (PEP) is then to find the worst-case performance: I.e., maximize

$$\mathcal{P}(x^*,x_0,x_1,\cdots,x_N,\mathcal{O}_f)$$

subject to

$$f\in\mathcal{F}$$
$$x^*\text{ is optimal for }f$$
$$x_1,\cdots,x_N\text{ is generated by }\mathcal{A}\text{ with initial point }x_0$$

with some additional technical assumptions such that the problem becomes well-posed.

Calculating the worst-case performance is in general an infinite-dimensional optimization problem. Luckily, there exist standard techniques in the PEP literature that render the problem tractable by transforming it into a finite-dimensional semidefinite program and do so tightly via so called interpolation conditions. See [2] for additional details. Moreover, the framework allows to select "optimal" design parameters in the algorithm $\mathcal{A}$ by minimizing the worst-case performance of $\mathcal{A}$.

Recently in [3], this framework has been adapted to finding tight contraction factors of fixed-point operators used in splitting schemes to solve monotone inclusion problems.

Integral quadratic constraints (IQC)
The IQC framework, see [4], can be used to analyze the case of a linear system interconnected in feedback to a, possibly uncertain, nonlinear system. In particular, [5] noticed that the IQC framework can be used in the analysis and design of optimization algorithms and [6] highlighted the close connection to the PEP framework. Moreover, in [7], tools from the IQC framework were used to develop a method to certify asymptotic stability of neural network-controlled systems.

Current work

We are currently considering optimization algorithms and splitting schemes that:

- Can be written as a linear system with a nonlinear feedback given by some operator, in accordance with the IQC framework.
- The operators involved have interpolation conditions that only involve quadratic inequalities enabling the use of the PEP framework.
- A quadratic Lyapunov function ansatz to obtain worst-case performance guarantees and extract either linear or sublinear rates of convergence.

Concurrently, within the same framework, we are considering:

- Analyzing the stability of MPC schemes given an iteration count constraint on the optimization algorithm.
- The stability verification and training of neural network-controlled systems.

References

1. Y. Drori and M. Teboulle (2014). Performance of first-order methods for smooth convex minimization: a novel approach. Mathematical Programming 145: 451-482
2. A.B. Taylor, J.M. Hendrickx and F. Glineur (2017). Smooth strongly convex interpolation and exact worst-case performance of first-order methods. Mathematical Programming 161: 307-345
3. E. K. Ryu, A. B. Taylor, C. Bergeling, and P. Giselsson (2020). Operator splitting performance estimation: tight contraction factors and optimal parameter selection, SIAM Journal on Optimization 30:3, 2251-2271
4. A. Rantzer (1997). System analysis via integral quadratic constraints. IEEE Transactions on Automatic Control 42(6): 819-830
5. L. Lessard, B. Recht and A. Packard (2016). Analysis and design of optimization algorithms via integral quadratic constraints (2016). Siam Journal on Optimization 26(1): 57-95
6. A. Taylor, B. Van Scoy, L. Lessard (2018). Lyapunov functions for first-order methods - tight automated convergence guarantees. 35th International Conference on Machine Learning, PMLR 80: 4897-4906
7. H. Yin, P. Seiler and M. Arcak (2021). Stability analysis using quadratic constraints for systems with neural network controllers. IEEE Transactions on Automatic Control

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Vallin, Jonatan
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Geometry and Approximation of ReLU Networks

In our upcoming paper we study the geometry and approximation properties of fully-connected ReLU networks. We start by describing the structure of a standard ReLU layer by introducing a convenient partition of the input space. Using this structure, we characterize the geometry of the decision boundary for shallow networks. We end our analysis by deriving approximation results for deep ReLU networks (not presented in this poster).

AI MATH | Page 27 B

Vallin, Jonatan
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Geometry and Approximation of ReLU Networks

Jonatan Vallin, Umeå University
UMEÅ UNIVERSITY
Department of Mathematics and Mathematical Statistics

Abstract

In our upcoming paper we study the geometry and approximation properties of fully-connected ReLU networks. More precisely, we consider networks $F\colon \mathbb{R}^d \to \mathbb{R}$ of the form $F(x) = L \circ T_N \circ \cdots \circ T_1(x)$, where each mapping $T_i\colon \mathbb{R}^d \to \mathbb{R}_+^d$ is a ReLU layer and $L\colon \mathbb{R}^d \to \mathbb{R}$ is an affine functional. Since $\mathrm{ReLU}\,(x) = \max(x, 0)$ is a piecewise linear map, the network F defines a piecewise linear function subordinate to a polygonal partition of the input space. We start by describing the structure of a standard ReLU layer by introducing a convenient partition of the input space. Using this structure, we characterize the geometry of the decision boundary $\Gamma = \{x \in \mathbb{R}^d\colon F(x) = 0\}$ for shallow networks. We end our analysis by deriving approximation results for deep ReLU networks (not presented in this poster).

Structure of ReLU Layers

A ReLU layer is a mapping $T\colon \mathbb{R}^d \to \mathbb{R}_+^d$ of the form

$$T(x) = \mathrm{ReLU}(Ax + b)$$

where $A \in \mathbb{R}^{d\times d}$ is a matrix, assumed to have full rank, with rows $a_i \in \mathbb{R}^d$ and $b \in \mathbb{R}^d$ is a vector with elements $b_i \in \mathbb{R}$. To describe the action of T, we introduce a set of dual vectors $\{a_i^*\colon i \in I\}$, where $I = \{1, \dots, d\}$, defined by the equation $a_i^* \cdot a_j = \delta_{ij}$. It follows that these vectors constitute a basis of $\mathbb{R}^d$ which will be convenient when analyzing the structure of T.

To that end, consider a partition $I_+ \cup I_- \cup I_0$ of the index set I, denoted by the three-tuple (I_+, I_-, I_0). Some of the sets may be empty as long as their union is I. For each such partition, we define

$$S_{(I_+,I_-,I_0)} = \left\{x_0 + \sum_{i\in I_+} \alpha_i a_i^* - \sum_{i\in I_-} \alpha_i a_i^* : \alpha_i > 0\right\} \subset \mathbb{R}^d$$

where x_0 is the unique solution to $Ax = -b$. By construction, $\dim\left(S_{(I_+,I_-,I_0)}\right) = |I_+ \cup I_-|$ and if $\mathcal{I}$ is the set of all such three-tuples, the family $\mathcal{S} = \{S_I\colon \boldsymbol{I} \in \mathcal{I}\}$ is a partition of $\mathbb{R}^d$ with pairwise disjoint sets. In the special case when $A = I_d$ (I_d being the identity matrix) and $b = 0$, the sets reduce to

$$\hat{S}_{(I_+,I_-,I_0)} = \left\{0 + \sum_{i\in I_+} \alpha_i e_i - \sum_{i\in I_-} \alpha_i e_i : \alpha_i > 0\right\} \subset \mathbb{R}^d$$

where e_i is the i:th Euclidean basis vector. We call the corresponding partition $\hat{\mathcal{S}}$. Examples of the families $\hat{\mathcal{S}}$ and $\mathcal{S}$ are shown in Figure 1.

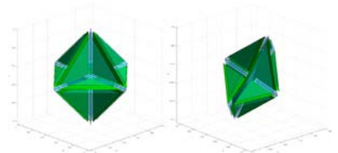

Figure 1. A visualization of the families $\hat{\mathcal{S}}$ (left) and $\mathcal{S}$ (right) when $d = 3$. Both families partition $\mathbb{R}^3$ in eight 3-dimensional sets (the transparent volumes), twelve 2-dimensional sets (the green faces), six 1-dimensional sets (the blue lines) and one 0-dimensional set (the black point in the center). The figure only shows slices of the sets and we have also intentionally added space between the sets for illustrative purposes.

Note, $\mathbb{R}_+^d = T(\mathbb{R}^d)$ can be partitioned as $\partial\hat{S} = \left\{\hat{S}_{(J,\emptyset,I\setminus J)}\colon J \subseteq I\right\}$ and it follows that the image of a set $S_{(I_+,I_-,I_0)} \in \mathcal{S}$ under T is exactly

$$T\left(S_{(I_+,I_-,I_0)}\right) = \hat{S}_{(I_+,\emptyset,I_0\cup I_-)} \in \partial\hat{S}$$

Thus, $T\left(S_{(I_+,I_-,I_0)}\right) = T\left(S_{(J_+,J_-,J_0)}\right)$ whenever $I_+ = J_+$ and T reduces to the affine map $x \mapsto Ax + b$ on the closure $\bar{S}_{(I,\emptyset,\emptyset)}$. Since $\dim\left(T\left(S_{(I_+,I_-,I_0)}\right)\right) \le \dim\left(S_{(I_+,I_-,I_0)}\right)$ it is clear that T has contracting properties. If ω is the preimage under the affine map of a set $\hat{\omega} \subset \hat{S}_{(J,\emptyset,I\setminus J)}$ then

$$T^{-1}(\hat{\omega}) = \left\{x - \sum_{i\in I\setminus} \alpha_i a_i^* : \alpha_i \ge 0, x \in \omega \cap S_{(J,\emptyset,I\setminus J)}\right\}$$

Examples of preimages can be seen in Figure 2.

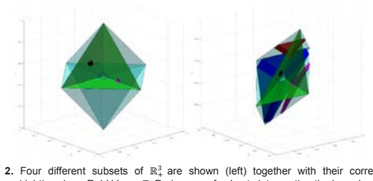

Figure 2. Four different subsets of $\mathbb{R}_+^3$ are shown (left) together with their corresponding preimages (right) under a ReLU layer T. Preimages of subsets intersecting the boundary $\partial\mathbb{R}_+^3$ will be spanned by a subset of the dual basis vectors.

Decision Boundaries

A shallow network has the form $F(x) = L \circ T(x)$ where $L\colon \mathbb{R}^d \to \mathbb{R}$ is an affine functional with a hyperplane $\hat{P}$ as its kernel. The decision boundary of F can be expressed as

$$\Gamma = T^{-1}\left(\hat{P} \cap \mathbb{R}_+^d\right) = \cup_{S_I\in\partial\hat{S}} T^{-1}\left(\hat{P} \cap \hat{S}_I\right)$$

Using the structure of T, we can expand each set in the union as

$$T^{-1}\left(\hat{P} \cap \hat{S}_{(J,\emptyset,I\setminus J)}\right) = \{x - \textstyle\sum_{i\in I\setminus} \alpha_i a_i^*\colon \alpha_i \ge 0, x \in P \cap S_{(J,\emptyset,I\setminus J)}\}$$

where P is the preimage of $\hat{P}$ under the affine map $x \mapsto Ax + b$. Thus, each non-empty intersection $\hat{P} \cap \hat{S}_{(J,\emptyset,I\setminus J)}$ will generate a unique linear piece of Γ spanned by a subset of the dual vectors. Moreover, Γ is completely determined by the preimages $T^{-1}\left(\hat{P} \cap \hat{S}_{(I\setminus\{i\},\emptyset,\{i\})}\right)$, $i \in I$, the intersections of $\hat{P}$ with the $d-1$ dimensional faces in $\partial\hat{S}$. The remaining pieces are essentially linear transitions between these parts. If n is a unit normal to P, then the signs of $n \cdot a_i^*$ indicate how Γ curves since the dual vectors are tangents to the pieces $T^{-1}\left(\hat{P} \cap \hat{S}_{(I\setminus\{i\},\emptyset,\{i\})}\right)$ and n is normal to the central piece $P \cap S_{(I,\emptyset,\emptyset)}$ to which all other pieces are connected. If the hyperplane $\hat{P}$ is in general position, that is, not parallel with any of the standard coordinate axes and $0 \notin \hat{P}$ then there are $t_i \in \mathbb{R}\setminus\{0\}$ such that $t_i e_i \in \hat{P}$ for each $i \in I$. It turns out that $\mathrm{sgn}(t_i) = \mathrm{sgn}(n \cdot a_i^*)$, thus the values $\{t_i\colon i \in I\}$ determine how Γ curves.

We show that the number of linear pieces of Γ is $2^d - 2^m$ where $m = |\{i \in I\colon t_i < 0\colon\}|$. Further, we show that for a shallow ReLU network $F\colon \mathbb{R}^d \to \mathbb{R}$ each possible decision boundary can be obtained by applying an invertible affine map to one of d canonical decision boundaries. Hence, it suffices to understand the properties of these canonical decision boundaries. Figure 3 shows the canonical decision boundaries when $d = 3$.

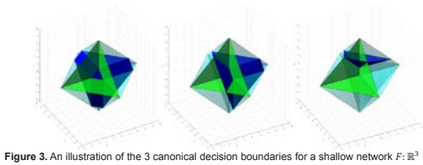

Figure 3. An illustration of the 3 canonical decision boundaries for a shallow network $F\colon \mathbb{R}^3 \to \mathbb{R}$.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A Stochastic Runge-Kutta Optimization Algorithm

Runge--Kutta--Chebyshev (RKC) methods are used to solve numerical differential equations. They have the advantage of being explicit methods with large stability regions. We propose a stochastic optimization scheme for machine learning problems based on the Runge-Kutta-Chebshev methods.

AI MATH | Page 28 B
Williamson, Måns
Lund university
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A Stochastic Runge-Kutta Optimization Algorithm

Måns Williamson
Centre for Mathematical Sciences

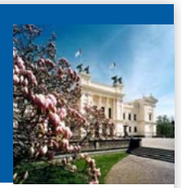

Runge–Kutta–Chebyshev methods

Runge–Kutta–Chebyshev (RKC) methods are methods used to solve numerical differential equations. They have the advantage of being explicit methods with large stability regions.

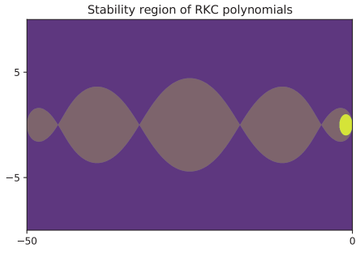

In the plot above we see the stability region of an RKC method with 5 stages and that of the explicit Euler scheme (in bright yellow).

Gradient flow & optimization

We can view the gradient descent algorithm as the explicit Euler scheme applied to the gradient flow equation

$$\dot{w} = -\nabla F(w).$$

As we saw above, the explicit Euler scheme (and thus the gradient descent) has a small stability region which puts a severe stepsize restricition on it.

Stochastic Runge–Kutta–Chebyshev descent

We here present a stochastic optimization algorithm that make use of the RKC methods for minimizing a cost function F:

- Choose an initial iterate w_1 and a sequence of jointly independent random variables $\{\xi_k\}$. For $k = 1, 2, \ldots$
- Set $w_{k0} \leftarrow w_k$ and recieve a stochastic approximation $g(\xi_k, \cdot)$ to $\nabla F(\cdot)$.
- Set $w_{k1} = w_{k0} + \tilde{\mu}_1 \alpha_k g(\xi_k, w_{k0})$.
  For $j = 2, \ldots, s$.
  $-$Set $w_{kj} = \mu_j w_{k,j-1} + \nu_j w_{k,j-2} + \tilde{\mu}_j \alpha_k g(\xi_k, w_{k,j-1})$.
- Set the new iterate as $w_{k+1} \leftarrow w_{ks}$.

Under various standard assumptions (such as strong convexity of the objective function F) we can show that the sequence $\{w_k\}_{k\geq 1}$ genereteted by the SRKCD algorithm converges sub-linearly in expectation. Under the assumption that F is twice differentiable we can show that the algorithm converges in expectation to a stationary point in the non-convex case:

$$\lim_{k\to\infty} \mathbb{E}\|\nabla F(w_k)\|^2 = 0.$$

Numerical experiments

Below we see the results from testing the SRKCD-scheme with 2 stages on a VGG-network using the Cifar-10 dataset.

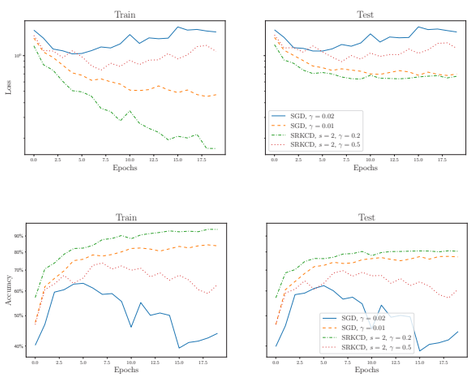

Poster made for WASP winter conference 2022. (Joint work with Tony Stillfjord.)

Regularised Weights in Statistical Models: A General Strategy for Bias Reduction and Increased Stability in Overparameterised Settings

Two challenging aspects of machine learning are label contamination in training data in supervised classification tasks and bias reduction in classical regularisation settings. Our research focuses on a general strategy to non-interactively deal with both problems by expanding the loss function with newly introduced weights. In the first article, we focus on reducing the impact of contaminated labels in training data by localising incorrect data points and reducing their contribution to the loss function. In the second article, we focus on reducing the added bias introduced by classical regularisation methods, like Lasso and Ridge, in a linear regression setting. By doing this, we can, under certain circumstances, keep key properties from the original regularisation penalty and reduce the bias giving us consistent estimators. This leads to the regularisation methods ”entropy weighted Lasso” (EWL) and ”entropy weighted Ridge” (EWR).

Regularised Weights in Statistical Models

A General Strategy for Bias Reduction and Increased Stability in Overparameterised Settings

Olof Zetterqvist
Chalmers University of Technology and University of Gothenburg
Olofze@chalmers.se

CHALMERS UNIVERSITY OF TECHNOLOGY

Introduction

Two challenging aspects of machine learning are label contamination in training data in supervised classification tasks and bias reduction in classical regularisation settings. Our research focuses on a general strategy to non-interactively deal with both problems. The material and experiments are distributed as follows:

1. (Article 1) Reduce the impact of contaminated labels in training data by localising incorrect data points and reducing their contribution to the loss function in a deep convolution neural network setting.
2. (Article 2) Reduce the added bias introduced by classical regularisation methods, like Lasso and Ridge, in a linear regression setting. This leads to the regularisation methods "entropy weighted Lasso" (EWL) and "entropy weighted Ridge" (EWR).

Our approach

Our approach is to expand the loss function with more parameters ω that can be considered weights of different terms. In the presence of label noise, the weights can be put on the data loss terms, and for bias reduction, they can be put on the regularisation terms. To find the "optimal" weight setting, we increase the minimisation task to also include the weights ω with an extra regularisation term $\tilde{g}(\omega) = \sum_i (\omega_i \log(\omega_i) - \omega_i + 1)$

Using weights to find contaminated labels

$$\tilde{\theta}, \tilde{\omega} = \underset{\theta, \omega \,:\, \bar{w}_c = \rho_c, \forall c}{\arg\min} \sum_i \omega_i \ell(X_i, Y_i; \theta) + \lambda g(\theta) + \alpha \tilde{g}(\omega)$$

$$\tilde{\theta} = \underset{\theta}{\arg\min} -\alpha \sum_k \rho_k |C_k| \log \left(\sum_{i \in C_k} e^{-\frac{\ell(x_i, y_i; \theta)}{\alpha}} \right) + \lambda g(\theta)$$

Using weights to reduce bias

$$\tilde{\theta}, \tilde{\omega} = \underset{\theta, \omega}{\arg\min} \frac{1}{2} ||Y - X\theta||_2^2 + \lambda \sum_i \omega_i g(\theta_i) + \gamma \tilde{g}(\omega)$$

$$\tilde{\theta} = \underset{\theta}{\arg\min} \frac{1}{2} ||Y - X\theta||_2^2 + \gamma \sum_i (1 - e^{-\frac{\lambda}{\gamma} g(\theta_i)})$$

Results (contaminated labels)

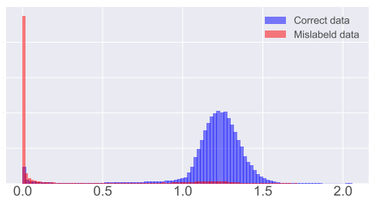

Figure 1: Histograms of the observation weight distributions of both correct and mislabelled training data.

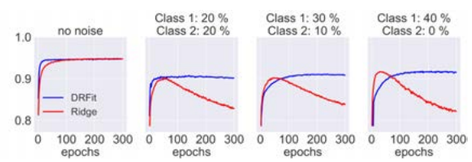

Figure 2: Validation accuracy during training in four different noise settings.

Results (bias reduction)

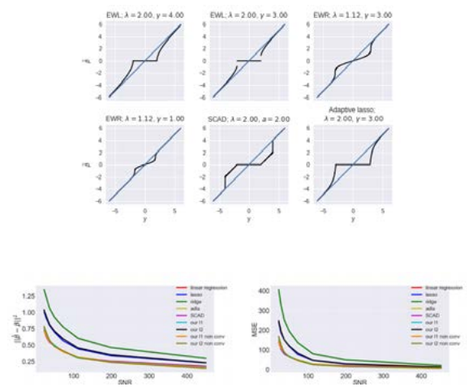

Figure 3: The average L_2 distance between the estimated parameters $\hat{\beta}$ and the true parameters β (**left**) and the mean squared error of predictions on test data (**right**) over 100 runs as functions of the signal to noise ratio (SNR) for nine models on uncorrelated covariates.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

WASP WINTER CONFERENCE 2022

AI MLX

Likelihood-free Out-of-Distribution Detection with Invertible Generative Models

Likelihood of generative models has been used traditionally as a score to detect atypical (Out-of-Distribution, OOD) inputs. However, several recent studies have found this approach to be highly unreliable, even with invertible generative models, where computing the likelihood is feasible. In this paper, we present a different framework for generative model--based OOD detection that employs the model in constructing a new representation space, instead of using it directly in computing typicality scores, where it is emphasized that the score function should be interpretable as the similarity between the input and training data in the new space. In practice, with a focus on invertible models, we propose to extract low-dimensional features (statistics) based on the model encoder and complexity of input images, and then use a One-Class SVM to score the data. Contrary to recently proposed OOD detection methods for generative models, our method does not require computing likelihood values. Consequently, it is much faster when using invertible models with iteratively approximated likelihood (e.g. iResNet), while it still has a performance competitive with other related methods.

Likelihood-free Out-of-Distribution Detection with Invertible Generative Models

Amirhossein Ahmadian and Fredrik Lindsten
Department of Computer and Information Science, Linköping University

WASP Winter Conference 2022 / Published at IJCAI 2021

Introduction

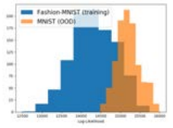

iResNet loglikelihoods on two datasets
In-distribution: FashionMNIST (blue)
OOD: MNIST (orange)

- *Out-of-Distribution (OOD) Detection* : **detecting inputs that are considerably different from the data on which a machine learning model is trained** (anomalies, novelties). Generative models can be used to solve this in a completely unsupervised setting.
- **Generative models may assign much higher likelihoods to OOD test data than in-distribution data**. Also, **likelihood can be too costly to compute.**
- We propose a method to detect OOD inputs using **pretrained invertible generative** models that **does not need to compute the model likelihood** (the log-determinant term). It is based on fast computation of very **low-dimensional statistics** (features) using the invertible model and an **image complexity measure**, and employing a **One-class SVM (OSVM)** in this new feature space.

OOD Detection: Similarity Score and Null Hypothesis Testing

$$u(x) = \sum_{x_i \in D_{tr}} \phi_i \exp\left(\frac{-\|T(x) - T(x_i)\|^2}{\sigma_0^2}\right) - \rho$$

OSVM is trained on in-distribution data in the space of statistics T, to obtain {Φ,ρ}

- Idea: OOD Detection should be based on measuring a **similarity/distance between the input and in-distribution (training) data**. We use a one-class SVM (**OSVM**) to define **a score function** measuring this similarity of interest.
- The representation space is particularly important. **Features called 'statistics' are extracted from the input (image) and the generative model state**. Hence, **the generative model is involved indirectly through the statistics**, and not its likelihood.
- Classical **null hypothesis testing** is used to adjust the threshold on similarity, given the training data and a significance level.
- The idea of using statistics with OSVM is also in **Density of States Estimation (DoS)** [Morningstar et al,2021]. However, we use different statistics that are faster to compute, while having a competitive performance. Our theoretical motivation is also different.

OOD Statistics (Features) for Invertible Models

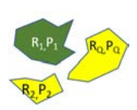

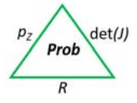

- *Simple but neglected fact*: the **probability** of a class under a generative model depends both on the model **likelihood** and the **region** (particularly the volume) corresponding to the class.
- With **invertible models**, the probability is a function of: **prior density, determinant of encoder Jacobian, integration region**. Some are costly or impossible to obtain, but we may use *any function related to them* as our *statistics*.
- We replace the determinant term, which can be costly to compute, by the **average of column-sum of the encoder Jacobian matrix**, that is another function of Jacobian matrix.
- Heuristic: more **complex inputs** are likely to come from a region with high volume. Assuming image data, we use the **compression size of image using FLIF format** as a proxy for the **region volume**.

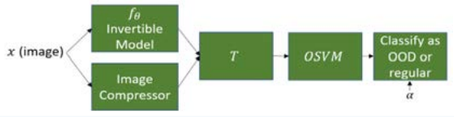

Performance Results and Comparison

- Evaluated on some popular OOD detection tasks such as *FashionMNIST vs. MNIST* and *CIFAR-10 vs. SVHN* with the models *Glow*, *iResNet*, and *ResFlow*, and compared to 4 other methods in terms of **Area Under ROC curve (AUROC)**
- Two combinations of the following statistics:

$$T_1 = log\ p_Z(f(x)), T_2 = C_{FLIF}(x),\ T_3 = \text{mean}\left|\sum_i J_{ij}\right|$$

where *f(.)* is the encoder of invertible model, and *J* is its Jacobian matrix.
- Our method is **competitive with the other ones overall**. It **ranks first in 5 cases**, although it is outperformed by another method in 2 cases. It is better than S-Score [Serra et al,2020] in most cases, which also uses FLIF image compression size complexity measure.
- Most near OODs, e.g. CIFAR-10 vs. CIFAR-100 are not detected well by any of the methods

Model/Data	T_1,T_2	T_1,T_3	DoS	S-Score	Simple LL
iResNet trained on MNIST					
OOD: FashionMNIST	0.99	0.99	0.99	0.97	0.99
iResNet trained on FashionMNIST					
OOD: MNIST	**0.96**	0.89	0.88	0.95	0.07
OOD: Vertical Flipped	0.62	0.60	0.63	0.66	0.55
ResFlow trained on CIFAR10					
OOD: SVHN	**0.96**	0.92	0.94	0.89	0.10
OOD: Vertical Flipped	0.50	0.54	0.52	0.54	0.51
Glow trained on CIFAR10					
OOD: SVHN	0.96	0.91	0.95	0.88	0.09
OOD: CIFAR100	0.57	0.56	0.57	0.49	0.52

Computational Cost Advantage

Model/Data	Log-likelihood Approximation	T_3 statistic
iResNet - MNIST	209ms	3ms
ResFlow – CIFAR10	378ms	20ms
GLOW-CIFAR10	-	18

T_2 statistic: 13ms per image
(Linux core-i7 machine)

- All the compared methods **are dependent on model likelihood** values. In some types of invertible models such as iResNet and ResFlow, likelihood is quite costly to compute since the Jacobian determinant term needs iterative approximation.
- **None of the statistics we suggest depend on explicit evaluation of likelihood** (Jacobian determinant). Consequently, our method can be about **10 times faster at test time**. T_2 comes from a fast image compression algorithm and T_3 from a typical automatic differentiation.
- Considering the close AUROC performance with the other methods, the primary advantage of our method is its speed.

Global Explanation by Characteristic Rules Extraction

Current techniques for explaining black-box predictions rarely produce explanations that generalize beyond the explained instances and hence do not allow for verification or prediction. A method for generating global explanatory rules by aggregating multiple local explanations is proposed. The generated rules can be used to understand how the black-box model operates in general and also emulate the model. The proposed method is applied to several different explanation techniques, individually and in combination. Experimental results show that the method produces high fidelity rules that are often as accurate as, and sometimes even more accurate than, the underlying black-box model.

AI MLX | Page 31 B
Alkhatib, Amr
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Global Explanation by Characteristic Rules Extraction

Amr Alkhatib and Henrik Boström

KTH Royal Institute of Technology
School of Electrical Engineering and Computer Science (EECS)

Abstract

Current techniques for explaining black-box predictions rarely produce explanations that generalize beyond the explained instances and hence do not allow for verification or prediction. A method for generating global explanatory rules by aggregating multiple local explanations is proposed. The generated rules can be used to understand how the black-box model operates in general and also emulate the model. The proposed method is applied to several different explanation techniques, individually and in combination. Experimental results show that the method produces high fidelity rules that are often as accurate as, and sometimes even more accurate than, the underlying black-box model.

Introduction

Techniques for explaining black box models can be classified into local and global methods. A local explanation method explains a single prediction made by a model, while global interpretations provide an understanding of how the model behaves in general and which features are globally important.
The local explanations are not empirically verifiable and cannot be easily used to understand the general behavior of the underlying model. We propose a method for generating explanations in the form of general rules, by aggregating multiple specific explanations. The rules can be combined to explain and emulate the underlying black-box model.

From Local Explanations to Global Rules

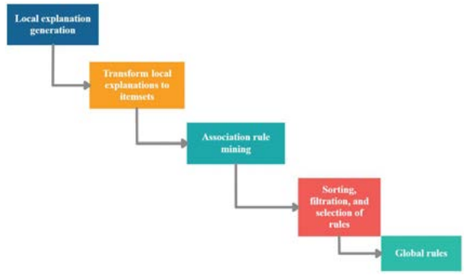

Discriminative vs. Characteristic Rules

The discriminative rules are a way to distinguish the given class from other classes using the minimum possible number of features, while a characteristic rule is the conjunction of all features that are common to all instances in the class.
Extracting the characteristic explanations to a representative subset of a class of instances provides a good approximation to that class's true set of global explanations.

The Accuracy of The Rules

Dataset	XGBoost Model*	Characteristic Explanatory Rules**	The Rules' Fidelity***
Adult	0.92	0.84	0.90
Compass	0.83	0.75	0.81
Spambase	0.99	0.95	0.95
Blood-transf.	0.71	0.69	0.93
German-credit	0.81	0.78	0.83

The explanation rules' performance in terms of area under ROC curve
*The accuracy of the black-box model.
**The accuracy of the rules.
***The accuracy of the rules in predicting the XGBoost model.

Conclusion and Future Work

This work proposes a method to aggregate local explanations and extract characteristic global explanations that we can test and measure their fidelity to the underlying model and apply to new data instances. We also show that the characteristic global explanations have high fidelity and can be more accurate than the black-box model. The method was used to compare different local explanation techniques and also to combine them for more accurate global explanations. An interesting direction for future work is to include the relative importance of each feature in the local explanations to compute a global explanation. It could also be interesting to use conformal prediction frameworks with explanation techniques to measure and quantify the provided explanations' good.

References

1. Scott M Lundberg and Su-In Lee. "A Unified Approach to Interpreting Model Predictions". In: Advances in Neural Information Processing Systems 30. Ed. by I. Guyon et al. Curran Associates, Inc., 2017, pp. 4765–4774. url: http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf.
2. Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. ""Why Should I Trust You?": Explaining the Predictions of Any Classifier". In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016. 2016, pp. 1135–1144.
3. Rakesh Agrawal and Ramakrishnan Srikant. "Fast Algorithms for Mining Association Rules in Large Databases". In: Proceedings of the 20th International Conference on Very Large Data Bases. VLDB '94. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1994, pp. 487–499. isbn: 155860-153-8. url: http://dl.acm.org/citation.cfm?id=645920.672836.

Mental well-being awareness for Human-Robot Interaction

Psychological assistive robots are the next generation of service robots due to the continuous increase of mental illnesses in modern societies. In order to provide the best human-robot experience, the robot should be able to perceive and predict future human behaviors. Therefore, this research stands out for the study of Machine Learning methodologies for learning people's behavior and conduct during daily life. The hypothesis is that the human's state (a latent variable describing the particular human state) is a powerful feature for predicting future behaviors, and a central point for designing a better Human-Robot Interaction.
In this way, we intend to delve into this problem by finding behavioral patterns (ways of moving, daily life activities, physical reactions, etc.) induced by different human states and then learn the best robot interaction.

AI MLX | Page 32 B
Almeida, Tiago
Örebro University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Mental well-being awareness for Human-Robot Interaction

Tiago Almeida, Örebro University
Center for Applied Autonomous Sensor Systems (AASS)

Motivation & Research Goals

Psychological assistive robots are the next generation of service robots due to the continuous increase of mental illnesses in modern societies. In order to provide the best human-robot experience, the robot should be able to **perceive** and **predict future human behaviors**. Therefore, this research stands out for the study of **Machine Learning** methodologies for learning people's behavior and conduct during daily life. The hypothesis is that the human's state (a latent variable describing the particular human state) is a powerful feature for predicting future behaviors, and a central point for designing a better **Human-Robot Interaction**. In this way, we intend to delve into this problem by finding behavioral patterns (ways of moving, daily life activities, physical reactions, etc.) induced by different human states and then learn the best robot interaction.

Overview

Mental well-being

- happiness
- sadness
- neutral
- angriness
- relaxation
- stress

Awareness

- perceive the mental well-being [1]
- find behavioral patterns that indicate the mental well-being
- predict behaviors conditioned on different mental states
- use evidences from Psychology

Human-Robot Interaction

- how the robot should navigate in a human-centred environment [2, 3]
- how the robot should approach the human [2]
- adaptive behaviors [3]

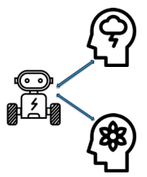

First Ideas

Trajectories conditioned on different mental states

1. Design mental well-being constrained experiments
2. Collect different trajectories of different people under different mental states
3. Assess the mental state of people during the experiments

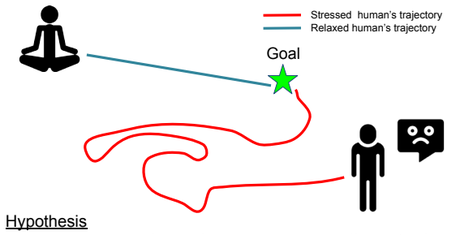

Hypothesis

- There is a correlation between behavioral patterns (motion trajectories) and mental states

Mental well-being aware Human Motion Prediction

1. Observe *(x, y)* position for person *i* during period *t*, so O_t^i in context *C*
2. Estimate mental-well being, M_t^i, *and C*
3. Predict the next coordinates *(x', y')* conditioned on M^i over a period τ

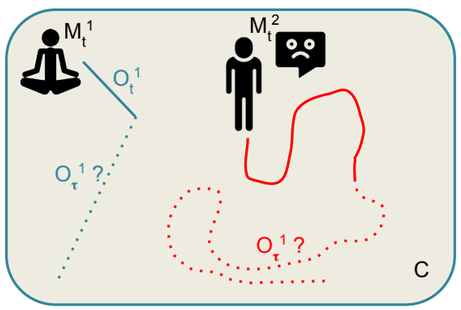

References

1. T. Randhavane, U. Bhattacharya, K. Kapsaskis, K. Gray, A. Bera and D. Manocha, "Learning Perceived Emotion Using Affective and Deep Features for Mental Health Applications," *2019 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, 2019, pp. 395-399, doi: 10.1109/ISMAR-Adjunct.2019.000-2.
2. V. Tolani, S. Bansal, A. Faust and C. Tomlin, "Visual Navigation Among Humans With Optimal Control as a Supervisor," in IEEE Robotics and Automation Letters, vol. 6, no. 2, pp. 2288-2295, April 2021, doi: 10.1109/LRA.2021.3060638.
3. P. Scales, O. Aycard and V. Aubergé, "Studying Navigation as a Form of Interaction: a Design Approach for Social Robot Navigation Methods," 2020 IEEE International Conference on Robotics and Automation (ICRA), 2020, pp. 6965-6972, doi: 10.1109/ICRA40945.2020.9197037.
4. V. Narayanan, B. M. Manoghar, V. Sashank Dorbala, D. Manocha and A. Bera, "ProxEmo: Gait-based Emotion Learning and Multi-view Proxemic Fusion for Socially-Aware Robot Navigation," *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 8200-8207, doi: 10.1109/IROS45743.2020.9340710.

Banerjee, Sourasekhar
Umeå university

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Fed-FiS : A Novel Information-theoretic Federated Feature Selection for Learning Stability

In the era of big data and federated learning, traditional feature selection methods show unacceptable performance for handling heterogeneity when deployed in federated environments. We propose Fed-FiS, an information-theoretic federated feature selection approach to overcome the problem occur due to heterogeneity. Fed-FiS estimates feature-feature mutual information (FFMI) and feature-class mutual information (FCMI) to generate a local feature subset in each user device. Based on federated values across features and classes obtained from each device, the central server ranks each feature and generates a global dominant feature subset. We show that our approach can find stable features subset collaboratively from all local devices. Extensive experiments based on multiple benchmark iid (independent and identically distributed) and non-iid datasets demonstrate that Fed-FiS significantly improves overall performance in comparison to the state-of-the-art methods.

Banerjee, Sourasekhar
Umeå university

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Fed-FiS: A Novel Information-theoretic Federated Feature Selection for Learning Stability

UMEÅ UNIVERSITY
Sourasekhar Banerjee, Umeå University
Dept. of Computing Science
Supervisors: Monowar Bhuyan, Erik Elmroth
sourasb@cs.umu.se, monowar@cs.umu.se, elmroth@cs.umu.se

Autonomous Distributed Systems Lab

Abstract

In the era of big data and federated learning, traditional feature selection methods show unacceptable performance for handling heterogeneity when deployed in federated environments. We propose Fed-FiS, an information-theoretic federated feature selection approach to overcome the problem occur due to heterogeneity. Fed-FiS estimates feature-feature mutual information (FFMI) and feature-class mutual information (FCMI) to generate a local feature subset in each user device. Based on federated values across features and classes obtained from each device, the central server ranks each feature and generates a global dominant feature subset. We show that our approach can find stable features subset collaboratively from all local devices. Extensive experiments based on multiple benchmark iid (independent and identically distributed) and non-iid datasets demonstrate that Fed-FiS significantly improves overall performance in comparison to the state-of-the-art methods. This is the first work on feature selection in a federated learning system to the best of our knowledge.

Introduction

- **Privacy preserving**, **collaborative machine learning** technique.
- Trains **local models** on data samples of each edge device without exchanging raw data.
- Server receives local models from edge devices
- Aggregate the models and produce global model.
- This process continues until the global model converges.

Problems

- Edge devices are normally low performing devices with limited resources. Therefore, computing models from terabytes of data is difficult.
- Sending terabytes of unprocessed data to server is costly and also violate privacy of users.

Objective

- Feature selection is essential and paramount to process such data and uncover useful knowledge for developing low-cost models.

Contributions

- Fed-FiS introduces a local feature subset selection method by using mutual information and clustering.
- We develop a score function based on FCMI and aFFMI for global feature subset selection.
- Fed-FiS finds a most relevant features set from all devices where data is distributed in iid and non-iid manner.

Fed-FiS: Proposed Approach

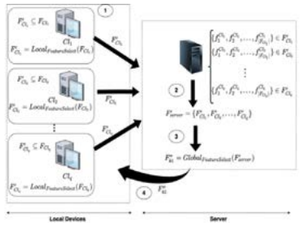

Local feature selection

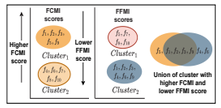

Global feature selection

- **Feature triplet**
  - feature_id,
  - fcmi score
  - affmi score
- **Global score function**
  - Where,

Simulation setup

Parameter(s)	Value	Description
local devices	5	Local feature selection
Server	1	Global feature selection
local device's participation	100%	Full participation of each local device
Cluster Validity Index (CVI)	Silhouette Coefficient	Identify the quality of Cluster
Clustering algorithm	k-means	Cluster with nearest mean
Learning algorithm	2	KNN-5, Federated forest
Performance metrics	1	Accuracy

- **Datasets**
  - NSL-KDD99
  - Anonymized creditcard transaction dataset
- **Data division**

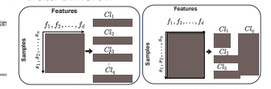

Results and discussions

- Cluster analysis and local feature subset (for Cl_1)

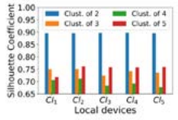

FCMI cluster analysis for NSL-KDD99 iid dataset

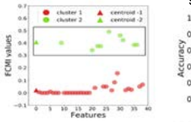

Clusters of features of NSL-KDD99 dataset with respect to the FCMI values.

- Global feature subset selection

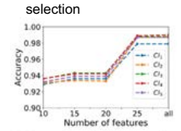

Feature subset for NSL-KDD99 iid dataset

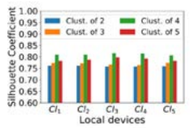

aFFMI cluster analysis for NSL-KDD99 iid dataset

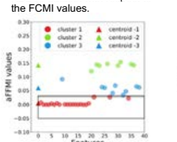

Clusters of features of NSL-KDD99 dataset with respect to the aFFMI values .

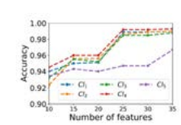

Feature subset for NSL-KDD99 non-iid dataset

- Steady learning ability

Datasets	Fed-FIS		without FS		classical MI based FS	
	Feature subset	Accuracy (%)	Feature set	Accuracy (%)	Feature subset	Accuracy (%)
NSL-KDD99 iid	23	99.3	41	99.5	23	99.24
NSL-KDD99 non-iid	9	89.33	15	54.07	7	92.75
ACC iid	10	99.95	29	99.96	6	99.3
ACC non-iid	9	99.83	14	99.81	10	99.83

- **Inference**
  - For iid datasets the performance of Fed-FiS is equivalent to classical MI based Feature selection method.

Conclusion

- We propose Fed-FiS, a mutual information-based federated feature selection method to select strongly relevant features set for stable and low-cost federated learning.
- Fed-FiS achieved expected model performance with lower number of features set, verified with federated forest algorithm.
- For IID dataset Fed-FiS gives expected performance but for non-IID dataset, size of overlapping feature set is important

Future Work

- More extensive experiments to check robustness of Fed-FiS for different type of dataset, such as, Feature distribution skew, Label distribution skew, Same label different features, Same features different labels, quantity skew or unbalancedness.
- Application specific implementation of Fed-FiS. Ex. Anomaly detection.

References

1. Manikandan, G., Abirami, S.: Feature selection is important: state-of-the-art methods and application domains of feature selection on high-dimensional data. In: Kumar, R., Paiva, S. (eds.) Applications in Ubiquitous Computing. EICC, pp. 177–196. Springer, Cham (2021).
2. Hoque, N., et al.: MIFS-ND: a mutual information-based feature selection method. Expert Syst. Appl. 41(14), 6371–6385 (2014)
3. Liu, G., et al.: Feature selection method based on mutual information and support vector machine. Int. J. Pattern Recogn. Artif. Intell. 35, 2150021 (2021)
4. Zheng, L., et al.: Feature grouping and selection: a graph-based approach. Inf. Sci. 546, 1256– 1272 (2021)
5. Gui, Y.: ADAGES: adaptive aggregation with stability for distributed feature selection. In: Proceedings of the ACM-IMS on Foundations of Data Science Conference, pp. 3–12 (2020)
6. Sohelli, M., et al.: DQPFS: distributed quadratic programming based feature selection for big data. J. Parallel Distrib. Comput. 138, 1–14 (2020)
7. Moran-Fern ´ andez, L., et al.: Centralized vs. distributed feature selection methods based on ´ data complexity measures. Knowl.-Based Syst. 117, 27–45 (2017)
8. Tavallaee, M., et al.: A detailed analysis of the KDD CUP 99 data set. In: IEEE Symposium on Computational Intelligence for Security and Defense Applications, pp. 1–6 (2009)

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Parameter Estimation for Non-linear Differential-Algebraic Equation Models with Unknown Disturbances

Differential-algebraic equations (DAEs) arise naturally as a result of equation-based object-oriented modeling. Such models often contain unknown parameters that have to be identified using measured data. A challenge with the identification of physical systems is the effect of unknown disturbances. If such disturbances are ignored during the identification procedure, one can obtain poor parameter estimates. To the best of the authors' knowledge, there are no general methods successfully dealing with parameter estimation for this type of model. In this work, we propose a simulation-based prediction error method for non-linear DAEs where disturbances are modeled as continuous-time stochastic processes. We assume that the model can be simulated using available DAE solvers. Our method is tested on a simulated pendulum example, which suggests that our method provides consistent parameter estimates.

AI MLX | Page 34 B
Bereza, Robert
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Parameter Estimation for Non-linear DAE Models with Unknown Disturbances

Robert Bereza[1]
Mohamed R.-H. Abdalmoaty[1], Oscar Eriksson[2], David Broman[2] and Håkan Hjalmarsson[1]
[1]KTH, Division of Decision and Control Systems
[2]KTH, Division of Software and Computer Systems

Motivation & Research Goals

Differential-algebraic equations (DAEs) appear naturally when modeling many physical systems, such as mechanical or electrical systems. They are also the underlying model in high-level modeling languages such as Modelica or Simscape. Such models can contain unknown parameters that have to be **estimated using experimental data**. However, disturbances present in real-life systems makes such estimation difficult, as neglecting the influence of such disturbances can lead to biased parameter estimates. Therefore, the goal of this research is to develop **computationally tractable** parameter estimation methods for **nonlinear DAEs**. These methods should model **process disturbances** and take their effect into account to obtain **consistent parameter estimates**.

Differential-Algebraic Equations

- Differential equations with algebraic constraints
- **Example:**

$$\ddot{x}_1 + \theta x_2 = w(t)$$
$$x_1^2 = u(t)$$

with disturbance $w(t)$, control input $u(t)$, and parameter θ

- **General form:**

$$F(\dot{x}(t), x(t), u(t), w(t); \theta) = 0$$
$$y(t; \theta) = q(x(t), u(t); \theta) + noise \quad (*)$$

- Appear naturally when modeling many physical systems
- Underlying model in high level component based modeling languages (e.g. Modelica and Simscape)
- Non-trivial to solve and/or re write as ODEs

Problem and Method

- Identify unknown parameters θ using only knowledge of input $u(t)$ and samples of output $\{y(t_1), y(t_2), \ldots, y(t_N)\}$
- **Challenge:** System influenced by unknown disturbances, which can lead to biased parameter estimates
- **Solution:** Minimize difference between measured output and expected value of model output w.r.t. disturbances [2]

$$\theta_o = \underset{\theta}{\operatorname{argmin}} J_N(\theta) = \underset{\theta}{\operatorname{argmin}} \frac{1}{N} \sum_{k=1}^{N} \|y(t_k) - \mathrm{E}_w[y(t_k; \theta)]\|^2$$

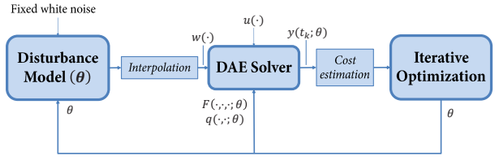

Figure 1. Parameter estimation process

References

1. M. R.-H. Abdalmoaty, O. Eriksson, R. Bereza, D. Broman, and H. Hjalmarsson. (2021). Identification of Non-Linear Differential-Algebraic Equation Models with Process Disturbances. Proceedings The 60th IEEE Conference on Decision and Control (CDC)
2. M. R.-H. Abdalmoaty, "Identification of stochastic nonlinear dynamical models using estimating functions," Ph.D. dissertation, KTH Royal Institute of Technology, 2019.

Simulation Results

Method tested on simulated example of a pendulum

$$m\ddot{x}_1(t) = x_3(t)x_1(t) - k|\dot{x}_1(t)|\dot{x}_1(t) + u(t) + w^2(t)$$
$$m\ddot{x}_2(t) = x_3(t)x_2(t) - k|\dot{x}_2(t)|\dot{x}_2(t) - mg$$
$$L^2 = x_1^2(t) + x_2^2(t)$$

Results when identifying parameter k [1]

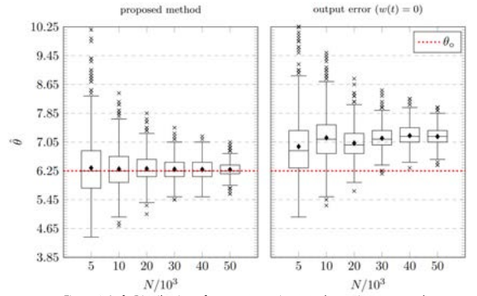

Figure 2. Left: Distribution of parameter estimates when using proposed method. Right: Distribution of parameter estimates when neglecting the effects of the process disturbance during parameter estimation

Stochastic Approximation

- For above results, cost function was approximated using Monte Carlo simulations
- This is time-consuming for large models

$$\nabla_\theta J_N(\theta) = \frac{2}{N} \sum_{k=1}^{N} (y(t_k) - \mathrm{E}_w[y(t_k; \theta)])(-\nabla_\theta \mathrm{E}_w[y(t_k; \theta)])$$

- For **stochastic gradient descent**, we need an **unbiased estimate** of $\nabla_\theta J_N(\theta)$

$$E_w[y(t_k; \theta)] \approx y^{(1)}(t_k; \theta)$$
$$\nabla_\theta E_w[y(t_k; \theta)] \approx \nabla_\theta y^{(2)}(t_k; \theta)$$

where $y^{(1)}(t_k; \theta)$ and $y^{(2)}(t_k; \theta)$ are two independent solutions to $(*)$

- $\nabla_\theta y^{(2)}(t_k; \theta)$ can be obtained while solving DAEs through **forward sensitivity analysis**

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Blöcker, Christopher
Umeå University

Map Equation Centrality

To measure node importance, network scientists employ centrality scores that typically take a microscopic or macroscopic perspective, relying on node features or global network structure. However, traditional centrality measures, such as degree centrality and PageRank, neglect the community structure found in real-world networks. To study node importance based on network flows from a mesoscopic perspective, we exploit the coding principles behind the map equation framework, and derive a community-aware information-theoretic centrality score analytically. Applied to artificial and real-world networks, we demonstrate that our approach enables a more fine-grained differentiation between nodes than node-local or network-global measures, and highlight the role that local network context plays in determining node importance.

AI MLX

Blöcker, Christopher
Umeå University

ZZ-net: A Universal Rotation Equivariant Architecture for 2D Point Clouds

In this work, we investigate a novel neural network architecture for 2D point clouds, guaranteeing rotation equivariance and invariance to permutations of the points.

AI MLX | Page 36 B
Bökman, Georg
Chalmers
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

CHALMERS UNIVERSITY OF TECHNOLOGY

ZZ-net: A Universal Rotation Equivariant Architecture for 2D Point Clouds

There has been much recent work on group equivariance and group invariance of neural networks. Refer to the figure on the right for an example where the group is the rotation group in 2D. Equivariance means that when the input of the network is acted on by a group element, then the output is acted on by the same group element. Invariance means that the output of the network stays constant when the input is acted on by the group.

In this work, we investigate a novel neural network architecture for 2D point clouds, guaranteeing rotation equivariance and invariance to permutations of the points. Investigations have already been done on the 3D case and the general nD case. We find here that the ability to use complex numbers as representations for both the points and the rotations, makes the 2D case stand out and renders it possible for us to find a universality result for neural network architectures that we can't easily generalize to higher dimensions.

There are three main ideas in the paper. First, we discuss how rotation equivariant, permutation invariant functions can be decomposed into a sum of evaluations of a rotation invariant function – see the "Flavour of the results" to the right. Second, we explain a way to approximate this decomposition with neural networks and prove the approximation to be universal. Third, we describe how to apply our framework to the case of an input of correspondences between two point clouds. To illustrate the last idea, we perform experiments on the estimation of essential matrices in stereo vision and find that our framework outperforms state of the art methods when the test data contains rotations unseen in the training data.

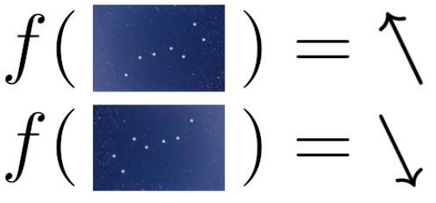

A simple example of rotation equivariance. The illustration shows the task of determining the direction to the North Star given other stars in the night sky. The input is the set of locations of the visible stars in some 2D coordinate frame. Rotation equivariance of the point cloud processor f means here that the determined direction should rotate if the night sky (or the observer) rotates.

Image modified from https://commons.wikimedia.org/wiki/File:Big_Dipper_20210116.jpg, used under Creative Commons Attribution-ShareAlike 4.0 International license. Stars in main constellation brightened.

Flavour of the results

Let Z denote a 2D point cloud with m points, viewed as a vector in $\mathbb{C}^m$. Denote the index set $\{0, 1, 2, \ldots, m-1\}$ by $[m]$ and let τ_i be the transposition of i and 0. Further, let $\mathrm{Stab}(0)$ be the stabilizer of 0, i.e. the group of all permutations of the m indices that fix 0.

Theorem 2

The set of functions on the form

$$f(Z) = \sum_{i \in [m]} \gamma(\tau_i(Z)) Z_i, \tag{1}$$

where γ is an arbitrary continuous, rotation invariant and $\mathrm{Stab}(0)$-invariant function, is dense in the set of continuous rotation equivariant and permutation invariant functions from $\mathbb{C}^m$ to $\mathbb{C}$.

Reference:
Bökman, Georg, Fredrik Kahl, and Axel Flinth. November 2021. 'ZZ-Net: A Universal Rotation Equivariant Architecture for 2D Point Clouds' http://arxiv.org/abs/2111.15341.

Georg Bökman
bokman@chalmers.se
Computer Vision
Department of Electrical Engineering
Chalmers University of Technology

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Automatic Static transformation of Probabilistic Programs

Probabilistic program languages (PPLs) provide tools to write a model and do statistical inference on the model ideally without considering the internals of the inference algorithm. While simulation is the strong aspect of the Monte Carlo methods in PPLs to generate samples, it needs to be done in a smart way such that the variance of the sampler reduces in a computationally feasible time. Especially, the models, having high number of random variables, require manual transformation of the model to get an efficient sampler. The aim of this research is to automatically analyse and transform the model given by the user to a more efficient equivalent representation using analytical relations between random variables in compile time.

AI MLX | Page 37 B
Caylak, Gizem
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Automatic Static Transformations of Probabilistic Programs

Gizem Caylak, KTH Royal Institute of Technology
Department of Computer Science
Main Advisor: David Broman

Abstract

Probabilistic program languages (PPLs) provide tools to write models and do statistical inference on the model ideally without considering the internals of the inference algorithm. While simulation is an important aspect of Monte Carlo methods in PPLs to generate samples, it needs to be done in a effective way that the variance of the sampler reduces in a computationally feasible time. Especially, the problem is that models, which have high number of random variables, require manual transformations to get an efficient sampler. The aim of this research is to **automatically analyze and transform** the model given by the user to a more **efficient** equivalent representation. We use analytical relations between random variables **at compile-time** in transformations.

Research Problem

PPLs provide an expressive interface to represent probabilistic models ideally without dealing with the inference. However, making the sampler efficient requires manual transformation of the model. Run-time algorithms, such as Delayed Sampling [1], are suggested to do the transformation automatically. To reduce the run-time overhead, we propose an algorithm to do the transformation **automatically at compile time.**

Method

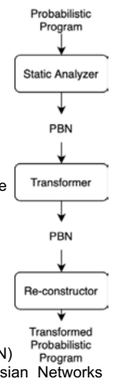

Static Analyzer
Creates a Programmatic Bayesian Network (PBN) from the probabilistic program. Vertices are either random variables or code blocks and edges represent the dependency between vertices.

Transformer
Takes the PBN as input and uses conjugate prior relationships between random variables to derive the posterior. We guarantee that the transformation is conservative, always producing a correct program.

Re-constructor
Takes the transformed PBN and reconstructs the probabilistic program.

Main Contributions
- We extend the concept of Bayesian Networks (BN) new kind of graph called Programmatic Bayesian Networks (PBN). A PBN encapsulates random variables and the code structures in a probabilistic program as well as reconstructing a probabilistic program from a PBN.
- We transform the PBN based on analytical relations, such as conjugate prior relations, between the model parameters.
- All steps are done **automatically at compile time.**
- We implement this method in the meta language framework Miking and demonstrate the efficiency of our algorithm on non-trivial models such as LDA

References

1. Murray, L., Lundén, D., Kudlicka, J., Broman, D. & Schön, T. (2018). Delayed Sampling and Automatic Rao-Blackwellization of Probabilistic Programs. *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, in *Proceedings of Machine Learning Research* 84:1037-1046

Transformations

We use conjugate prior relations between random variables to derive the posterior hyperparameters. This gives a closed-form expression for the posterior.

$$P(\theta \mid Y) = \frac{P(Y \mid \theta) P(\theta)}{P(Y)}$$

Posterior — Likelihood — Prior

Example: A simple coin toss mode

$$P(x) \sim Beta(5,5)$$

$$P(Y \mid x) \sim Bernoulli(x)$$

$$P(x \mid Y = true) \sim ?$$

```
1 let x = assume (Beta 5.0 5.0) in
2 let obs = true in
3 observe obs (Bernoulli x);
4 x
```

```
1 let x = assume (Beta 6.0 5.0) in
2 x
```

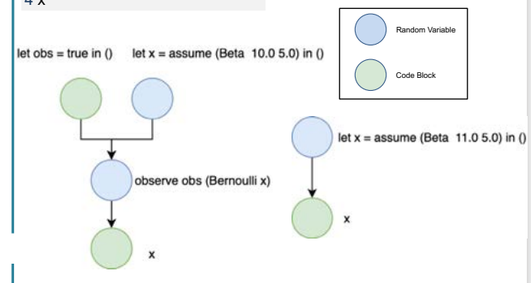

Limitations and work in progress
- Since the transformation is conservative, it may not output a more efficient representation, e.g. in **stochastic branching,** it may not capture different relations discovered between random variables.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Dimensionality reduction for Attributed Graphs (on-going work)

Embedding nodes in a large graph using vectors but most solutions rely on optimizing embedded representations. This yields problems; adding new nodes require re-training the embeddings, and training is costly and time-consuming. To mitigate this, we propose a simple, static and lossy way of compressing and representing higher-order connections in a network. We make use of Random Indexing, embedding the nodes through static vectors in a euclidean vector space using aggregation methods, showing a performance increase in comparison to baselines, needing less dimensions for more expressivity.

AI MLX | Page 38 B
Cornell, Filip
Gavagai AB
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Dimensionality reduction for Attributed Graphs (on-going work)

Filip Cornell
KTH Royal Institute of Technology & Gavagai

Introduction

Embedding nodes in a large graph using vectors but most solutions rely on **optimizing** embedded representations. This yields problems:
- Adding new nodes require re-training the embeddings
- Training is costly and time-consuming

We propose:
- A simple, static and lossy way of compressing and representing higher-order connections in a network.
- We make use of **Random Indexing,** embedding the nodes through static vectors in a euclidean vector space

Two methods:
- Linear Graph Neural Network aggregation (**RIA**)
- Build up vectors from random walks (**RI + Random Walks**)

Data

Five benchmark datasets; citation & social networks

Dataset	$\|V\|$	$\|E\|$	$\|F\|$	$\|C\|$
Cora	2,708	5,429	1,433	7
Citeseer	3,312	4,715	3,703	6
Pubmed	19,717	44,338	500	3
Wiki	2,405	17,981	4,973	17
BlogCatalog	5,196	343,486	8,189	6

Datasets: Citation & social networks

Method

Empirical investigation[2]
1. Embed the graphs using our methods
2. Evaluate using linear regression on the embeddings using **node** and **link** prediction

Random hyperparameter search
For all methods, perform a hyperparameter search

Baselines:
- QUINT[1] – recent hashing method baselines
- Attributes only (binary one-hot encodings)

Experiment pipeline

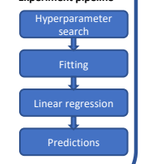

Results

Link prediction
- We beat our baselines in all cases
- Embedding higher order neighborhood is **crucial**

Method/Dataset	Cora		Citeseer		Pubmed		Wiki		BlogCatalog	
	Micro	Macro	Micro	Macro	Micro	Macro	Micro	Macro	Micro	Macro
Feature baseline	.768	.726	.756	.707	.8	.796	.706	.659	.925	.923
QUINT	.804	.768	.766	.713	.876	.875	.797	**.69**	**.968**	**.967**
RIW (ours)	.862	.841	.746	.705	.762	.735	.815	.712	.749	.749
RIA (ours)	**.89**	**.88**	**.798**	**.723**	**.89**	**.887**	**.8**	.669	.966	.966

Node classification experiments. Our method is superior to our baselines in most cases

Method/Dataset	Cora	Citeseer	Pubmed
	with/out	with/out	with/out
Features only	.747	.826	.808
QUINT	.865/.722	.92/.669	.914/.892
RI + Random walks	.72/.696	.719/.628	.891/.868
Aggregation, RI	**.895/.814**	**.968/.809**	**.975/.891**

Link prediction. We beat our baselines in all cases

- Our method needs **less dimensions** to produce better results than other hashing methods -> higher expressiveness

Future work

- Test features' quality when combined with a Graph Neural Network (GNN)
- Can be useful to initialize representations when graphs have **binary** or **no attributes**
- Initial experiments show up to an 80 % reduction of parameters with **slight** performance decrease

Future pipeline

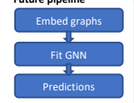

Acknowledgements

This work was partially supported by Wallenberg Autonomous Systems Program (WASP).

1 QUINT: Node embedding using network hashing. arXiv preprint arxiv:2109.04206, 2021.

Doostmohammadi, Ehsan
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Verb Understanding in Video Grounded Language Models

Training a language model only on textual data does not give the model a comprehensive understanding of the world. One way to amend this and make more capable models is to train them on more modalities, for example, images. There is also a vast literature on video grounding, which usually means further training (most of the times transformer-based) language models on videos and their captions. Training on videos is usually more expensive than images and requires 3D convnet encoders, which raises the question, how does it benefit a language model? In this work, we try to investigate this further and see if grounding in videos improves model's understanding of actions (verbs) and objects (nouns), over only text, text and images, and possibly sparsely sampled video frames.

AI MLX | Page 39 B

Doostmohammadi, Ehsan
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Does Grounding in Videos Help with Verb Understanding?

Ehsan Doostmohammadi, Linköping University
Department of Computer and Information Science
Supervisors: Marco Kuhlmann (LiU) and Richard Johansson (Chalmers)

What is Grounding?

Grounding is defined as "Learning language representations from explicit visual associations" (Sileo, 2021), which basically means using other knowledge sources (e.g., images or knowledge graphs) in addition to text for training language models. We expect such models to perform better in some aspects and have a better understanding (?) of the world.

Grounding in Videos; How Does it Help?

Have you seen those perfect pictures in image-caption datasets that are self-contained and obvious? Unfortunately, the world is not that perfect. Can you guess what is being added in the left picture and what they are doing in the right one?

Or what are they doing in this one? Should you plug in before doing a certain task or unplug? In this case, the information is available in text, but that is not always the case.

Ideally, one would create a dataset (maybe similar to Hendricks et al., 2021) and be done with it. But creating datasets is expensive and even more complicated for videos.

Imageability

As it is not currently feasible to create a dataset, we rely on a body of work in psychology about an easy-to-grasp concept called imageability (Bird et al., 2001). Words like "(to) walk" and "(a) wasp" are more imageable than "(to) think" and "jealousy".
We hypothesize that a language model that is grounded in videos will perdict highly imageable words more accurately than less imageable ones and also compared to a model that is not grounded.

References

- Bird, Helen, Sue Franklin, and David Howard. "Age of acquisition and imageability ratings for a large set of words, including verbs and function words." Behavior Research Methods, Instruments, & Computers 33.1 (2001): 73-79.
- Hendricks, Lisa Anne and Aida Nematzadeh. "Probing Image-Language Transformers for Verb Understanding." FINDINGS (2021).
- Miech, Antoine, et al. "Howto100m: Learning a text-video embedding by watching hundred million narrated video clips." Proceedings of the IEEE/CVF International Conference on Computer Vision. 2019.
- Sileo, Damien. "Visual Grounding Strategies for Text-Only Natural Language Processing." Proceedings of the Third Workshop on Beyond Vision and LANguage: inTEgrating Real-world kNowledge (LANTERN). 2021.
- Sun, Chen, et al. "Videobert: A joint model for video and language representation learning." Proceedings of the IEEE/CVF International Conference on Computer Vision. 2019.
- A. -M. Oncescu, et al., "QUERYD: A Video Dataset with High-Quality Text and Audio Narrations," ICASSP 2021 - 2021 ICASSP, 2021, pp. 2265-2269.

Data

We need video data paralleled with text, such as video captions. But such datasets are expensive to create and small in size. Therefore we use **Howto100m** dataset, which is a large-scale dataset of narrated instructional videos (Miech et al., 2019). As this dataset is automatically collected, there is a considerable amount of noise in it, which is partly compensated by the sheer size of it, ∼12 TB.

How to train such a model?

We used a pretrained distil-BERT language model and fine-tuned on parallel video-text data. Random words and video inputs are masked and the model learns to predict them. Additionally, the text and video in some random training samples are not aligned and the model is trained to distinguish them. For more details see **VideoBERT** (Sun et al., 2019).

Preliminary Results

We train the language model as described above and test it on a held-out test set of 5 thousand videos. The words that are available in the Bird dataset are masked for the model to predict them. Only words with high and low imageablity are kept, as interpreting the results on medium imageability would be difficult.

Train/Test	Img.	Acc. (Δ)	V. Acc. (Δ)	N. Acc. (Δ)
T/T	Low	34.3	35.7	24.0
	High	16.8	17.5	16.6
TV/TV	Low	33.7 (-0.6)	35.0 (-0.7)	23.7 (-0.3)
	High	17.7 (0.9)	18.9 (1.4)	17.2 (0.6)
TV/T	Low	34.1 (-0.2)	35.5 (0.2)	24.0 (0.0)
	High	17.1 (0.3)	18.0 (0.5)	16.7 (0.1)

In the first column, T stands for text and V for video. In the first row, Acc. is accuracy, V. is verb, and N. stands for noun. Delta is the difference between that cell and the corresponding cell in the T/T row.

The results show 1.4% increase in the accuracy when we train on videos and text, and test on both. When only testing on text, we see a 0.5% increase for the highly imageable words. The results also show that there is a higher increase for verbs compared to nouns.

What's next?

Train a better model, test on high quality data, such as **QueryD** (Oncescu et al., 2021) (as only 50% of the text is refering to an object/action in the scene in Howto100m), and analyze and interpret the results.
Have any qustion? Don't hesitate to contact me:
ehsan.doostmohammadi@liu.se :)

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning Language-Based Representations for Efficient and Intelligent Acting

We learn compact and reusable control knowledge as representations of a target language. The control knowledge captures the subgoal structure in a planning domain. We can formally prove that the learned representations can be used to solve any instance from the domain efficiently. In this setting, we are restricted to tractable domains. Those are domains, where it is easy to find a suboptimal solution. We address the learning problem using answer set programming. We present some results of the learning and directions for future work.

Learning Language-Based Representations for Efficient and Intelligent Acting

Dominik Drexler, PhD Student, Linköping University

Supervisors: Hector Geffner and Jendrik Seipp

LINKÖPINGS UNIVERSITET

Motivation & Research Goals

- Learn reusable control knowledge for intelligent acting
- Learned target is language representation that
  - is compact and simple
  - represents the subgoal structure
  - generalizes over target class $\mathcal{Q}$ of planning problems
  - can be used to solve any problem $P \in \mathcal{Q}$ efficiently
- Goals in future work: scalability, learn hierarchies, learn from arbitrary inputs

Methods

Classical Planning Problem.

- **Given:** A **planning problem** P models a world and consists of
  - **Domain** D, i.e., **Predicates** over **objects**, **Actions** over predicates
  - **Instance information**, i.e, sets of literals made from ground atoms over predicates and objects:
    - Fully observable **initial state** I
    - **Goal** description G
- **Objective:** find action sequence that leads from I to G

Width.

- **Width** $w(P)$ measures difficulty of problem P
- **Theorem:** If $w(P) \leq k$ then $IW(k)$ algorithm solves P in time $\exp(k)$

Serialization & Policy Sketches.

- **Policy sketch** R_Φ is set of rules of form $C \mapsto E$ over state features Φ
- R_Φ defines subgoals in $P \in \mathcal{Q}$ over common domain
- R_Φ decomposes $P \in \mathcal{Q}$ into subproblems
- **Sketch width** $w_{R_\Phi}(\mathcal{Q})$ is largest width of any subproblem in all $P \in \mathcal{Q}$
- **Serialization** SIW_{R_Φ}: search greedily to closest subgoal and repeat
- **Theorem:** if $w_{R_\Phi}(\mathcal{Q}) \leq k$ then SIW_{R_Φ} solves $P \in \mathcal{Q}$ in time $\exp(k)$

Learning Sketches.

- As combinatorial optimization problem: answer set programming
- Unsupervised from example problems $P_1, \ldots, P_n$
- Input parameter $k \in \mathbb{N}_0$ controls difficulty of subproblems
- Current limitations: tractable domains, scalability, requires PDDL like inputs

Selected Results

Gripper. An agent must move all balls from room a to room b.

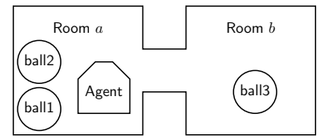

Consider **features** $\Phi = \{g_a, g_b\}$ where

- g_a is number of balls in room a
- g_b is number of balls in room b

The **learned width-2 sketch** R^2 over features $\{g_b\}$ is

$$r = \{\} \mapsto \{g_b\uparrow\}$$

- r: increasing number of balls in room b ($g_b\uparrow$) is good

The **learned width-1 sketch** R^1 over features $\{g_a, g_b\}$ is

$$r_1 = \{\} \mapsto \{g_a\downarrow, g_b?\}$$
$$r_2 = \{\} \mapsto \{g_b\uparrow\}$$

- r_1: decreasing number of balls in room a ($g_a\downarrow$) and arbitrarily changing number of balls in room b ($g_b?$) is good
- r_2: increasing number of balls in room b ($g_b\uparrow$) and keeping number of balls in room a the same (no effect) is good

Experimental Results.

	$w=1$				LAMA		BFWS	
Domain	S	T	AW	MW	S	T	S	T
Blocks-clear (30)	30	3	0.78	1	30	4	30	8
Blocks-on (30)	30	36	1.00	1	30	4	30	52
Childsnack (30)	30	4	0.11	1	9	3	4	10
Delivery (30)	30	1	1.00	1	30	3	30	2
Gripper (30)	30	5	0.50	1	30	3	30	7
Miconic (30)	30	140	0.54	1	30	7	30	28
Reward (30)	30	3	1.00	1	30	3	30	4
Spanner (30)	30	1	0.25	1	0	–	0	–
Visitall (30)	30	1687	0.01	1	29	507	18	113

References

[1] Expressing and Exploiting the Common Subgoal Structure of Classical Planning Domains Using Sketches
Dominik Drexler, Jendrik Seipp and Hector Geffner
KR2021, 3-12 November, 2021, Hanoi, Vietnam

[2] Learning Sketches for Decomposing Planning Problems into Subproblems of Bounded Width
Dominik Drexler, Jendrik Seipp and Hector Geffner
submitted to ICAPS2022, 19-24 June, 2022, Singapore, Singapore

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Ek, Sofia
Uppsala University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning Language-Based Representations for Efficient and Intelligent Acting

We learn compact and reusable control knowledge as representations of a target language. The control knowledge captures the subgoal structure in a planning domain. We can formally prove that the learned representations can be used to solve any instance from the domain efficiently. In this setting, we are restricted to tractable domains. Those are domains, where it is easy to find a suboptimal solution. We address the learning problem using answer set programming. We present some results of the learning and directions for future work.

Ek, Sofia
Uppsala University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

UPPSALA UNIVERSITET

Learning Pareto-Efficient Decisions with Confidence

Sofia Ek, Dave Zachariah, Petre Stoica
Uppsala University, Sweden

Introduction

- Interested in problems where the objective functions are unknown but we have access to data from past outcomes of decisions. The goal is to balance trade-offs between several outcome objectives using Pareto-efficiency.
  - Decision $\mathbf{x}$ is *Pareto-efficient* if it is not strictly dominated by any alternative decision $\mathbf{x}'$.
  - Efficient decisions $\mathcal{X}^{\alpha}$ form a 'Pareto-frontier' $\mathcal{Y}^{\alpha}$ in reward space.

Main features of proposed method:
- Considers the lower tail rewards at any specified level.
- Provides finite-sample coverage guarantees for the rewards even when the data-generating process is unknown.
- Enables the study of decision trade-offs using a Pareto-frontier that takes reward uncertainties into account.

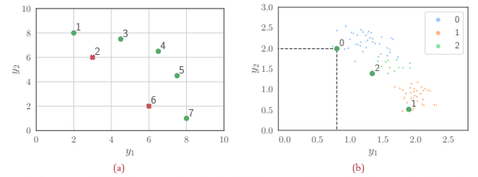

Figure: (a) Illustration of efficient decisions $\mathcal{X}^{\alpha} = \{1, 3, 4, 5, 7\}$ shown as solid green dots. Decisions 2 and 6 (red crosses) are thus strictly dominated. (b) Decision $\mathbf{x} = 0$ and its corresponding outcome rewards $\mathbf{y}_1$ and $\mathbf{y}_2$.

Problem Formulation

We have observed data $\mathcal{D} = \{(\mathbf{x_i}, \mathbf{y_i}, \mathbf{z_i})\}_{i=1}^{n}$, where:
- $\mathbf{x} \in \{0, 1, 2, \dots\}$ are discrete decisions (treatments) in a space $\mathcal{X}$.
- $\mathbf{y}$ are reward vectors (utility/negative loss) in a reward space $\mathcal{Y} \subset \mathbb{R}^{\mathbf{m}}$. The number $\mathbf{m}$ is typically small, the examples here illustrate cases with $\mathbf{m} = 2$ rewards.
- $\mathbf{z}$ are $\mathbf{d}$-dimensional covariate vectors representing decisions taken in different contexts.

Goal: Learn covariate-specific efficient decisions $\mathcal{X}^{\alpha}(\mathbf{z})$ and the corresponding 'Pareto-frontier' $\mathcal{Y}^{\alpha}(\mathbf{z})$ in reward space that *lower bounds* random rewards with a high probability of at least $\mathbf{1} - \boldsymbol{\alpha}$:

$$\mathbb{P}\big(\, \mathbf{y}^{\alpha}(\mathbf{x}, \mathbf{z}) \preceq \mathbf{y} \,|\, \mathbf{x} \,\big) \;\geq\; \mathbf{1} - \boldsymbol{\alpha}, \qquad \forall \mathbf{x} \in \mathcal{X}. \tag{1}$$

If a nontrivial bound can be found, we use it to define efficient decisions that take into account the uncertainty of outcome. We also consider a covariate shift, i.e. the training distribution $\bar{\mathbf{p}}(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is not the same as the testing distribution $\mathbf{p}(\mathbf{x}, \mathbf{y}, \mathbf{z})$.

Method

We want to ensure (1) . Using the union-bound, we have that:

$$\mathbb{P}\big(\, \cup_{\mathbf{k}=1}^{\mathbf{m}} \mathbf{y_k} < \mathbf{y}_{\mathbf{k}}^{\alpha}(\mathbf{x}, \mathbf{z}) \,|\, \mathbf{x} \,\big) \leq \sum_{\mathbf{k}=1}^{\mathbf{m}} \mathbb{P}(\mathbf{y_k} < \mathbf{y}_{\mathbf{k}}^{\alpha}(\mathbf{x}, \mathbf{z}) \,|\, \mathbf{x}).$$

Building upon the 'weighted conformal quantile regression' method [4, 3, 2], we construct $\mathbf{y}_{\mathbf{k}}^{\alpha}$ such that:

$$\mathbb{P}(\mathbf{y_k} < \mathbf{y}_{\mathbf{k}}^{\alpha}(\mathbf{x}, \mathbf{z}) | \mathbf{x}) \leq \frac{\boldsymbol{\alpha}}{\mathbf{m}} \qquad \forall \mathbf{k},$$

given $\mathbf{p}(\mathbf{z}|\mathbf{x})$. This ensures that (1) is satisfied. The conditional *quantile function* of an individual continuous reward $\mathbf{y}$ can be expressed as:

$$\mathbf{q}^{\alpha}(\mathbf{x}, \mathbf{z}) = \inf \left\{ \tilde{\mathbf{y}} : \frac{\boldsymbol{\alpha}}{\mathbf{m}} = \mathbb{P}(\mathbf{y} < \tilde{\mathbf{y}} \,|\, \mathbf{x}, \mathbf{z}) \right\}.$$

This quantity is unknown and $\hat{\mathbf{q}}^{\alpha}(\mathbf{x}, \mathbf{z})$ is estimated from data using *any* flexible machine learning method. Determine an *adjustment* variable $\kappa^{\alpha}(\mathbf{x})$ and form:

$$\mathbf{y}^{\alpha}(\mathbf{x}, \mathbf{z}) := \hat{\mathbf{q}}^{\alpha}(\mathbf{x}, \mathbf{z}) - \kappa^{\alpha}(\mathbf{x}).$$

The adjustment variable is needed to guarantee coverage in finite samples, but should also be weighted to take distributional shift into account

$$\bar{\mathbf{p}}(\mathbf{x}, \mathbf{y}, \mathbf{z}) = \mathbf{p}(\mathbf{z})\mathbf{p}(\mathbf{x}|\mathbf{z})\mathbf{p}(\mathbf{y}|\mathbf{x}, \mathbf{z}), \quad \mathbf{p}(\mathbf{x}, \mathbf{y}, \mathbf{z}) = \mathbf{p}(\mathbf{z})\mathbb{1}(\mathbf{x} = \mathbf{x}^{*})\mathbf{p}(\mathbf{y}|\mathbf{x}, \mathbf{z}),$$

using weights $\mathbf{w}(\mathbf{x}, \mathbf{z})$ that must be estimated when the past decision policy is unknown.

Synthetic Data

One-dimensional covariate $\mathbf{z} \sim \mathcal{N}(\mathbf{60}, \mathbf{100})$. Five decision alternatives $\mathcal{X} = \{\mathbf{0}, \mathbf{1}, \mathbf{2}, \mathbf{3}, \mathbf{4}\}$. The rewards $\mathbf{y}$ are drawn as:

$$\begin{cases} \mathbf{y}_1 &= \mathbf{a}_{\mathbf{x}} + \mathbf{b}_{\mathbf{x}}\mathbf{f}\left(\frac{\mathbf{z}-\mathbf{55}}{\mathbf{7}}\right) + \mathbf{u}_0, \\ \mathbf{y}_2 &= \mathbf{c}_{\mathbf{x}} + \mathbf{d}_{\mathbf{x}}\mathbf{f}\left(\frac{\mathbf{z}-\mathbf{50}}{\mathbf{8}}\right) + \mathbf{u}_1, \end{cases}$$

where $\mathbf{f}(\cdot)$ is the sigmoid function, $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}$ are constants and $\mathbf{u}_0, \mathbf{u}_1$ are noise terms. We create an unbalanced scenario where some of the decisions $\mathbf{x}$ are rare and learn a (two-component) Gaussian mixture model of $\mathbf{p}(\mathbf{z}|\mathbf{x})$ in:

$$\mathbf{w}(\mathbf{x}, \mathbf{z}) = \frac{\mathbb{1}(\mathbf{x} = \mathbf{x}^{*})}{\mathbf{p}(\mathbf{x}|\mathbf{z})} = \frac{\mathbb{1}(\mathbf{x} = \mathbf{x}^{*})}{\mathbf{p}(\mathbf{z}|\mathbf{x})\mathbf{p}(\mathbf{x})} \sum_{\mathbf{x}' \in \mathcal{X}} \mathbf{p}(\mathbf{z}|\mathbf{x}')\mathbf{p}(\mathbf{x}').$$

Results: Synthetic Data

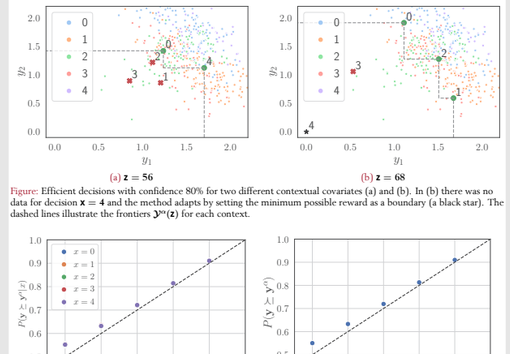

Figure: Efficient decisions with confidence 80% for two different contextual covariates (a) and (b). In (b) there was no data for decision $\mathbf{x} = 4$ and the method adapts by setting the minimum possible reward as a boundary (a black star). The dashed lines illustrate the frontiers $\mathcal{Y}^{\alpha}(\mathbf{z})$ for each context.

Figure: Evaluation of reward boundary $\mathbf{y}^{\alpha}(\mathbf{x}, \mathbf{z})$ using 500 Monte Carlo simulations. (a) Experimental data scenario with known random assignment probabilities $\mathbf{p}(\mathbf{x})$. (b) Observational data scenario where $\mathbf{p}(\mathbf{z}|\mathbf{x})$ is estimated.

Semi-real Data

- Teachers and students randomly assigned into:
  - Small class size: 13 to 17 students per teacher,
  - Regular class size: 22 to 25 students per teacher,
  - Regular-with-aide class size: 22 to 25 students per teacher plus a full-time teacher assistant.
- The study was a *randomized controlled trial*, with interventions assigned by a given random policy $\mathbf{p}(\mathbf{x})$.
- Outcome $\mathbf{y}_1$ is an achievement test score (the sum of standardized math, reading and listening test scores).
- Outcome $\mathbf{y}_2$, is a non-linear generated value [1] and could correspond to a (negative) cost per student.

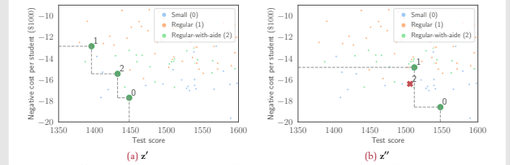

Figure: Frontiers $\mathcal{Y}^{\alpha}(\mathbf{z})$ for student groups with two different sets of covariates (a) and (b) and learned from real data with three different class size interventions. Confidence level 80%.

Conclusions

In this work we extended the concept of Pareto-efficiency to include statistical confidence. We propose a method that learns efficient decisions and a frontier that takes into account the lower tail rewards at any specified confidence level. Such a frontier can quantify trade-offs and provide valuable insight to decision makers due to its finite-sample statistical guarantees.

Acknowledgment

This research was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation.

References

[1] J. L. Hill.
Bayesian nonparametric modeling for causal inference.
Journal of Computational and Graphical Statistics, 20(1):217–240, 2011.

[2] L. Lei and E. J. Candès.
Conformal Inference of Counterfactuals and Individual Treatment Effects.
arXiv preprint arXiv:2006.06138, 2020.

[3] Y. Romano, E. Patterson, and E. Candès.
Conformalized quantile regression.
Advances in Neural Information Processing Systems, 32, 2019.

[4] R. J. Tibshirani, R. Foygel Barber, E. Candès, and A. Ramdas.
Conformal prediction under covariate shift.
Advances in Neural Information Processing Systems, 32, 2019.

sofia.ek@it.uu.se

Generalized Jensen-Shannon Divergence Loss for Learning with Noisy Labels

Prior works have found it beneficial to combine provably noise-robust loss functions e.g., mean absolute error (MAE) with standard categorical loss function e.g. cross entropy (CE) to improve their learnability. Here, we propose to use Jensen-Shannon divergence as a noise-robust loss function and show that it interestingly interpolate between CE and MAE with a controllable mixing parameter. Furthermore, we make a crucial observation that CE exhibit lower consistency around noisy data points. Based on this observation, we adopt a generalized version of the Jensen-Shannon divergence for multiple distributions to encourage consistency around data points. Using this loss function, we show state-of-the-art results on both synthetic (CIFAR), and real-world (e.g., WebVision) noise with varying noise rates.

AI MLX | Page 42 B
Englesson, Erik
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Generalized Jensen-Shannon Divergence Loss for Learning with Noisy Labels

Erik Englesson, Hossein Azizpour

1. Jensen-Shannon Loss Generalizes CE and MAE

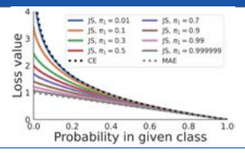

2. Learning with Label-Noise Reduces Consistency (CE)

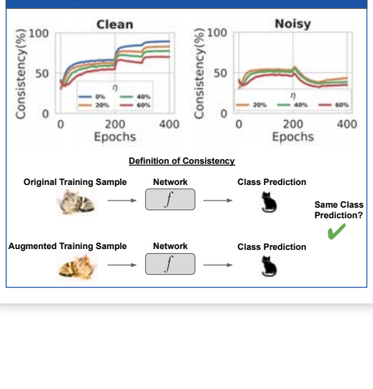

3. Generalized JS loss to Improve Consistency

y = One-hot label
p_1, p_2 = Predictive distributions of augmented versions of the sample

Label-dependent term | Consistency term

$$\mathrm{GJS}(y, p_1, p_2) := \mathrm{JS}_\pi(y, \frac{p_1+p_2}{2}) + (1-\pi)\mathrm{JS}_{\frac{1}{2}}(p_1, p_2)$$

$$\mathrm{JS}_\pi(p, q) := \pi \mathrm{KL}(p \| m) + (1-\pi)\mathrm{KL}(q \| m)$$

$$m := \pi p + (1-\pi) q$$

4. GJS Improves Consistency and Generalization

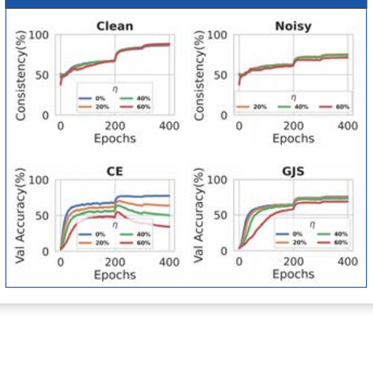

5. Results: Synthetic Noise - CIFAR-10 & CIFAR-100

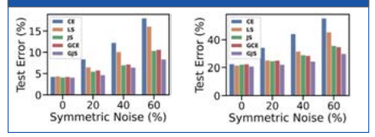

6. Results: Real-World Noise - WebVision

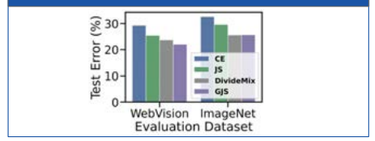

References

- Wei, Jiaheng, and Yang Liu. "When Optimizing f-divergence is Robust with Label Noise." ICLR 2021.
- Zhang et al. "Generalized cross entropy loss for training deep neural networks with noisy labels." *NeurIPS* 2018.
- Ghosh et al. "Robust loss functions under label noise for deep neural networks." *AAAI* 2017.
- Hendrycks, et al. "AugMix: A simple data processing method to improve robustness and uncertainty." ICLR 2020.

Acknowledgements

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

Eriksson, Hannes
Zenseact AB

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Safe Decision-Making for Autonomous Driving

For sequential decision-making problems such as autonomous driving it is imperative to consider the full range of outcomes as they might range from arriving to a location faster than expected, to being part of a catastrophic crash. In particular, we are mostly concerned with the left-tail properties of the distribution of outcomes.
By devising risk-aware agents that focus on performance in the worst outcomes, we can arrive at safer decision-makers.

Eriksson, Hannes
Zenseact AB

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Safe Decision-Making for Autonomous Driving

Hannes Eriksson
hannese@chalmers.se, hannes.eriksson@zenseact.com
Zenseact AB, Chalmers University of Technology

Abstract

For sequential decision-making problems such as autonomous driving it is imperative to consider the full range of outcomes as they might range from arriving to a location faster than expected, to being part of a catastrophic crash. In particular, we are mostly concerned with the left-tail properties of the distribution of outcomes.
By devising risk-aware agents that focus on performance in the worst outcomes, we can arrive at safer decision-makers.

Paper I: Epistemic Risk-sensitive Reinforcement Learning

In this work we studied the concept of *epistemic risk*, that is the risk that arises due to uncertainty about the model parameters μ. Typically this situation occurs when we have a belief over MDPs $\xi(\mu)$ and we want to optimize for a risk-sensitive objective w.r.t. the uncertainty due to ξ.
The main contributions were defining an entropic risk measure for epistemic risk and the delivery of two algorithms, one based on **Approximate Dynamic Programming** and one based on **Bayesian policy gradient**.

$$\pi^E(U) = \arg\max_\pi \frac{1}{\beta} \log \mathbb{E}\Big(\exp(\beta R)\Big). \quad (1)$$

Eq. 1. defines the objective over the utility function mirroring an entropic risk measure with some nice properties. By also considering the uncertainty induced by ξ we can arrive at the full objective in Eq. 2., by replacing U with the considered utility function in Eq. 1.

$$\pi^E(U, \xi) \triangleq \arg\max_\pi \int_{\mathcal{M}} U(\mathbb{E}^\pi_\mu[R]) \, \mathrm{d}\xi(\mu). \quad (2)$$

SENTINEL Schematic

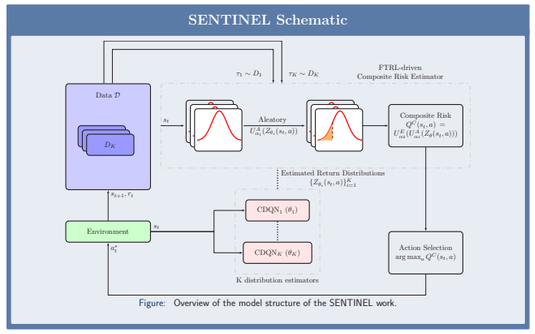

Figure: Overview of the model structure of the SENTINEL work.

Paper II: Inferential Induction: A Novel Framework for Bayesian Reinforcement Learning

In this work we introduced a novel Bayesian Reinforcement Learning framework that correctly infers value function distributions from data. From this framework, depending on what you marginalize over, gives arise to a whole new class of BRL algorithms. In particular, we develop and demonstrate comparable to state-of-the-art performance of **Bayesian Backwards Induction**.

$$\mathbb{P}^\pi_\beta(V_i \,|\, D) = \int_{\mathcal{V}} \mathbb{P}^\pi_\beta(V_i \,|\, V_{i+1}, D) \, \mathrm{d}\mathbb{P}^\pi_\beta(V_{i+1} \,|\, D). \quad (3)$$

Paper III: SENTINEL: Taming Uncertainty with Ensemble-based Distributional Reinforcement Learning

In *SENTINEL*, we study a novel kind of risk measure, in this work termed *composite risk*, which combines both the risk due to aleatory uncertainty and the risk due to epistemic uncertainty into one risk measure. We prove that this new risk measure better estimates the total risk than one that considers both risks separately.

$$\text{Comp. Risk} \triangleq \int_\Theta \int_Z Z \, \mathrm{d}(U^A_{\alpha} \circ \Pr)(Z|\theta) \, \mathrm{d}(U^E_{\alpha_0} \circ \beta)(\theta) \quad (4)$$

Further, we demonstrate how to design an agent that optimizes for this risk, using distribution estimators, as seen in the above schematic.

Paper IV: *In progress* Transfer Reinforcement Learning with Risk

In an ongoing work we are considering techniques that leverage knowledge transfer from a set of *source* domains to a *target* domain. This setting is interesting when you for instance have a task that you know how to solve and you now want to solve a different task but with similar structure. For instance, knowing how to drive in **Europe** should inform you to some extent how to drive in the **US**, but there are some differences, namely traffic rules, road signs and road behavior.

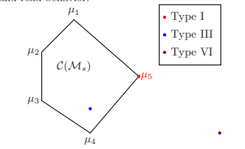

Figure: Overview some of the types of settings that may arise in Transfer Reinforcement Learning.

In particular, we consider three kinds of structures over model space. The first, which gives arise to *Type I*, assumes the target MDP μ_t is part of the finite set of source MDPs $\mathcal{M}_s$. The second, *Type III*, which considers the convex set of source MDPs $\mathcal{C}(\mathcal{M}_s)$, searches for the best representative $\hat{\mu}_t \in \mathcal{C}(\mathcal{M}_s)$. Finally, the last type of structure is the general case where the target MDP can be arbitrarily different from the source MDPs. In that case, you are studying problems of *Type V*.

Acknowledgements

I want to thank all my co-authors and the people I have worked with, some of these are, Emilio Jorge, Divya Grover, Debabrota Basu, Tommy Tram, Christos Dimitrakakis, Mina Ali-Beigi Nabi, Aristide Tossou.

Robot Learning of Symbol Grounding in Multiple Contexts Through Dialog

When a person talks to a robot about an object in the environment, the robot has to find which object the person is talking about. Symbol grounding is the task of finding a link between a word and one of the observed objects. It is a complex task: What information the robot can perceive (language and vision)? What links exist between language and vision? What links exist between the objects that can help symbol grounding? In (1) we build a platform with a Pepper robot that can see, listen, and talk about objects that it can see, and save this information. In (2) we extended (1) and assumed that the user wants to ground objects implicitly. In (3) we discuss a better vision model, that is not limited to recognizing the category of objects.

AI MLX | Page 44 B
Faridghasemnia, Mohamadreza
Örebro University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Robot Learning of Symbol Grounding in Multiple Contexts Through Dialog

Mohamadreza Faridghasemnia, Örebro University
Centre for Applied Autonomous Sensor Systems (AASS)
Supervisors: Prof. Alessandro Saffiotti, Prof. Lars Karlsson
Project partner: Bram Willemsen (KTH)

ÖREBRO UNIVERSITET

A brief outline

When a person talks to a robot about an object in the environment, the robot has to find which object the person is talking about. **Symbol grounding**: The task of finding a link between a word and one of the observed objects. It is a complex task: What information the robot can perceive (language and vision)? What links exist between language and vision? What links exist between the objects that can help symbol grounding? In 1. we build a platform with a Pepper robot that can see, listen, and talk about objects that it can see, and save this information. In 2. we extended (1) and assumed that the user wants to ground to objects implicitly. In (3) we discuss a better vision model, that is not limited to recognizing the category of objects.

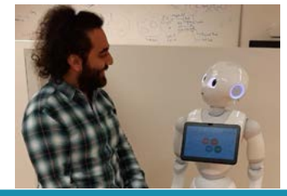

(1). Extract information from language and vision

Audio Speech Recognition: Audio signal to phrase.
Semantic mapping: Extract frame and frame element of language.
Vision: Extract object categories in the image.
AIML: Handling general dialog.
Manager: Holds detected object, Holds information from Semantic mapping module, Generate an utterance to confirm the user input
Manager works with symbols, while symbols are generated from applying deep neural networks on sensory inputs.

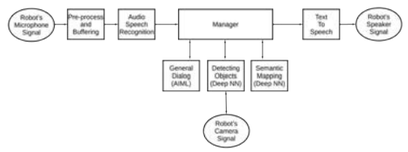

(2). Relationships that help symbol grounding

Ground to an object that is green and belongs to Mary. We have a mug that is green and kitchenware, a mug that is white and kitchenware, a scissor that is black, kitchenware, and belongs to Mary, and a ball and a car that are toys and are for Tom.

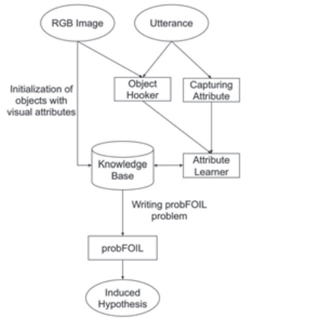

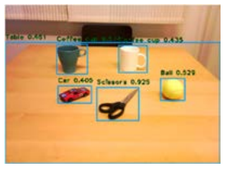

Robot ***Human***

Hi there! — Hello.
I see, the white mug is on the table — the white mug on the table
Got it! it belongs to mary — i guess it is for mary
Ok, I save it as kitchenware — its label is kitchenware

EAV data:
$mug(object001)$
$white(color001)$
$on-table(position001)$
$kitchenware(label001)$
$1.0 :: mary(object001, color001, position001, label001)$
$0.0 :: john(object001, color001, position001, label001)$
To: $mary(A, B, C, D, ...) : -kitchenware(C)$

(3). Extracting rich information from vision

The vision model was a deep Neural network object detector, and colors were recognized from RGB values. Why not a smarter vision model?

References

[1] Faridghasemnia, M. , et al., (2020). Towards Abstract Relational Learning in Human Robot Interaction.
[2] Faridghasemnia, M. (2020). Reducing Uncertainty in Contextualized Dialog.
[3] Faridghasemnia, M. , et al., (2019). Capturing Frame-Like Object Descriptors in Human Augmented Mapping.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Fay, Dominik
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Federated Learning for Smart Radiotherapy Systems

Today, treatment planning for radiotherapy involves several tedious and time-consuming tasks such as the segmentation of tumors in an MR image. Automating these tasks using machine learning could make radiotherapy substantially more cost-effective. Here, complications arise in cases where patient data must be stored locally at the clinic for privacy reasons. We investigate federated learning as a solution to learn a global privacy-preserving model. Aside from data storage requirements, we also aim for learning algorithms that satisfy statistical notions of privacy, adapt to local differences in clinical practice and continuously update the global model efficiently.

AI MLX | Page 45 B

Fay, Dominik
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Exploring the use of Phenotypic Screening data for Bioactivity Prediction

Developing new drugs is a long and costly process, both in terms of time and resources. Early stages of the process include High Throughput Screenings (HTS), probing the effects of hundreds of thousand or millions of compounds, on the target of interest. Predicting the outcome of such experiments could speed-up the drug discovery process and provide new insight into the underlying biological processes. Phenotypic screening data could potentially provide a high-resolution method for characterizing such compounds. In this work we explore the use of phenotypic screening data for predicting compound bioactivity. Our results indicate that the data captures information that can be linked to activity in a wide variety of assay targets, as well as information directly related to the underlying biology by confirming predictive performance across assay and target types as well as validating the predictive performance across assays.

AI MLX | Page 46 B
Fredin Haslum, Johan
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Deep Learning for Drug Discovery

Johan Fredin Haslum[1,2,3]
Supervisors: Kevin Smith[1,2], Erik Müllers[3], Karl-Johan Leuchowius[4], Riku Turkki[5]
[1] Science for Life Laboratory, Stockholm, Sweden
[2] School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden
[3] Bioscience Cardiovascular, Research and Early Development, Cardiovascular, Renal and Metabolism (CVRM), BioPharmaceuticals R&D, AstraZeneca, Gothenburg, Sweden
[4] Mechanistic Biology, Discovery Sciences, R&D, AstraZeneca, Gothenburg, Sweden , [5] Data Sciences & Quantitative Biology, Discovery Sciences, R&D, AstraZeneca, Gothenburg, Sweden

Exploring the use of Phenotypic Screening data for Bioactivity Prediction

Developing new drugs is a long and costly process, both in terms of time and resources. Early stages of the process include High Throughput Screenings (HTS), probing the effects of hundreds of thousand or even millions of compounds, on the target of interest. Predicting the outcome of such experiments could speed-up the drug discovery process and provide new insight into the underlying biological processes. Phenotypic screening data could potentially provide a high-resolution method for characterizing such compounds. In this work we explore the use of phenotypic screening data for predicting compound bioactivity. Our results indicate that the data captures information that can be linked to activity in a wide variety of assay targets, as well as information directly related to the underlying biology by confirming predictive performance across assay and target types as well as validating the predictive performance across assays focused on the same target.

Problem

Identifying compounds active towards a target of interest is an important step of early drug discovery. High Throughput Screening (HTS) Assays, are experiments focused on identifying compounds with biological activity towards a particular target and are designed in order to efficiently screen millions of compounds. Bioactivity Prediction is an attractive solution to limit the number of experiments that needs to be performed. Enabling the prioritization of compounds with higher likeliness of activity. Potentially reducing the size of such HTS Screens.

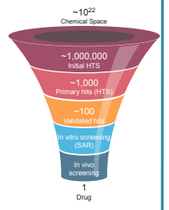

Method

Predicting compound bioactivity, based on phenotypic screening data is an attractive alternative to compound structure-based approaches[4], due to it's potential of capturing a compound's effect on a biologically system. Previous work have shown that such data contain information relevant for predicting bioactivity in orthogonal assays [1,2]. In this work we built a phenotypic screening dataset, using the Cell Painting protocol [3] to generate high resolution fluorescent microscopy images. In effect characterizing compound by their biological effect in a cell line and the morphological changes it induces. Combined with repurposed bioactivity readouts for a wide range of compounds and targets. A CNN is trained in a supervised multi-task manner to predict bioactivity in multiple assay at once.

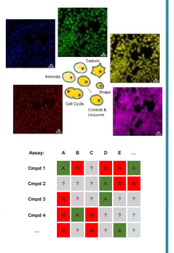

References

1. Simm, Jaak, et al. "Repurposing high-throughput image assays enables biological activity prediction for drug discovery." Cell chemical biology 25.5 (2018): 611-618.
2. Hofmarcher, Markus, et al. "Accurate prediction of biological assays with high-throughput microscopy images and convolutional networks." Journal of chemical information and modeling 59.3 (2019): 1163-1171.
3. Bray, Mark-Anthony, et al. "Cell Painting, a high-content image-based assay for morphological profiling using multiplexed fluorescent dyes." *Nature protocols* 11.9 (2016): 1757-1774.
4. Rogers, David, and Mathew Hahn. "Extended-connectivity fingerprints." *Journal of chemical information and modeling* 50.5 (2010): 742-754.
5. Carpenter, Anne E., et al. "CellProfiler: image analysis software for identifying and quantifying cell phenotypes." Genome biology 7.10 (2006): 1-11.

Results

Predicting Bioactivity over 140 different assays using Cell Painting images as input. 6-fold cross validation gives an average ROC-AUC prediction score of 0.744. With a majority (71%) of assays being predicted with a ROC-AUC above 0.7. Indicating that information relevant for assay activity prediction is captured in the Cell Painting images and that our network is able to recognize it.

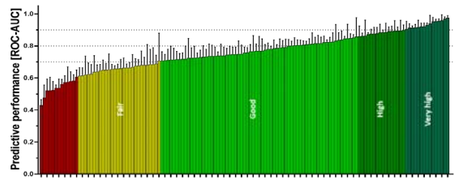

Screening compounds according to the models ranking, significantly enriches the number of actives identified. Yielding a much higher HIT-rate than expected by randomly sampling the compounds, but also better than competing methods using compound structure (ECFP4 [4]) or Cell Profiler features [5] as input.

Input modality	Phenotypic Deep Learning	Phenotypic Cell Profiler	Structure Based, ECFP4
ROC-AUC	0.744	0.714	0.687

Beyond higher ROC-AUC performance, we show that Cell Painting based predictions results in a more diverse set of compounds, compared to structure based. With a Wilcoxon two-sided test of structure diversity between the top ranked compounds in the test set and the known hits from the training set, we find that there is a significant difference between the two input modalities.

To further validate the performance of our model we also probe the activity of the top ranked compounds in several follow-up assays, meant to confirm the activity of each active compound under different experimental conditions. Five different assays with varying targets and model performances were selected and followed up in wet-lab experiments. Showing similar results in HTS and wet-lab assays, confirming that the network is capturing target relevant features when predicting bioactivity and enriching the compound set.

Assay	HTS Assay ROC-AUC	Enrichment at X percentile	
		90%	98%
A	0.95	6.3	30
B	0.68	3.6	15
C	0.94	3.8	16
D	0.63	4.2	7
E	0.81	3.5	3.9

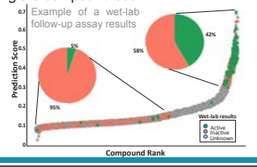

Generative Aging of Brain Images with Diffeomorphic Registration

Predicting subject-specific brain aging can be used for improving the diagnosis and prognosis in neurodegenerative diseases. Previous approaches have been restricted to group-level predictions, or yielded unreal results. This study addresses these issues by proposing a novel method that generates synthetic MRI images of the brain to simulate its changes due to aging. The method is based on diffeomorphic image registration, which can provide more accurate and fidelity-controllable subject-level predictions.

Generative Aging of Brain Images with Diffeomorphic Registration

Jingru Fu[1], Daniel Ferreira[2], Rodrigo Moreno[1]
1 KTH Royal Institute of Technology
2 Karolinska Institute

Research goals

Predicting subject-specific brain aging can be used for improving the diagnosis and prognosis in neurodegenerative diseases. Previous approaches have been restricted to group-level predictions [1], or yielded unreal results [4]. This study addresses these issues by proposing a novel method that **generates synthetic MRI images** of the brain to simulate its changes due to **aging.** The method is based on **diffeomorphic image registration**, which can provide more ***accurate*** and ***fidelity-controllable subject-level*** predictions.

Method

Diffeomorphic Registration:

- **Diffeomorphic registration (*DR*)** enjoys many advantages [2], such as the learned deformation field are differentiable and invertible, and thus preserve topology. According to [3], the contraction of the brain structures shows a linear change in old age. This gives us theoretical support to use the registration of two images at different ages to simulate the changes in brain aging.

Architecture:

- As shown in the figure below, we used images acquired at different time points from the same subject. The objective is to generate images in between these two time points using the velocity field estimated through deep learning-based diffeomorphic registration.

Two modules are introduced within the *DR* framework:

- ***Aging Generative Module (AGM)*** integrates the velocity fields at different time steps to generate several deformation fields that are used to generate the synthetic images.
- ***Quality Control Module (QCM)*** takes synthetic images quality into consideration, leading to adjusting the stopping point of the integration layer (s) in the *AGM*. It is possible to extrapolate images by using s > 1.

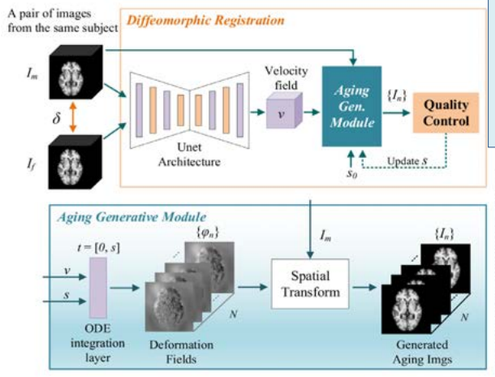

I_f: fixed image of subject
I_m: moving image of subject
$\{I_n\}$: set of generated images
δ: age difference
N: number of generated images
$\{\varphi_n\}$: set of deformation fields
t: integration interval
s: stopping point of integration layer

References

1. Sharmin Pathan and Yi Hong. *Predictive image regression for longitudinal studies with missing data*. arXiv preprint arXiv:1808.07553, 2018.
2. Adrian V Dalca et al. *Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces*. **Medical image analysis**, 57:226–236, 2019.
3. Kristine B. Walhovd et al. *Effects of age on volumes of cortex, white matter and subcortical structures*. **Neurobiology of Aging**, 26(9):1261–1270, October 2005. ISSN 0197-4580.
4. Viktor et al. *Generative Aging of Brain MR-Images and Prediction of Alzheimer Progression*.**GCPR**, 2019.

Results

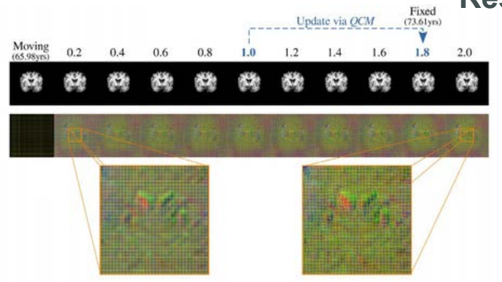

Generative aging of one sample subject from the non-demented subset of the OASIS-3 dataset. The first row represents the generated images from 0 to 2; the second row represents the corresponding deformation.

Results show in the left figure:

- Twenty images were generated for this subject with values of *t* between 0 and 2 (the moving image is at *t*=0). Although t=1 should theoretically match the fixed image, in practice it is not (*t*=1.8 in the example). The *QCM* adjusts *s* with quality measurements to make the generated images better correspond to chronological age.

Conclusions:

- As shown in the second row of the image, brain deformations are increasing with time.
- The entire generation for an individual can be completed in about five minutes using only CPU.
- We have generated synthetic images for three large-scale longitudinal datasets so far.

AIDA Analytic Imaging Diagnostics Arena

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

ResNet-based ECG Diagnosis of Myocardial Infarction in the Emergency Department

Myocardial infarctions (MIs) are often missed in the emergency department. In managed settings deep learning models have shown promise in electrocardiogram (ECG) classification. However, in a real-world scenario there is a lack of high performing models for classification of MIs. We developed a ResNet-based deep neural network to classify the ECG between non-ST-elevation MI (NSTEMI), ST-elevation MI (STEMI), and control status in the more challenging real-world setting. In a test set, our model discriminates STEMIs/NSTEMIs with an AUROC of 0.99/0.83 and a Brier score of 0.05/0.05. The model also generalizes well and obtains a similar performance on an additional test set collected in the months following the initial collection and that does not overlap temporally with the set used for developing the model. Our results are above human-level performance reported in previous studies for STEMIs and NSTEMIs.

AI MLX | Page 48 B
Gedon, Daniel
Uppsala University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Gillsjö, David
Lund University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

In Depth Semantic Scene Completion

This work studies Semantic Scene Completion which aims to predict a 3D semantic segmentation of our surroundings, even though some areas are occluded.
For this we construct a Bayesian Convolutional Neural Network (BCNN), which is not only able to perform the segmentation, but also predict model uncertainty.
This is an important feature not present in standard CNNs.
We show results for the Semantic Scene Completion task where a category is introduced at test time on the SUNCG dataset.
In this complex task the Bayesian approach outperforms the standard CNN, showing better Intersection over Union score and excels in mean Average Precision.
With the added benefit of having better calibrated scores and the ability to express model uncertainty.

AI MLX | Page 49 B

Gillsjö, David
Lund University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

In Depth Semantic Scene Completion

David Gillsjö, Kalle Åström
Centre for Mathematical Sciences, Lund University, Sweden

LUND UNIVERSITY

Introduction

Semantic Scene Completion (SSC) is a challenging task in which both visible and occluded surfaces are labeled semantically in 3D. In Figure 1 we see an illustration of the problem where a UAV would benefit from knowing what to expect in occluded areas.

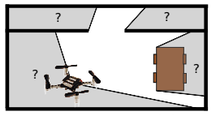

Figure 1: An UAV has some occluded areas in its surrounding and would like to have an idea about what to expect.

Our contributions include:

- An open source system for BSSC using Variational Inference released on `https://github.com/DavidGillsjo/bssc-net`.
- An extended SSC task on the SUNCG dataset with more occluded space.
- Experiments showing that the Bayesian approach is more robust to unseen data in the SSC task.
- Parameter studies on both MNIST and SUNCG.

Bayes by backprop

This method introduced by [1] is based on Variational Inference. Each weight in the network is sampled from a normal distribution, as illustrated in Figure 2.
We estimate the posterior $P(w|\mathcal{D})$ using a simpler model $q(w|\theta)$ with learnable parameters θ, which minimizes the approximate Kullback-Leibler (KL) divergence to the true posterior.

$$\theta^* = \arg\min_\theta \sum_{i=1}^{n} \frac{\beta}{n} \underbrace{\left[\log q(w^{(i)}|\theta) - \log P(w^{(i)})\right]}_{\text{Complexity}} \underbrace{- \log P(\mathcal{D}|w^{(i)})}_{\text{Likelihood}}.$$

where $w^{(i)}$ is a sample from the variational posterior $q(w^{(i)}|\theta)$. The scale factor $\frac{\beta}{n}$ with β as design parameter is introduced to tune the regularization.

(a) Standard (b) Bayesian

Figure 2: In 2a we see a filter bank from a standard 2D CNN, each weight is a scalar. In 2b we see a filter bank in a Bayesian Variational Inference 2D CNN, here each weight represented as a distribution which is sampled from at inference time.

Prediction & Uncertainty

An unbiased estimation of the expectation is given [2] by

$$\mathbb{E}_{q(w|\theta)}[P(\hat{y}|\hat{x}, w)] = \int q(w|\theta) p_t \, dw \approx \frac{1}{T} \sum_{t=1}^{T} p_t,$$

where $p_t := P(\hat{y}|\hat{x}, w^{(t)})$ is the softmax output from forward pass t.
For uncertainty we use Predictive Entropy,

$$H = -\sum_{t=1}^{T} p_t \log p_t.$$

For metrics we use mean Average Precision (mAP), Intersection over Union (IoU) for performance. For separation metric we use the Bhattacharyya coefficient (BC)

$$BC(p, q) = \frac{1}{N} \sum_{i=1}^{N} \sqrt{p_i q_i},$$

where N is the number of categories, q_i and p_i are the number of TP and FN. Lower score indicates better separation.

Model

We have explored two network architectures. The first network architecture is inspired by the original SUNCG article [3]. We call it **SSC-Net**. The second architecture is a **UNet**. We chose softplus as activation functions instead of relu to have more active weights in the network [2]. The architecture is displayed in Figure 3.

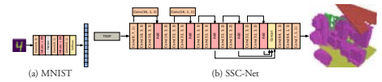

(a) MNIST (b) SSC-Net

Figure 3: Architecture of SSC-Net used for MNIST and SUNCG experiments. Conv(d, k, l) stands for a 3D convolution filter stack of depth d and kernel size k and dilation l. Batch normalization and softplus activation is performed after every Conv layer. Softmax in the final layer.

Selected References

[1] C. Blundell, J. Cornebise, K. Kavukcuoglu and D. Wierstra, 'Weight uncertainty in neural networks,' *arXiv preprint arXiv:1505.05424*, 2015.
[2] K. Shridhar, F. Laumann and M. Liwicki, *A comprehensive guide to bayesian convolutional neural network with variational inference*, 2019. arXiv: 1901.02731 [cs.LG].
[3] S. Song, F. Yu, A. Zeng, A. X. Chang, M. Savva and T. Funkhouser, 'Semantic scene completion from a single depth image,' *Proceedings of 30th IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

MNIST Experiment

In Figure 4 we see output distributions from MNIST test set for digits 0 and 1 when 0 is removed from the training data. The Bayesian Score is more better calibrated and the Entropy is higher for 0.

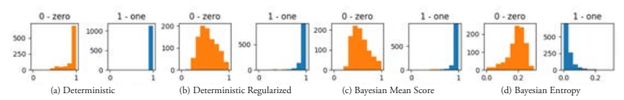

(a) Deterministic (b) Deterministic Regularized (c) Bayesian Mean Score (d) Bayesian Entropy

Figure 4: Here we see true (blue) and false (orange) predictions for 0 and 1.

SUNCG Experiment

SUNCG [3] is a large dataset with manually created and labeled synthetical indoor scenes. We've used a subset of 2000 training and 1000 testing scenes for the experiments. In Figure 5 we show a parameter study on β.Figure 6 shows example output.
We also conducted an experiment when category *bed* was removed from training, the result is presented in Table 1.

Table 1: BC, mAP and mIoU for different network architectures when the *bed* class is removed from training. S=Score, E=Entropy. We observe that Bayesian SSC-Net has the best score in all metrics.

CNN	mIoU	mAP: S	mAP: E	BC: S	BC: E
SSC-Net ω=0	0.19	0.2		0.31	
SSC-Net ω=0.01	0.14	0.23		0.29	
B-SSC-Net	**0.21**	**0.26**	**0.19**	**0.27**	**0.28**

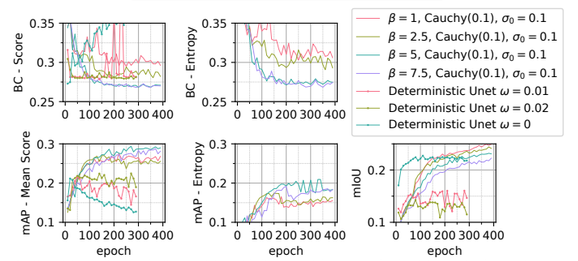

Figure 5: BC, mAP and mIoU for the Bayesian UNet with different weights β and ω for the SUNCG mini dataset. We observe that $\beta = 5$ is better in all metrics but mIoU, where $\beta = 1$ is best.

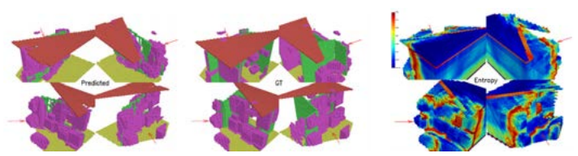

Figure 6: Example from the SUNCG test set. From the left we have predicted, true labels and entropy.

Govindarajan, Hariprasath
Linköping University / Arriver Sweden AB

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Self-Supervised Representation Learning for Content Based Image Retrieval of Complex Scenes

Although Content Based Image Retrieval (CBIR) is an active research field, application to images simultaneously containing multiple objects has received limited research interest. For such complex images, it is difficult to precisely convey the query intention, to encode all the image aspects into one compact global feature representation and to unambiguously define label similarity or dissimilarity. Motivated by the recent success on many visual benchmark tasks, we propose a self supervised method to train a feature representation learning model. We propose usage of multiple query images and use an attention based architecture to extract features from diverse image aspects that benefits from this.

Govindarajan, Hariprasath
Linköping University / Arriver Sweden AB

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Self-Supervised Representation Learning for Content Based Image Retrieval of Complex Scenes

LINKÖPING UNIVERSITY
Hariprasath Govindarajan, Peter Lindskog, Dennis Lundström,
Amanda Olmin, Jacob Roll, and Fredrik Lindsten
Department of Computer and Information Science (IDA)
Arriver

Introduction

Although **Content Based Image Retrieval (CBIR)** is an active research field, application to images simultaneously containing multiple objects has received limited research interest. For such complex images, it is difficult to precisely convey the query intention, to encode all the image aspects into one compact global feature representation and to unambiguously define label similarity or dissimilarity. Motivated by the recent success on many visual benchmark tasks, we propose a **self- supervised method** to train a feature representation learning model. We propose usage of **multiple query images** and use an **attention-based architecture** to extract features from diverse image aspects that benefits from this.

Method

Attention Based Ensemble architecture (ABE-M) [1]

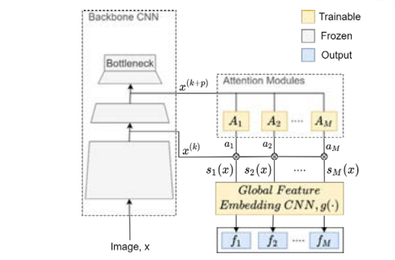

Self-supervised Training

We propose a self-supervised training method based on a contrastive pretext task using frames from videos. We expect a frame to be more similar to an adjacent frame in the video than to frames from other videos.

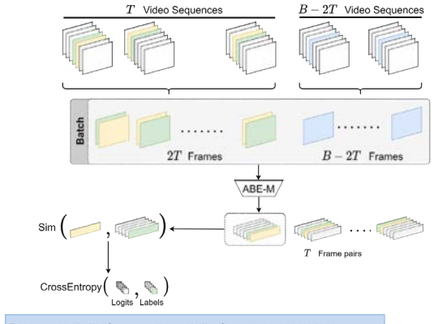

Encourages similar frames to have similar feature representations.

$$\mathcal{L}_{\text{NCE},t} = -\log \frac{\exp(f(x_t) \cdot f(x_{t'})/\tau)}{\sum_{x \in X\_t} \exp(f(x_t) \cdot f(x)/\tau)}$$

NCE Loss

$$\mathcal{L}_{\text{NCE}} = \frac{1}{T} \sum_{t=1}^{T} \mathcal{L}_{\text{NCE},t}$$

$$\mathcal{L}_{\text{div}}^{(b)} = \frac{1}{M(M-1)} \sum_{1 \le p < q \le M} \max\left(0, \delta - \mathrm{d}\left(f_p^{(b)}, f_q^{(b)}\right)\right)$$

Divergence Loss

$$\mathcal{L}_{\text{div}} = \frac{1}{B} \sum_{b=1}^{B} \mathcal{L}_{\text{div}}^{(b)}$$

Encourages different attention heads to focus on different spatial regions.

Overall Loss

$$\mathcal{L} = \mathcal{L}_{\text{NCE}} + \lambda \mathcal{L}_{\text{div}}$$

Results

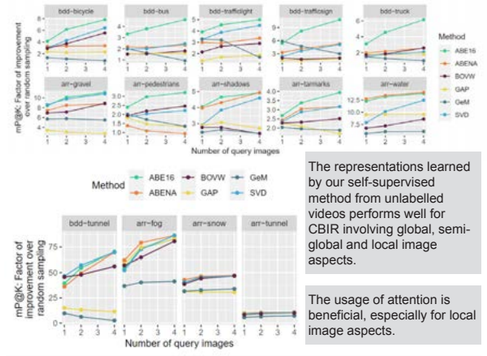

The representations learned by our self-supervised method from unlabelled videos performs well for CBIR involving global, semi-global and local image aspects.

The usage of attention is beneficial, especially for local image aspects.

The usage of multiple query images is useful for conveying query intention for CBIR involving images with multiple image aspects. It also consistently improves the performance of our method on all image aspects.

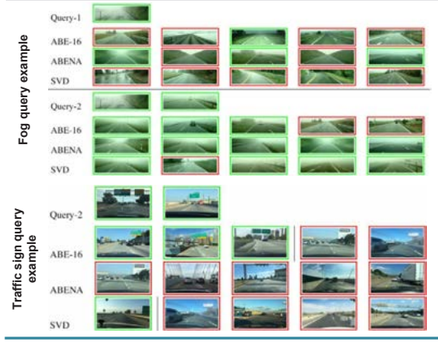

References

1. W. Kim, B. Goyal, K. Chawla, J. Lee, and K. Kwon, "Attention-based ensemble for deep metric learning," in ECCV 2018.

Automation of scientific discovery in systems biology using active learning

In our research, we are seeking to automate the discovery of scientific knowledge by building a robot scientist – a combination of laboratory automation hardware and artificial intelligence (AI) capable of closed-loop cycles of experimentation. This means the robot scientist will design experiments, execute them, analyse results and generate new scientific knowledge without human intervention. We direct our robot scientists, Eve and Genesis, toward generating new scientific knowledge about the metabolism of the yeast Saccharomyces cerevisiae. Two crucial parts of the scientific process are: hypothesis generation and applying results of experiments to refine theory. We implement experiment selection algorithms to decide gene knockouts based on influence profiles of genes over certain sections of metabolism. The aim is to select experiments that will provide maximum information gain for the area of metabolism we seek to understand.

AI MLX | Page 51 B
Gower, Alexander
Chalmers
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Automation of scientific discovery in yeast systems biology using active learning

Alexander Gower[a], Daniel Brunnsåker[a], Ievgeniia Tiukova[a] and Ross King[abc]

[a] Chalmers University of Technology, Gothenburg, Sweden
[b] Cambridge University, Cambridge, United Kingdom
[c] Alan Turing Institute, London, United Kingdom

CHALMERS UNIVERSITY OF TECHNOLOGY

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

The case for a robot scientist

Our research aims to automate the discovery of scientific knowledge by building a **robot scientist** – a combination of laboratory automation hardware and artificial intelligence (AI) capable of closed-loop cycles of experimentation. This means the robot scientist will design experiments, execute them, analyse results and **generate new scientific knowledge without human intervention**. We direct our robot scientists, Eve and Genesis, toward generating new scientific knowledge about the metabolism of the yeast *Saccharomyces cerevisiae*.

A common approach to generate new knowledge about cell metabolism is through **factorial experimentation**, where the scientist creates a mutant strain by removing one or more genes and cultivates the yeast in various conditions. The potential number of experiments is vast–*Saccharomyces cerevisiae* has approximately 6000 genes. As such, heuristics are employed to select experiments that will generate the most knowledge. A robot scientist has many advantages: **selecting experiments likely to yield the most information for the least cost**; greater precision and reproducibility; greater capacity; and high-throughput data analysis.

An integrated model of a yeast cell

The background knowledge of yeast cell biology that our robot scientists use is in the form of an integrated computational model. Here we use a **modular framework**–we model cell signalling, gene regulation and metabolism separately and integrate them during simulation calculations. This means we can choose model structures that suit the underlying biological processes.

The model for our robot scientists is comprised of: a boolean cell-signaling network; a directed hypergraph of co-regulatory relationships between genes; and a genome-scale metabolic model (Yeast8). The regulatory model is reduced to a dynamic Bayesian network for simulation. As all these processes happen on different timescales and with different magnitudes, we are experimenting with several simulation protocols.

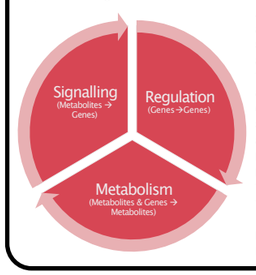

A modular model has other advantages for a robot scientist: machine learning algorithms can be selected to match the model structure and best make use of experimental data; the robot scientist can isolate changes a specific section of the model; and we can select a mathematical formulation that is suited to representing the biological processes (e.g. stoichiometric matrix for metabolic reactions, Boolean rules for signaling).

The robot scientist's research pipeline

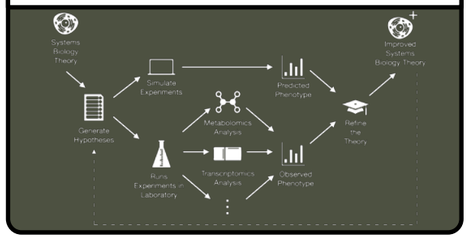

Automated model improvement

A model of yeast cell metabolism can be improved in many different ways. For a robot scientist to be able to make improvements algorithmically, a specific procedure to reach a goal must be specified. We define these goals as **units of discovery**, for example:

- a new regulatory interaction;
- updating an existing regulatory interaction;
- a regulatory effect between groups of genes; and
- a new signaling interaction (link between some condition and some response).

Algorithms can be defined within the structure of our integrated yeast model in order that the **robot scientist can suggest revisions to the model**. These algorithms exploit an optimisation objective that represents the **quality of the model**. There are many such objectives one could use, however we focus on a model's predictive power. In practice this means the robot scientist compares simulations based on different prospective models against experimentally obtained transcriptomic and metabolomic data, along with fermentation profiles.

Metrics of quality for scientific models

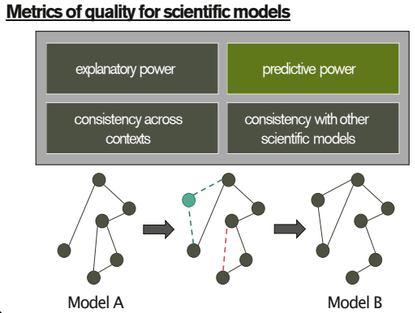

[1] A. Coutant et al., 'Closed-loop cycles of experiment design, execution, and learning accelerate systems biology model development in yeast', *PNAS*, vol. 116, no. 36, pp. 18142–18147, Sep. 2019, doi: 10.1073/pnas.1900548116.
[2] D. T. Banos, P. Trébulle, and M. Elati, 'Integrating transcriptional activity in genome-scale models of metabolism', *BMC Systems Biology*, vol. 11, no. 7, p. 134, Dec. 2017, doi: 10.1186/s12918-017-0507-0.

This work is partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

Grönqvist, Johan
Lund University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Closed Loop Guarantees with Neural Networks in Control

We use methods from Robust Control to obtain guarantees of closed loop stability for a dynamical system controlled using a neural network.

AI MLX | Page 52 B

Grönqvist, Johan
Lund University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Closed Loop Guarantees with Neural Networks in Control

Johan Grönqvist, Lund University
Department of Automatic Control

LUNDS UNIVERSITET

Charge while Driving

- Pick-Up Contact under vehicle
- PID controller to track Charging Rail
- Camera and Neural Network Estimate State
- Master Thesis at Automatic Control LTH

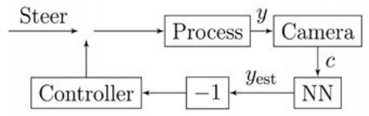

Closed Loop System

- Linear Time-Invariant Process
- Camera and Neural Network Estimate State
- Inaccurate Estimate: Nonlinear Disturbance
- Nonlinear Blocks: Quadratic Forms

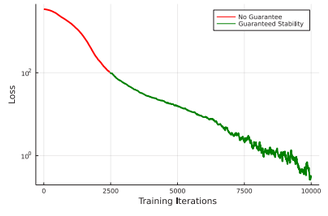

Neural Network and Stability Guarantee

- Prediction Accuracy as Loss Function
- NN challenges: Overfitting and Adversaries
- Neurons described by Quadratic Forms
- IQC Proof of Global Stability
- Several Kinds of Guarantees Possible

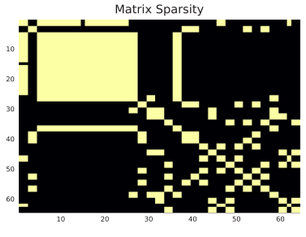

Sparse Semi-definite Programming

- Problem Dependent Structure
- Depends on Choices in IQC Description
- Challenge: Deep and Large Networks
- Sparsity to the Rescue

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Can Industrial Transport Applications be improved using AI planning?

AI planning methods are fundamental in industrial transport applications. These methods typically rely on manually-specified knowledge to derive plans. The goal of my PhD is to use Machine Learning to enhance AI planning methods by learning from human planning experts. The purpose is to develop learning algorithms to help humans build domains, aiming at

1. reducing the knowledge engineering effort for humans

2. providing better quality plans in many domains

3. allowing automated planning and scheduling systems to learn how to execute plans or policies from previous experience.

In this poster I present an approach to address this problem with its the main challenges and I highlight the industrial need.

AI MLX | Page 53 B
Gugliermo, Simona
Örebro University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Can Industrial Transport Applications be improved using AI planning?

Simona Gugliermo, Örebro University
Center of Applied and Autonomous Sensors (AASS)
Advisors: Federico Pecora, Christos Koniaris

Motivation

- AI planning methods are useful for industrial transport applications.
- Current AI planning methods rely on manually-specified knowledge, encoded in the *planning domain*.
- AI planning domains are difficult to write for humans.
- Aim: use Machine Learning to
  1. Reduce the knowledge engineering effort for humans
  2. Provide better quality plans in many domains
  3. Allow AI planning systems to learn how to execute plans or policies from previous experience.

Proposed Approach

Main Features

- Involve fleet transport managers in the learning loop
- Make the learning process more transparent and tractable
- Increase the generality step by step
- Create an application-independent method, usable with different scenarios.

1. Learn a Behavior Tree (BT) from executed traces (data mining techniques).
2. Revise and correct (if needed) the learned Behavior Tree by humans to fully control the resulting behavior.
3. Approval of the final Behavior Tree.
4. Learn a domain description from the Behavior Tree

Research Questions

1. **How to learn behaviors from execution traces?**
   Related work
   - Learn Behavior Trees from traces in the real-time strategy gaming context [1].
     1. A maximally specific BT is created from the given trace
     2. The BT is iteratively reduced in size by finding and combining common patterns of actions.
     3. If no new satisfactory patterns are found, the process stops.

   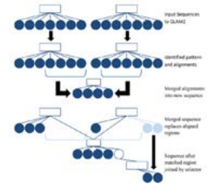

   - Learn Behavior Trees and transfer new skiils using programming by demonstration (PbD) [2,3]

   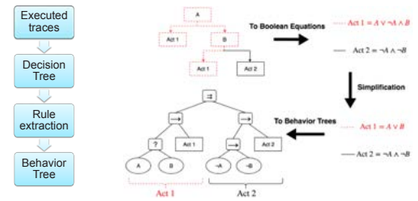

2. **How to facilitate human modelling of BTs?**

3. **How to obtain a domain theory from BTs?**
   Research studies [4,5] have been conducted to extract BTs from planning domains. The idea is instead to extract a planning domain from a BT.

References

1. Robertson and Watson, (2015). Building behavior trees from observations in real-time strategy games. In *2015 International Symposium on Innovations in Intelligent SysTems and Applications (INISTA)*.
2. Sagredo-Olivenza et al., (2019). Trained Behavior Trees: Programming by Demonstration to Support AI Game Designers. *IEEE Transactions on Games*.
3. French et al., (2019). Learning Behavior Trees From Demonstration. In *2019 International Conference on Robotics and Automation (ICRA)*.
4. Colledanchise et al. (2019a). Towards Blended Reactive Planning and Acting using Behavior Trees. *In IEEE International Conference on Robotics and Automation*, Montreal, Canada.
5. Martín et at., (2021). Optimized Execution of PDDL Plans using Behavior Trees. In *International Foundation for Autonomous Agents and Multiagent Systems (AAMAS)*

Industrial need

AI planning can introduce disruptive changes in industrial transport applications.

A decision support tools can increase the capabilities of fleet transport managers.

Better company performance and better satisfaction of their customers.

Overcome the challenge of transferring the fleet transport managers knowledge

Sequential Decision-making in Drug Discovery

Two new paradigms have emerged in the pharmaceutical industry to increase the productivity in drug design: (1) AI-augmented molecular design that utilizes generative models for sampling the chemical space; (2) Automated laboratories together with machine learning to make, test and analyze potential drug candidates without human intervention. This work focuses on how to select what molecules to make in order to explore and exploit the chemical space in an efficient way. The results can enable a faster and better way to find drug candidates, ultimately providing new drugs for unmet medical needs faster to the benefit of patients worldwide.

AI MLX | Page 54 B
Gummesson Svensson, Hampus
Chalmers
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Sequential Decision-making in Drug Discovery

Hampus Gummesson Svensson[†,‡],
Esben Jannik Bjerrum[‡], Christian Tyrchan[‖], Alexander Schliep[♯],
Ola Engkvist[‡], Morteza Haghir Chehreghani[†]

UNIVERSITY OF GOTHENBURG
CHALMERS UNIVERSITY OF TECHNOLOGY
AstraZeneca

Motivation & Research goals

Two new paradigms have emerged in the pharmaceutical industry to increase the productivity in drug design: (1) AI-augmented molecular design that utilizes generative models for sampling the chemical space; (2) Automated laboratories together with machine learning to make, test and analyze potential drug candidates without human intervention. This work focuses on how to select what molecules to make in order to explore and exploit the chemical space in an efficient way. The results can enable a faster and better way to find drug candidates, ultimately providing new drugs for unmet medical needs faster to the benefit of patients worldwide.

Background

Drug optimization is a **complex**, **multiparameter optimization** with dozens of non-correlating and even **opposing parameters**. Therefore, the process of finding drug candidates is an **iterative process**.

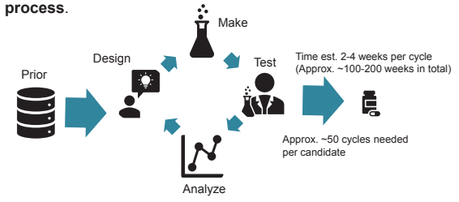

Generative models such as REINVENT [1] can be used to **design** thousands of molecules by sampling the chemical space, which is estimated to consist of up to 10^{60} drug-like molecules. Subsequently, to understand the true properties of a molecule, we need to **make** it. If a molecule is successfully made, we can **test** and **analyze** its properties to better steer the generation of molecules. On the other hand, it is only possible to make a few molecules since each experiment is **costly** and **time-consuming**, which limits the amount of information that can be acquired in each iteration.

Goal: Optimize the selection of molecules to test, to explore and exploit desired areas of the chemical space in an optimal way.

Previous contributions (see [2]):
- **Active learning** to help to decrease the amount of data needed to **develop robust models for reaction yield predictions**, helping to successfully make molecules.

References

1. Blaschke, Thomas, et al. "REINVENT 2.0: an AI tool for de novo drug design." *Journal of Chemical Information and Modeling* 60.12 (2020): 5918-5922.
2. Johansson, Simon Viet, et al. "Using Active Learning to Develop Machine Learning Models for Reaction Yield Prediction." *ChemRxiv* (2021).
3. Slivkins, Aleksandrs. "Contextual bandits with similarity information." *Proceedings of the 24th annual Conference On Learning Theory*. JMLR Workshop and Conference Proceedings, 2011.
4. Chen, Lixing, Jie Xu, and Zhuo Lu. "Contextual combinatorial multi-armed bandits with volatile arms and submodular reward." *Advances in Neural Information Processing Systems* 31 (2018): 3247-3256.
5. Nika, Andi, Sepehr Elahi, and Cem Tekin. "Contextual combinatorial volatile multi-armed bandit with adaptive discretization." *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020.

Approach

Challenges:
1. "Infinite" space (generated)
2. New molecules in each cycle
3. Select several molecules to make and test

Proposed solution:
To study this problem in the context of **sequential decision-making**, where the goal is to **adaptively** compute the most informative decisions.

Formerly, we currently seek a multi-armed bandit that can handle the following settings:
- **Contextual**
- **Infinite** action and/or context space
- **Volatile** arms
- **Combinatorial**

Previous works show approaches for several or all these settings using similarity information [3,4,5]. We have a **high-dimensional** and complicated context and large **unfixed** dataset (generated). Hence, we need to adapt and extend these methods to our problem.

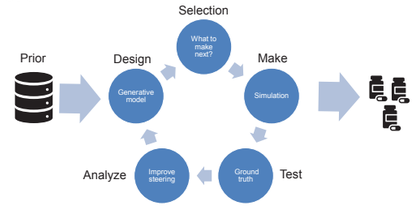

A framework for **simulating** the different steps in the drug discovery process is being developed. This will help fast **evaluation and comparison** of different sequential decision-making strategies for selecting what molecules to make.

† Department of Computer Science and Engineering, Chalmers University of Technology
‡ Molecular AI, Discovery Sciences, R&D, AstraZeneca
‖ Medicinal Chemistry, Research and Early Development, Respiratory and Immunology (R&I), BioPharmaceuticals R&D, AstraZeneca
♯ Department of Computer Science and Engineering, University of Gothenburg

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

AI-based Mental Well-being monitoring during activities of daily life

The way technology has evolved in the last decades has changed our vision on how to face the problems of our daily life. Artificial Intelligence (AI) is taking an important role, being present in many fields, as for example Health Care. This research aims to fuse the different types of laboratory measures (including behavioral, peripheral, and central nervous systems) with cutting-edge technology based on predictive algorithms. The main objective is to discover patterns in the data that give us cues of how a person is feeling with the final goal of designing models that are able to monitor mental well-being, avoiding depressive, anxious, or stressful states during daily life activities.

AI MLX | Page 55 B
Gutierrez Maestro, Eduardo
Örebro University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

AI-based Mental Well-being monitoring during activities of daily life

Eduardo Gutierrez-Maestro, Örebro University
Center for Applied Autonomous Sensor Systems (AASS)

Motivation & Research goals

The way technology has evolved in the last decades has changed our vision on how to face the problems of our daily life. Artificial Intelligence (AI) is taking an important role, being present in many fields, as for example Health Care. This research aims to fuse the different types of laboratory measures (including behavioral, peripheral, and central nervous systems) with cutting-edge technology based on predictive algorithms. The main objective is to discover patterns in the data that give us cues of **how a person is feeling** with the final goal of designing models that are able **to monitor mental well-being**, avoiding **depressive, anxious, or stressful states** during daily life activities.

Approach

- **Well-being:** individuals' own abilities to cope with normal stresses of life, to work productively and to be able to contribute his or her community (World Health Organization)
  - Examples: stress, happiness, calm, sad, depressed, etc.

- **Multi-modal** scenario:
  - Smartphones, wearables, presence sensors.
- Hypothesis: context information can help to understand well-being states → **Activities of Daily Living (ADL)**
- **ADL**: non-controlled environment
  - Routine analysis
  - Real-world solution

Current work

- Well-being **prediction** based on **physiological** data [1] [2]

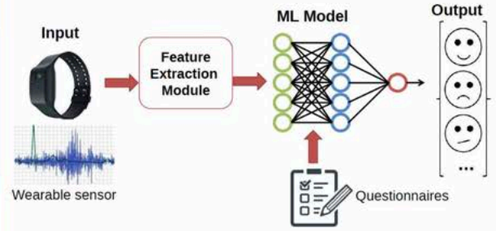

- How can we **label** physiological data with **emotions**? [3]

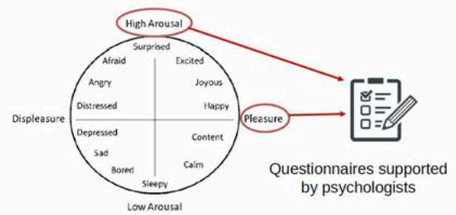

- First results on **mood prediction** [1]

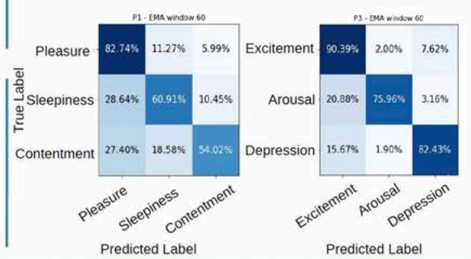

References

1. Calatrava-Nicolás, F.M.; Gutiérrez-Maestro, E.; Bautista-Salinas, D.; Ortiz, F.J.; González, J.R.; Vera-Repullo, J.A.; Jiménez-Buendía, M.; Méndez, I.; Ruiz-Esteban, C.; Mozos, O.M. Robotic-Based Well-Being Monitoring and Coaching System for the Elderly in Their Daily Activities. Sensors 2021
2. Bautista-Salinas, D.; Gonzalez, J.R.; Mendez, I.; Mozos, O.M. Monitoring and Prediction of Mood in Elderly People during Daily Life Activities. In Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Berlin, Germany, 23–27 July 2019
3. Russell et al., "A circumplex model of affect.", 1980

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Hagström, Lovisa
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Transferring Knowledge from Vision to Language: How to Achieve it and how to Measure it?

Large language models are known to suffer from the hallucination problem in that they are prone to output statements that are false or inconsistent, indicating a lack of knowledge. A proposed solution to this is to provide the model with additional data modalities that complements the knowledge obtained through text. We investigate the use of visual data to complement the knowledge of large language models by proposing a method for evaluating visual knowledge transfer to text for uni- or multimodal language models. The method is based on two steps, 1) a novel task querying for knowledge of memory colors, i.e. typical colors of well-known objects, and 2) filtering of model training data to clearly separate knowledge contributions. Additionally, we introduce a model architecture that involves a visual imagination step and evaluate it with our proposed method. We find that our method can successfully be used to measure visual knowledge transfer capabilities in models and that our novel model architecture shows promising results for leveraging multimodal knowledge in a unimodal setting.

Hagström, Lovisa
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Transferring Knowledge from Vision to Language: How to Achieve it and how to Measure it?

Tobias Norlund*, Lovisa Hagström*, Richard Johansson
Chalmers University of Technology

Motivation & Summary

Despite the ability of language models to learn and hold large quantities of structural knowledge [1], LMs are are also known to suffer from the **hallucination problem** in that they are prone to output statements that are false or inconsistent, indicating significant knowledge gaps [2]. A hypothesis is that this knowledge in many cases simply is missing in the large text corpora typically used for training the models, due to e.g reporting bias [3]. In such cases certain types of knowledge might also be more readily available in a different data modality.

In this work, we investigate **visual knowledge transfer**, i.e. to which extent language models can incorporate and *textually* express knowledge originating from a visual modality.

We investigate this by constructing a novel cloze-style task testing knowledge of **memory colors** for common objects (such as blood is red, a lemon is yellow etc). We also build a large vision-and-language dataset used for self-supervised training, and by careful filtering we make sure the color information is only available through the images and not the text.

Finally, we compare two strategies for how to most effectively query the trained multimodal language model for this visual knowledge.

We show that language models are able to texually express knowledge obtained from a visual modality, as a result from multimodal self-supervised training.

References
[1] Fabio Petroni et al. Language Models as Knowledge Bases? EMNLP (2019).
[2] Robert Logan et al. 2019. Barack's wife Hillary: Using knowledge graphs for fact-aware language modeling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy. Association for Computational Linguistics.
[3] Jonathan Gordon and Benjamin Van Durme. 2013. Reporting bias and knowledge acquisition. In *Proceedings of the 2013 workshop on Automated knowledge base construction*.
[4] Alec Radford et al. 2021. Learning transferable visual models from natural language supervision.

Multimodal Self-supervised Training: CLIP-BERT

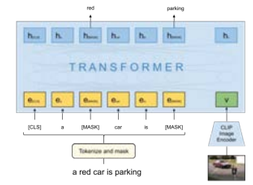

We propose a novel visual-and-language model denoted **CLIP-BERT**, where the image encoder of the pre-trained CLIP [4] model is used to represent the image before appended to the input of a BERT-base model. We train this model using MLM on our vision-and-language dataset, and seek to evaluate how this affects performance on the memory colors task.

Task: Memory colors

With the help of 11 human annotators, we have created a dataset of 109 common objects and their memory colors, with high annotator agreement.

* Equal contribution

Querying strategies

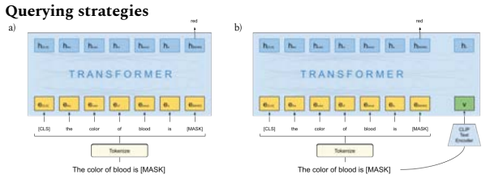

We compare two strategies for querying the language models for this visual knowledge, through [MASK] token prediction.
a) Implicit: The visual knowledge is retrieved from the trained parameters of the language model.

b) Explicit: We use the text encoder of the CLIP model to "imagine" a visual representation of the query text and append to the transformer input. This way the visual knowledge can partly be retrieved from this input, as well as from the trained model parameters.

Results

- ①: The original (text-only) BERT-base performs poorly on this task, close to majority baseline
- ① vs ⑥: Continue MLM training of BERT-base on the text part of our multimodal dataset improves performance slightly despite filtering
- ⑥ vs ⑧: Adding images to training improves performance significantly, showing effective visual knowledge transfer!
- ③ vs ④ and ⑦ vs ⑧: The explicit querying strategy performs better than the implicit

Training	Model	Accuracy	
	Random baseline	0.091 ± 0.026	
	Majority baseline	0.229 ± 0.000	
	Human baseline	0.937 ± 0.051	
None	BERT-base	0.252 ± 0.102	①
Unfiltered	BERT-base	0.724 ± 0.112	②
	CLIP-BERT		
	implicit	0.744 ± 0.080	③
	explicit	0.870 ± 0.086	④
	images	0.876 ± 0.063	⑤
Filtered	BERT-base	0.460 ± 0.083	⑥
	CLIP-BERT		
	implicit	0.541 ± 0.060	⑦
	explicit	0.733 ± 0.098	⑧
	images	0.785 ± 0.055	⑨

Recorded Future

Adaptive Control of Data Center Cooling using Reinforcement Learning

Data centers (DCs) have complex thermal environments which traditional cooling controllers are not able to fully capture.
These controllers are tuned using simple heuristics, which can result in inefficient operation and suboptimal cooling.
A Reinforcement Learning (RL) agent is developed for controlling the cooling fans, featuring system awareness from sensors on servers, in the room, and outdoors.
To realistically run the training and evaluate agent effectiveness, a model of a DC environment is developed.
The model features a CFD simulation of the DC room, heat-generating servers, and fans as well as heat exchangers, compressors, and dry coolers.
Experiments show that the RL agent can outperform baseline agents modeling current best practices in a simple setting where some external disturbances act on the system. Additionally, the RL agent can adapt to larger changes in the environment, such as systems breaking down.

AI MLX | Page 57 B
Heimerson, Albin
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Adaptive Control of Data Center Cooling using Reinforcement Learning

Albin Heimerson, Johannes Sjölund, Rickard Brännvall, Jonas Gustafsson, Johan Eker

Collaborations

Ericsson Reserach Data Center and RICE SICS North are research data centers collaborating with us to develop better methods for data center control. The work was supported by Vinnova grant ITEA3-17002 (AutoDC).

Problem

The total energy consumption by data centers is expected to grow from 1.15% of global energy consumption in 2016 to around 1.86% in 2030 [a] . The International Energy Agency states that one of the innovation gaps in the IT-sector that needs to be filled is applying AI to data centers.

We want to create smarter control algorithms for the cooling systems that are both adaptive and can reason based on more information available in the datacenter.

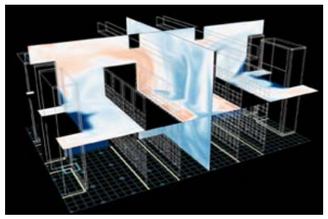

[a] M. Koot and F. Wijnhoven, "Usage impact on data center electricity needs: A system dynamic forecasting model, 2021"

Modelling

A conceptual schematic of the physical model is shown below where the different parts are modeled using different strategies. The "IT space" is modeled using an CFD method called Lattice Boltzmann Method, using an algorithm called Single Relaxation Time Bhatnagar-Gross-Krook, and the model was first presented by Sjölund[a]. Boundary conditions of the servers and cooling systems are based on mathematical modeling of hardware such as the power used by the IT load, the fan speeds, the vapor compression, and the heat exchangers, which all affect the temperatures and air velocities. In previous work[b] a similar model was used with a simpler room simulation.

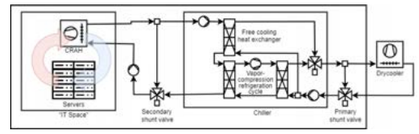

A typical cooling loop for a data center. It can utilize a drycooler to unload the compressor and reduce energy consumption.

[a] J. Sjölund, "Real-time thermal flow predictions for data centers: Using the lattice boltzmann method on graphics processing units for predicting thermal flow in data centers," 2018
[b] A. Heimerson, R. Brännvall, J. Sjölund, J. Eker, and J. Gustafsson, "Towards a Holistic Controller: Reinforcement Learning for Data Center Control", 2021

Reinforcement Learning

The states used is server loads, server outlet temperatures and outdoor temperature. The action is temperature and flow setpoints for the cooling units.
The reward was a weighted sum of two objectives, the energy consumption in the cooling system and a penalty on breaking a temperature threshold of 27° C for the server inlets.

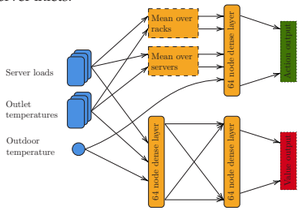

The algorithm needed to be a bit stable while learning, so PPO[a] was used throughout this work.
PPO is an actor-critic algorithm, where the actor maps state to action and the critic maps state to expected reward. The critic tries to learn the bellman equation, $V(s_t) = r_{t+1} + \gamma V(s_{t+1})$, while the actor tries to increase the probability of actions that will have a good future reward according to the critic.
The special thing about PPO is that it doesn't allow the actor to change too much during each update, before it has been evaluated on new data.

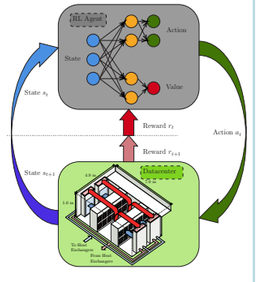

[a] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal Policy Optimization Algorithms", 2017

Results

We compare the RL agent to two baselines, one that is very good at the energy objective by keeping a higher outlet temperature of 22° C, and another that fixes the temperature threshold by keeping the minimal temperature of 18° C.

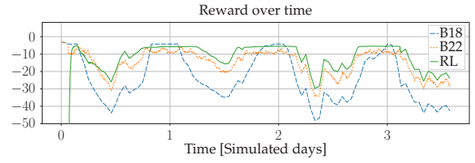

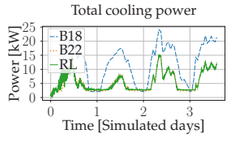

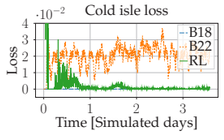

We see that the RL agent matches the energy efficient baseline quite well in energy, while still managing to keep the cold isle loss to a minimum, and achieves the highest reward.
Next we want to see how the agent can adapt to changing conditions such as a cooling unit breaking down $CRAH_0$ loses efficiency, and the cold isle loss will be harder to uphold, but the RL agent will adapt with time to remove the loss (left). The RL agent increases the flow of $CRAH_0$ after the inefficiency is introduced, as well as the flow of $CRAH_1$ which share the same cold aisle.

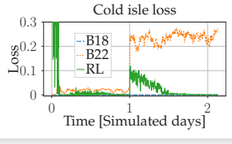

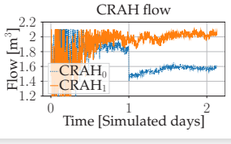

Augmenting Acquisition Functions with User Beliefs for Bayesian Optimization

Bayesian optimization (BO) has become an established framework and popular tool for hyperparameter optimization (HPO) of machine learning (ML) algorithms. While known for its sample-efficiency, vanilla BO can not utilize readily available prior beliefs the practitioner has on the potential location of the optimum. Thus, BO disregards a valuable source of information, reducing its appeal to ML practitioners. To address this issue, we propose \method, an acquisition function generalization which incorporates prior beliefs about the location of the optimum in the form of a probability distribution, provided by the user. In contrast to previous approaches, πBO is conceptually simple and can easily be integrated with existing libraries and many acquisition functions. We provide regret bounds when \method is applied to the common Expected Improvement acquisition function and prove convergence at regular rates independently of the prior. Further, our experiments show that \method outperforms competing approaches across a wide suite of benchmarks and prior characteristics. We also demonstrate that πBO improves on the state-of-the-art performance for a popular deep learning task, with a $12.5\times$ time-to-accuracy speedup over prominent BO approaches.

Jorge, Emilio
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Inferential Induction: A Novel Framework for Bayesian Reinforcement Learning

Bayesian Reinforcement Learning (BRL) offers a decision-theoretic solution to the reinforcement learning problem. While “model-based” BRL algorithms have focused either on maintaining a posterior distribution on models, BRL “model-free” methods try to estimate value function distributions but make strong implicit assumptions or approximations. We describe a novel Bayesian framework,\emph {inferential induction}, for correctly inferring value function distributions from data, which leads to a new family of BRL algorithms. We design an algorithm, Bayesian Backwards Induction (BBI), with this framework. We experimentally demonstrate that BBI is competitive with the state of the art. However, its advantage relative to existing BRL model-free methods is not as great as we have expected, particularly when the additional computational burden is taken into account.

Jorge, Emilio
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

CHALMERS UNIVERSITY OF TECHNOLOGY

zenseact

UiO : University of Oslo

Inria

Inferential Induction:
A Novel Framework for Bayesian Reinforcement Learning

Emilio Jorge*,1 Hannes Eriksson*,1,2 Christos Dimitrakakis*,1,3 Debabrota Basu1,4 Divya Grover1

*Equal contribution [1]Chalmers university of Technology [2]Zenseact [3]University of Oslo [4]Scool, Inria Lille-Nord Europe

The problem

- **Setting**: Bayesian reinforcement learning (BRL).
- **Model-based BRL**: Straightforward formalisation by model distributions.
- **Model-free BRL**: Value function distributions via implicit approximations.

This work

- **Solution**: Derive correct value function distributions directly.
- **Expectation**: Improved modeling could lead to better performance
- **Reality**: BBI is competitive, but more complicated and slower.

Reinforcement learning

An unknown Markov Decision Process (MDP) μ with state s_t, action a_t, reward $r_t \sim P_\mu(r_t \mid s_t, a_t)$, next state $s_{t+1} \sim P_\mu(s_{t+1} \mid s_{t+1}, a_t)$.

Objective: Maximize utility $u_t = \sum_{k=t}^{T} \gamma^t r_t$

The value function V^π of a policy π is

$$V_\mu^\pi(s) \triangleq E_\mu^\pi[u_t \mid s_t = s_0], \qquad a_t \sim P^\pi(a \mid s_t)$$

Bayesian reinforcement learning

The Bayes-optimal solution is

$$\max_\pi E^\pi(u|D)$$

Two main Bayesian approaches

- **Model based:** Belief $\beta \triangleq P(\mu \mid \mathrm{D})$. We can then obtain

$$V_\beta^\pi(s) = \int_\mu V_\mu^\pi(s) dP(\mu \mid \mathrm{D})$$

- **Model free:** Estimate $P(V \mid D)$ directly.

Bayesian value function estimates

Existing model-free BRL algorithms follow the GPTD[1] framework.

- Gaussian process prior over the $P(V)$
- Likelihood function

$$P(D \mid V) \approx \prod_{i=1}^{t} \exp\{-|V(s_i) - r_i - \gamma V(s_{i+1})|^2\}, \; s_i \in D.$$

- At a high level, the inference is :

$$P(V \mid D) = \frac{P(V) P(D \mid V, \hat{\mu}(D))}{P(D)}$$

- Implicitly assumes the empirical MDP $\hat{\mu}(D)$ is correct
  ⇒ ignores model uncertainty.

Inferential induction

We propose a framework, **Inferential Induction**, to calculate the value function distribution $P^\pi(V \mid D_t)$ for policy π, correctly.

Data $D_t = s_1, a_1, r_1, \ldots, s_t, a_t, r_t$
⇒VF posterior $P(V_T|D_t), \ldots, P(V_t|D_t), \ldots, P(V_0|D_t)$.

Calculate the value functions with the inductive integral

$$P^\pi(V_t \mid D_t) = \int_V P^\pi(V_t \mid V_{t+1}, D_t) \, \mathrm{d}P^\pi(V_{t+1} \mid D_t) \qquad \text{(induction)}$$

$$P^\pi(V_t \mid V_{t+1}, D_t) = \int_{\mathcal{M}} \underbrace{P^\pi(V_t \mid \mu, V_{t+1})}_{\text{Bellman operator}} \, \mathrm{d}\overbrace{P^\pi(\mu \mid V_{t+1}, D_t)}^{\text{link distribution}}. \quad \text{(marginalisation)}$$

We introduce **Bayesian Backwards Induction** for calculating $P^\pi(V \mid D_t)$.

- Calculate integral through Monte Carlo sampling of V_{t+1} and μ.
- Define Gaussian kernel relating V_t and utility samples from μ to calculate link distribution $P^\pi(\mu \mid V_{t+1}, D_t)$.
- Importance sampling weights on $P(V_t \mid \mu, V_{t+1})$
- Utilising link distribution may above all be useful when true μ not in model class.

Experiments

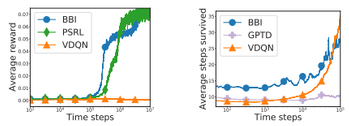

(a) Maze (b) Inverted Pendulum

CPU-time (in seconds) used for each algorithm with 100 policy updates for Chain and 5 updates for Maze.

	BBI	PSRL
Chain	14	5
Maze	921	6

Conclusion

- New framework for Bayesian RL.
- Beats GPTD, but results are not as impressive as we hoped.
- BBI uses $P(\mu \mid D)$ to obtain $P(\mu|V, D)$. Additional methods suggested in the work avoid this.
- It does not appear to be possible to do purely “model-free” Bayesian value function estimation.

References

[1] Yaakov Engel, Shie Mannor, and Ron Meir.
Bayes meets Bellman: The Gaussian process approach to temporal difference learning.
In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 154–161, 2003.

https://arxiv.org/abs/2002.03098 I Can't Believe It's Not Better! Workshop at NeurIPS 2020 emilio.jorge@chalmers.se

Karlsson, Alexander
KTH / SAAB

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Model-Aided Drone Classification Using Convolutional Neural Networks

Abstract - Classifiers using convolutional neural networks (CNNs) often yield high accuracies on samples that come from the same distribution as the training data. In this study we evaluate a CNN classifier's ability to discriminate drones from non-drone targets, such as birds, when they are not represented in the training data. We found that the mean accuracy on such out-of-distribution drones was 78%. By introducing a synthetic drone class, generated from a mathematical model, the out-of-distribution drone accuracy was improved to 86%. When trained on all drone types the mean accuracy over all classes was 90%, and greater than 95% for signal to noise ratios of at least 17.5 dB. The data was collected with a 77 GHz mechanically scanning radar with only 9 ms dwell time.

Karlsson, Alexander
KTH / SAAB

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Model-Aided Drone Classification Using Convolutional Neural Networks

Alexander Karlsson*†, Magnus Jansson*, Mikael Hämäläinen†
*Division of Information Science and Engineering KTH,
†Electronic Warfare Systems, Saab.

Abstract - Classifiers using convolutional neural networks (CNNs) often yield high accuracies on samples that come from the same distribution as the training data. In this study we evaluate a CNN classifier's ability to discriminate drones from non-drone targets, such as birds, when they are not represented in the training data. We found that the mean accuracy on such out-of-distribution drones was 78%. By introducing a synthetic drone class, generated from a mathematical model, the out-of-distribution drone accuracy was improved to 86%. When trained on all drone types the mean accuracy over all classes was 90%, and greater than 95% for signal to noise ratios of at least 17.5 dB. The data was collected with a 77 GHz mechanically scanning radar with only 9 ms dwell time

Background

A drone, or unmanned aerial vehicle (UAV), can be used to deliver harmful payloads such as bombs, cause disturbances at e.g. airports, and collect footage of sensitive sites to name just a few potential threats. The ability to accurately detect and classify drones is therefore of great importance in present and future radar surveillance systems. By classification we refer to the process of distinguishing drones from non-drone targets.

In the literature, the non-drone class is mainly constituted by birds, but depending on what area is under surveillance humans, animals, flying frisbees etc. may also become relevant non-drone targets. Airplanes and other aerial and ground vehicles can often be discarded as false alarms by their radar cross section (RCS) values alone. Birds, humans and drones however may very well have comparable RCS.

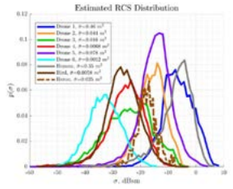

Our main objective is to evaluate a neural network based classifier on drone types not present in the training data. To make the classifier more robust to such out-of-distribution drones we introduce training on both synthetic and real data. The radar we use in this study is a SAAB SIRS 1600 frequency modulated continuous wave (FMCW) radar operating at 77 GHz. This is a mechanically scanning radar and the dwell time is limited to the scan rate and beamwidth, in this case 9 ms.

Method

Data was collected from six different drones, birds (mostly seagulls) and humans. We trained the neural network on each drone, yielding six different scenarios. For reference we also trained on all six drones. The number of classes is N_c.

Layer type	Weight Shape	Strides	Output Shape	activation
Convolutional	F: 10, K: 5×8	[1,1]	$1 \times 143 \times 10$	ReLu
Convolutional	F: 10, K: 1×8	[1,2]	$1 \times 68 \times 10$	ReLu
Convolutional	F: 10, K: 1×8	[1,3]	$1 \times 21 \times 10$	ReLu
Fully Connected	32×210	-	32×1	ReLu
Fully Connected	$N_c \times 32$	-	$N_c \times 1$	Softmax

F: Number of filters/kernels. K: Kernel size. Total number of parameters: $8782 + N_c(32 + 1)$.

The input data is a segment from a scan corresponding to 5 m in range and 2.3° in azimuth.

Each segment is then preprocessed by taking the discrete Fourier transform over azimuth (yielding a Doppler spectrum) and then normalized. Synthetic drone data at range r and FM-sweep p was generated as

$$\tilde{s}_{r,p} = G_{r,p}\left(1 + \rho z_p/\bar{z}\right) + w_{r,p}$$

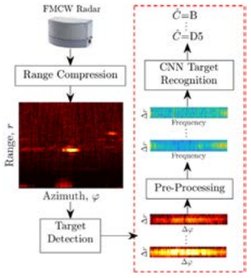

where z_p is the combined propeller return at pulse/FM-sweep p, $w_{r,p}$ is complex white Gaussian noise, $G_{r,p}$ is the combined pulse gain at range index r and azimuth beam gain at index p, and $\bar{z}$ is the rms value of z over all p pulses. For each pulse

$$z_p = \sum_{n=1}^{N_r}\sum_{i=1}^{N_s}\sum_{k=1}^{2} \alpha_k \exp\left(j(-1)^k I(\theta_{n,i})\theta_{n,i}\right) \mathrm{sinc}(\theta_{n,i})$$

$$\theta_{n,i} = 2\pi d_i \sin\left(2\pi\nu p + \phi_n\right)/\lambda$$
$$N_s = \lceil 10d/\lambda \rceil$$
$$d_i = di/N_s$$
$$I(\theta) = \begin{cases} 1, & \text{if } \cos(\theta) \geq 0 \\ -1, & \text{otherwise} \end{cases}$$

where N_p is the number of propellers, λ is the carrier wavelength, ν is the propeller's rotation rate in rounds per second divided by the pulse repetition frequency, φ_n is the initial phase of propeller n, d is the blade length and α_1 and α_2 are weights that determine the symmetry of the Doppler spectrum.

Results

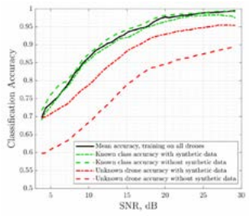

We see that for SNR> 17.5 dB the mean accuarcy is 95% or more when all drones are used in training. When only one drone class is used, the mean accuracy on the known classes is also > 95% for SNR> 17.5 dB whether synthetic data is used or not. The accuracy on the out-of-distribution drones is only > 82% without the synthetic data class and > 91% with synthetic data when the SNR> 17.5 dB.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Kidane, Lidia
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Novel Data Selection Strategies and Associated Machine Learning Algorithms for Cloud Management

Cloud management systems are increasingly using machine learning models for autonomous resource provisioning. These systems need to be frequently calibrated and their models re trained to capture and understand the changing behaviors in the cloud system.
Fundamental assumptions include: workload volume may drastically increase from initial deployment to normal use years later, and that major changes to the machine learning models are expected at times of software and hard ware upgrades and with changing user trends. Models need also be able to deal with periods of volatility in various metrics. A notable problem that arises during those conditions is the change in the statistical properties of the monitoring data. This condition known as concept drift results in incorrect predictions and reduced efficiency of the models. Focus of this research is investigate and propose efficient method that adapt the dynamic behavior of cloud systems. Moreover, since the amount of monitoring data available in these systems is virtually unlimited, our second challenge includes developing new methods and algorithms for selecting important subsets of data for efficient training while being in control of prediction uncertainty.

AI MLX | Page 61 B

Kidane, Lidia
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Novel Data Selection Strategies and Associated Machine Learning Algorithms for Cloud Management

Lidia Kidane, Umeå University
Computing Science
Main Supervisor: Erik Elmroth
Co-Supervisor: Paul Townend

Motivation and Research Goals

Cloud management systems are increasingly using machine learning models for autonomous resource provisioning. These systems need to be frequently calibrated and their models retrained to capture and understand the changing behaviors in the cloud system. Fundamental assumptions include: workload volume may drastically increase from initial deployment to normal use years later, and that major changes to the machine learning models are expected at times of software and hardware upgrades and with changing user trends. Models need also be able to deal with periods of volatility in various metrics. A notable problem that arises during those conditions is the change in the statistical properties of the monitoring data. This condition known as concept drift results in incorrect predictions and reduced efficiency of the models. Focus of this research is investigate and propose efficient method that adapt the dynamic behaviour of cloud systems. Moreover, since the amount of monitoring data available in these systems is virtually unlimited, our second challenge includes developing new methods and algorithms for selecting important subsets of data for efficient training while being in control of prediction uncertainty.

Challenges

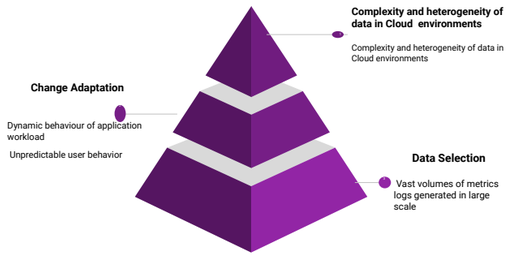

Methods

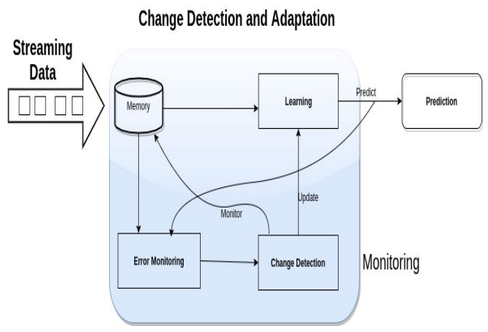

Workload Predictions based on only historical logs does not consider changes in usage patterns or resources.

Concept Drift:

$$p(c_i|X) = \frac{P(c_i)p(X|c_i)}{p(X)}$$

Class definition change p(c|X) while p(X) remains the same. (real drift / concept drift / functional relation change)

- Sudden shifts in workload
- Change in user usage pattern
- Software or hardware upgrade

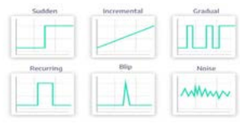

Change Adaptation:

- We implement State-of-the-art concept drift detection algorithms for time series analysis for workload prediction within cloud environment.
- We utilize both **machine learning** time series prediction techniques incorporated with **stream processing** algorithms to update workload prediction models on the fly.

Future Work

Smart data selection strategies

- **Uncertainty Sampling** targets data that is obviously confusing to your model in its current state
- **Diversity sampling** targets data that are gaps in your model's knowledge.
  - Eliminating overlapping of information
- Find metrics that collectively provide accurate, contextual, and insightful information on various aspects of model performance.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Källström, Johan
Linköping University / Saab Aeronautics

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Utility-Based Reinforcement Learning in Support of Simulation-Based Training

Team training in complex domains often requires a substantial amount of resources, e.g., instructors, role-players and vehicles. For this reason, it may be difficult to realize efficient and effective training scenarios in a real-world setting. Instead, intelligent agents can be used to construct synthetic, simulation-based training environments. However, building behavior models for such agents is challenging, especially for the users of the training systems, who typically do not have expertise in artificial intelligence. In this project, we study how reinforcement learning can be used to simplify the process of constructing agents for simulation-based training. By constructing smarter synthetic agents the dependency on human training providers can be reduced, while the availability and quality of training is improved.

AI MLX | Page 62 B

Källström, Johan
Linköping University / Saab Aeronautics

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Utility-Based Reinforcement Learning in Support of Simulation-Based Training

Johan Källström, Linköping University, johan.kallstrom@liu.se
Department of Computer and Information Science

Abstract

Team training in complex domains often requires a substantial amount of resources, e.g., instructors, role-players and vehicles. For this reason, it may be difficult to realize efficient and effective training scenarios in a real-world setting. Instead, intelligent agents can be used to construct synthetic, simulation-based training environments. However, building behavior models for such agents is challenging, especially for the users of the training systems, who typically do not have expertise in artificial intelligence. In this project, we study how reinforcement learning can be used to simplify the process of constructing agents for simulation-based training. By constructing smarter synthetic agents the dependency on human training providers can be reduced, while the availability and quality of training is improved.

The Utility-Based Approach

In this work, we use Multi-Objective Markov Decision Processes (MOMDPs) to model problems with multiple objectives, represented by multiple reward signals. We develop Multi-Objective Reinforcement Learning (MORL)[1] algorithms to solve these problems. MORL allows synthetic agents to learn how to prioritize among multiple, possibly conflicting objectives. The priorities among the objectives of the learning agent are defined by a utility function. Some advantages of MORL compared to standard single-objective reinforcement learning algorithms is that complex non-linear utility functions can be used, a greater degree of flexibility in adapting to changes in goals or utility is achieved, and a more diverse set of solutions can be found[2]. In MORL, there are two major types of optimization criteria: scalarized expected returns (SER) and expected scalarized returns (ESR).

$$V_u^{\pi}(s) = u\left(E\left[\sum_{t=0}^{\infty}\gamma^t \mathbf{r}_t | s_0 = s\right]\right), for\ SER\ criterion$$

$$V_u^{\pi}(s) = E\left[u\left(\sum_{t=0}^{\infty}\gamma^t \mathbf{r}_t\right) | s_0 = s\right], for\ ESR\ criterion$$

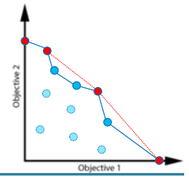

Automating Training Systems

Synthetic, learning agents can be used to automate simulation-based training by replacing or assisting human instructors and role-players. In this work, we use two types of learning agents: A Synthetic Trainer Agent and a Scenario Adaptation Agent[3]. The Synthetic Trainer Agent participates as an actor in training scenarios. By adjusting the utility function of the agent the dynamics of the simulation can be adapted to current training needs, The Scenario Adaptation Agent considers trainees' proficiency in relation to training goals, and then populates training scenarios with human and synthetic agents to maximize improvement in trainees' performance.

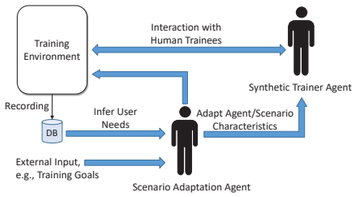

References

1. Hayes, C.F., Rădulescu, R., Bargiacchi, E., Källström, J., et al. (2021). A practical guide to multi-objective reinforcement learning and planning. Submitted to the Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS).
2. Vamplew, P., Smith, B.J., Källström, J., et al. (2021). Scalar reward is not enough. Submitted to the Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS).
3. Källström, J. (2020). Adaptive Agent-Based Simulation for Individualized Training. In Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS).
4. Källström, J., Granlund, R., & Heintz, F. (2021). Design of Simulation-Based Pilot Training Systems using Machine Learning Agents. Submitted to the Aeronautical Journal (AER).
5. Källström, J., & Heintz, F. (2020). Agent Coordination in Air Combat Simulation using Multi-Agent Deep Reinforcement Learning. In proceedings of the 2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC).

Case Study

As a case study, we use a simulation-based air combat training system[4]. In this system, in addition to dealing with multiple training goals, learning agents also need to consider multiple conflicting objectives related to the simulated scenario, e.g., tactical mission goals, resource consumption, and safety. To learn adaptable policies, we use deep neural networks that are conditioned on parameters of the agent's utility function[5]. By letting the agent explore with different utility functions, the agent can learn a set of Pareto optimal policies. Then, after learning, the policy that maximizes the user's current utility can be selected.

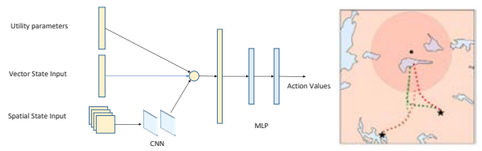

LINKÖPING UNIVERSITY | SAAB | VINNOVA Sweden's Innovation Agency | WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Posterior linearisation smoothing with robust iterations

This work considers the problem of robust iterative Bayesian smoothing in nonlinear state-space models with additive noise using Gaussian approximations. Iterative methods are known to improve smoothed estimates but are not guaranteed to converge, motivating the development of more robust versions of the algorithms. The poster presents Levenberg-Marquardt (LM) and line-search extensions of the classical iterated extended Kalman smoother (IEKS) as well as the iterated posterior linearisation smoother (IPLS). The IEKS has previously been shown to be equivalent to the Gauss-Newton (GN) method. We derive a similar GN interpretation for the IPLS.
Furthermore, we show that an LM extension for both iterative methods can be achieved with a simple modification of the smoothing iterations, enabling algorithms with efficient implementations.

Posterior linearisation smoothing with robust iterations

Jakob Lindqvist

Iterative smoothing

The work presented here is based on an article submitted to Transactions on Signal Processing and is also available as a pre-print [3]
Smoothing is a form of state estimation where a sequence of latent states in a Markov process $x_{1:K} \coloneqq (x_1, x_2, \ldots, x_K)$ are estimated from noisy measurements $y_{1:K} \coloneqq (y_1, y_2, \ldots, y_K)$.
The system is described by a state-space model with additive Gaussian noise

$$\begin{aligned} x_{k+1} &= f_k(x_k) + q_k & q_k &\sim \mathcal{N}(0, Q_k) \\ y_k &= h_k(x_k) + r_k & r_k &\sim \mathcal{N}(0, R_k). \end{aligned} \tag{1}$$

- Rauch-Tung-Stribel (RTS) smoothing computes a closed form solution for linear (affine) motion $f(\cdot)$ and measurement $h(\cdot)$ models.
- General Gaussian RTS smoothers are a family of methods which linearise the state-space models and then perform exact RTS smoothing of the approximate system. Members of this family are extended Kalman smoothing (EKS) and SLR-smoothers.
- Linearisation is done around the best available estimate of the state, commonly the predicted estimates $\hat{x}_{k|k-1}$.
- Iterative smoothers repeatedly perform smoothing and use the posterior estimates of the previous iteration for linearisation points. General Gaussian smoothers have natural iterative extensions, e.g. the IEKS and the IPLS.
- The smoothing problem can be viewed as optimisation problems. The IEKS is equivalent to Gauss-Newton (GN) opt. of the neg. log. likelihood of the state-space in eq. (1) [1].

Posterior linearisation

The choice of linearisation method defines a particular smoother with different properties.

- The EKS and IEKS use first-order Taylor approximation around the estimated mean.
- The IPLS use SLR based on the full distribution [2]:

$$F_k(\hat{x}_k, \hat{P}_k) = \Psi_{f_k}^\top \hat{P}_k^{-1} \tag{4a}$$
$$b_k(\hat{x}_k, \hat{P}_k) = \bar{x}_k - F_k \hat{x}_k \tag{4b}$$
$$\Omega_k(\hat{x}_k, \hat{P}_k) = \Phi_{f_k} - A\hat{P}_k A^{\mathrm{T}}, \tag{4c}$$

$$\begin{aligned} \bar{x}_k &= \int f_k(x_k) p(x_k) \mathrm{d}x_k \\ \Psi_{f_k} &= \int (x_k - \hat{x}_k)(f_k(x_k) - \bar{x}_k)^\top p(x_k) \mathrm{d}x_k \\ \Phi_{f_k} &= \int (f_k(x_k) - \bar{x}_k)(f_k(x_k) - \bar{x}_k)^\top p(x_k) \mathrm{d}x_k. \end{aligned} \tag{5}$$

Iterative smoothers can select a better linearisation point.

- Non-iterative methods use the predicted state $(\hat{x}_{k|k-1}, \hat{P}_{k|k-1})$, which do not take the measurements $y_{k:K}$ into account.
- Ideally, linearisation would be done w.r.t. the posterior distribution of the state.
- Iterative smoothers repeatedly refine the estimates until the linearisation point is approximately chosen w.r.t. the posterior.

Smoothers with robust iterations

Iterative smoothers might diverge and more robust versions are necessary. We take inspiration from general optimisation and use the connection to GN optimisation to propose such modified smoothers.

- We show that the iterations of IPLS of eq. (1) is equivalent to the iterations of GN optimisation of the cost function

$$\begin{aligned} L_{IPLS}^{(i)}(x_{1:K}) = \frac{1}{2}\Big(& (x_1 - \hat{x}_{1|0})^\top \hat{P}_{1|0}^{-1} (x_1 - \hat{x}_{1|0}) \\ & + \sum_{k=1}^{K-1} (x_{k+1} - \bar{x}_k(x_k))^\top \left(Q_k + \Omega_k^{(i)}\right)^{-1} (x_{k+1} - \bar{x}_k(x_k)) \\ & + \sum_{k=1}^{K} (y_k - \bar{y}_k(x_k))^\top \left(R_k + \Gamma_k^{(i)}\right)^{-1} (y_k - \bar{y}_k(x_k))\Big), \end{aligned} \tag{2}$$

- We extend the results of [4] and show that Levenberg–Marquardt (LM) regularised versions of the IEKS and IPLS can be achieved by extending eq. (1) with a pseudo-measurement of the state:

$$\hat{x}_k^{(i)} = x_k + \epsilon_k, \qquad \epsilon_k \sim \mathcal{N}(0, (\lambda^{(i)})^{-1} S_k^{(i)})$$

- We propose line-search versions of the IEKS and IPLS.

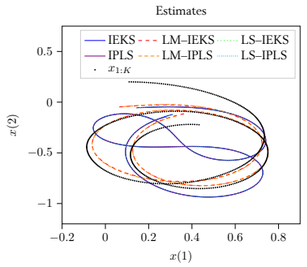

(3) Visualisation of a single realisation of a coordinated turn (CT) experiment with varying bearings only measurements. The measurements at $k = 50, 100, \ldots, 500$ are low noise bearings measurements from a single sensor at $(1, 1)^\top$ Note that for this particular realisation, it is only the LM-regularised smoothers, LM–IEKS and LM–IPLS that estimate the general shape of the true trajectory.

Experimental results

We propose the robust smoothers LM–IEKS and LM–IPLS as well as the line-search based LS–IEKS and LS–IPLS.

We report root mean square error (RMSE) and normalised estimation error squared (NEES) for the CT experiment above. The plots show the metrics averaged over 100 independent simulations with the standard error in the error bars.

- The IPLS based smoothers perform better overall.
- The large spread in the IEKS metrics are caused by some divergent realisations.
- The regularised version perform better than the original methods. In particular, they exhibit faster convergence.

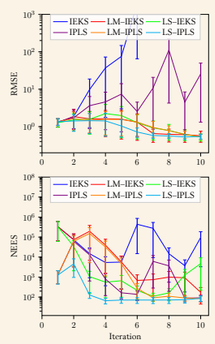

Discussion

- We present robust versions of the IEKS and IPLS smoothers. LM regularisation is achieved by a simple modification of the state-space model and the line-search versions are conceptually simple.
- The IPLS is robust in its original form since it linearises w.r.t. to the full distribution and not just a point estimate, like the IEKS, and benefits less from the robustness modifications.
- Making the linearisation based on the full distribution is more computationally expensive but the added complexity is ameliorated somewhat by the observation that this group of smoothers seems to require fewer iterations to converge to an acceptable trajectory.

References

[1] B. M. Bell. The iterated Kalman smoother as a Gauss–Newton method. *SIAM Journal on Optimization*, 4(3):626 – 636, 1994.

[2] Á F. García-Fernández, L. Svensson, and S. Särkkä. Iterated posterior linearization smoother. *IEEE Transactions on Automatic Control*, 62(4):2056–2063, 2017.

[3] Jakob Lindqvist, Simo Särkkä, Ángel F. García-Fernández, Matti Raitoharju, and Lennart Svensson. Posterior linearisation smoothing with robust iterations. *arXiv e-prints*, page arXiv:2112.03969, December 2021.

[4] S. Särkkä and L. Svensson. Levenberg-Marquardt and line-search extended Kalman smoothers. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5875–5879, 2020.

Lourenço, Inês
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A teacher-student framework for online correctional learning

A classical learning setting is one in which a student collects data, or observations, about a system, and estimates a certain quantity of interest about it. Correctional learning is a type of cooperative teacher-student framework where a teacher, who has knowledge about the system, has the possibility to observe and alter (correct) the observations received by the student in order to improve its estimation. In this poster, we present our formulation of both the batch and online correctional learning problem - while the former is approached as an optimisation problem, for the latter we formulate the MDP and solve it using dynamic programming. Our results show that in both cases the variance of the estimate of the student is reduced with the help of the teacher.
An the end of the poster there is a short summary of our other research areas that we would be very glad to share with you. These go beyond the cooperative setting to adversarial, biologically-inspired, and medical decision-making scenarios.

AI MLX | Page 64 B

Lourenço, Inês
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

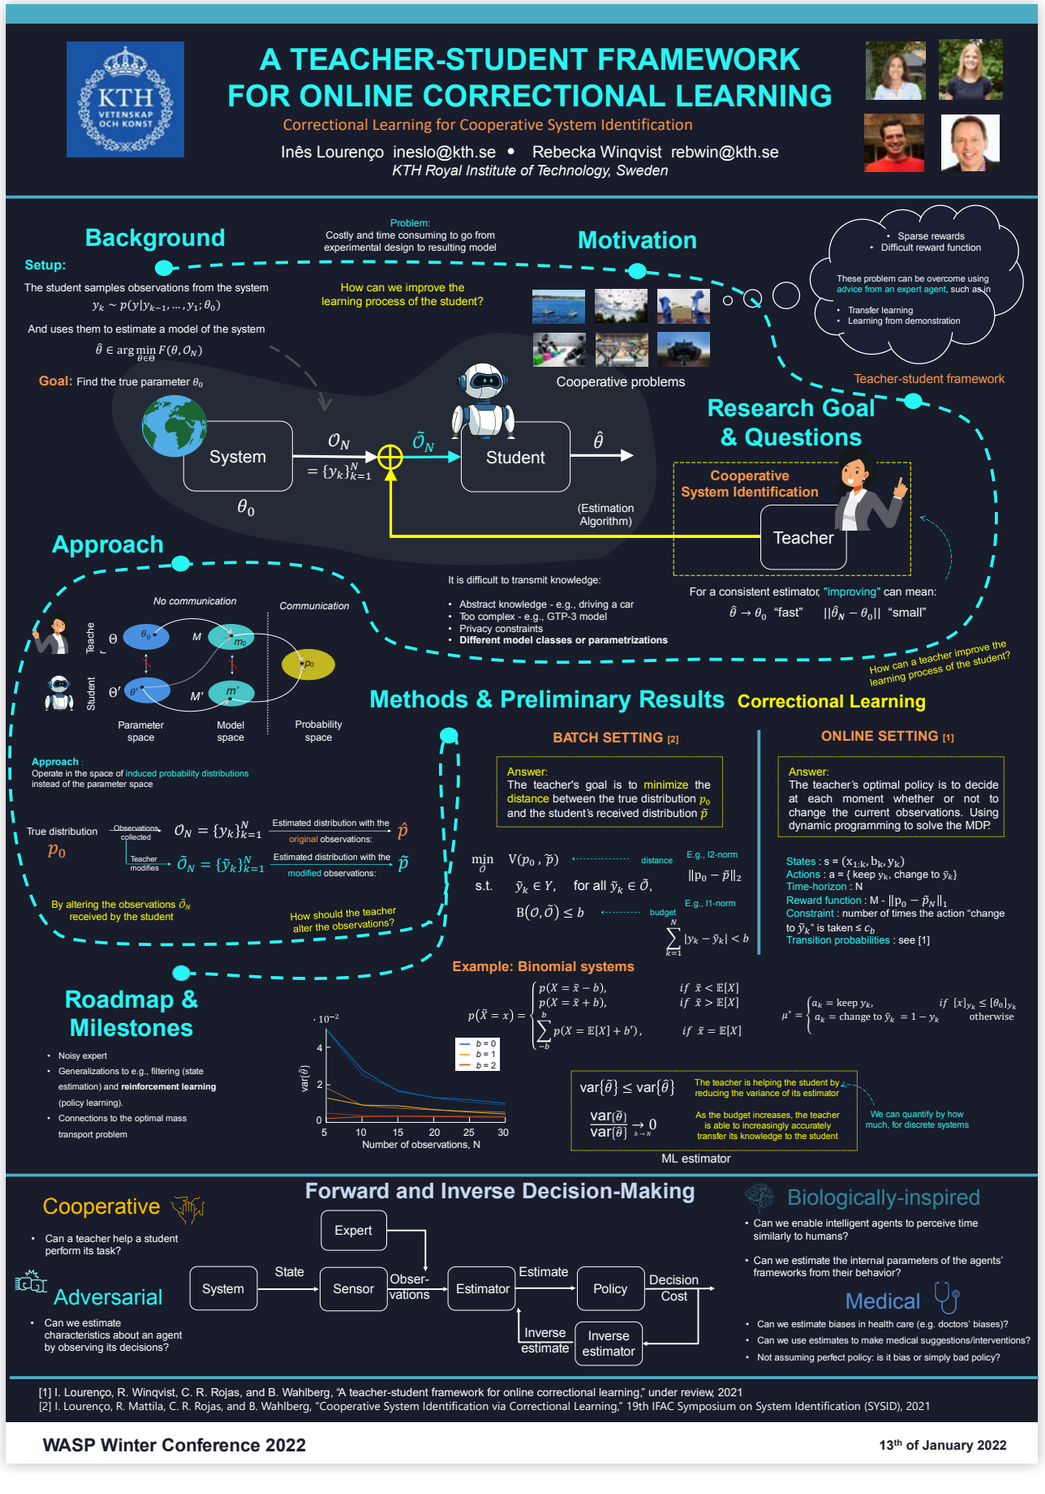

Marti, Miquel
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

An analysis of over-sampling labeled data in semi-supervised learning with FixMatch

Most semi-supervised learning methods over-sample labeled data when constructing training mini-batches. This paper studies whether this common practice improves learning and how. We compare it to an alternative setting where each mini-batch is uniformly sampled from all the training data, labeled or not, which greatly reduces direct supervision from true labels in typical low-label regimes. However, this simpler setting can also be seen as more general and even necessary in multi-task problems where over-sampling labeled data would become intractable. Our experiments on semi-supervised CIFAR-10 image classification using FixMatch show a performance drop when using the uniform sampling approach which diminishes when the amount of labeled data or the training time increases. Further, we analyse the training dynamics to understand how over-sampling of labeled data compares to uniform sampling. Our main finding is that over-sampling is especially beneficial early in training but gets less important in the later stages when more pseudo-labels become correct. Nevertheless, we also find that keeping some true labels remains important to avoid the accumulation of confirmation errors from incorrect pseudo-labels.

Marti, Miquel
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

An analysis of over-sampling labeled data in semi-supervised learning with FixMatch

Miquel Martí i Rabadán[1,2], Sebastian Bujwid[1], Alessandro Pieropan[2], Hossein Azizpour[1] and Atsuto Maki[1]
[1]KTH Royal Institute of Technology, Division of Robotics, Perception and Learning, Sweden
[2]Univrses AB, Sweden

KTH VETENSKAP OCH KONST

UNIVRSES

Motivation & Goals

Most semi-supervised learning methods construct mini-batches sampling its parts from a set of labeled samples and a set of unlabeled ones, then repeating the first set that gets depleted [1,2,3,4], effectively over-sampling labeled data. We aim at studying how this common practice affects learning and compare it to the simplest alternative: uniform sampling of all training data, labeled or not. Alternative approaches to over-sampling can be a requirement in certain learning problems, e.g. in multi-task learning (MTL) samples are said to be labeled with respect to each of the tasks considered. In such case, a sample can not be thought as being "labeled" in general. Thus there is no single, direct approach to over-sample "labeled" data in MTL.

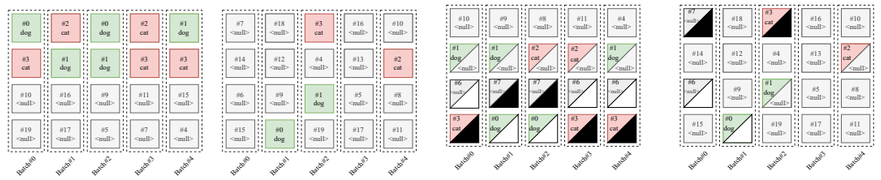

(a) Over-sampling labeled data (b) Uniform sampling (c) MTL ext. of over-sampling (d) MTL uniform sampling

Experiments

We compare both sampling approaches in semi-supervised image classification with CIFAR-10/100 using FixMatch [4] and analyze their training dynamics. (O)ver-sampling - (U)niform sampling - 6x is longer training. Mean/std.dev. of top-1 test accuracy (%) over 5 runs.

CIFAR-10	40 labels	250 labels	4000 labels	All labels
Supervised	36.77 ± 4.48	$59.88 \pm .73$	$87.39 \pm .20$	$96.54 \pm .11$
FixMatch(O)	83.44 ± 6.68	$93.01 \pm .58$	$94.89 \pm .16$	-
FixMatch(U)	73.90 ± 8.04	$91.42 \pm .98$	$94.84 \pm .05$	-
FixMatch(O) 6x	85.40 ± 2.83	$94.40 \pm .78$	-	-
FixMatch(U) 6x	86.59 ± 4.14	$94.00 \pm .54$	-	-
CIFAR-100		2500 labels	10000 labels	All labels
Supervised		47.26	67.58	82.38
FixMatch(O)		62.86	74.61	-
FixMatch(U)		55.53	71.83	-

- **Obs. 1** Both approaches outperform strong supervised baselines.
- **Obs. 2** Over-sampling outperforms uniform sampling. More labels or longer training improve both and reduce the performance gap.
- **Obs. 3** With few labels, uniform sampling is largely inefficient early in training - does not use unlabeled samples, which are the large majority.
- **Obs. 4** Over-sampling is especially beneficial early in training. High-confidence predictions to be used as pseudo-labels are available earlier.

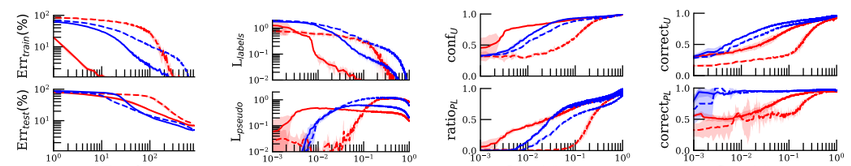

Mean/std.dev. over 5 runs for CIFAR-10: 250 labels, 4000 labels, solid line for over-sampling, dashed for uniform sampling.

Training without true labels in models already producing high rates of correct pseudo-labels eventually leads to collapse to random accuracy, likely due to the accumulation of errors. This hints to the presence of *some* samples with true labels remaining important throughout training to anchor learning to the true task.

Conclusions

The relative success and generality of the simple uniform sampling approach enables its direct use in learning problems where over-sampling labeled data is not possible, such as semi-supervised multi-task learning. Future research should be directed at 1) validating these results for methods substantially different to FixMatch, 2) exploring other sampling approaches for semi-supervised learning and especially in MTL.

References

[1] Semi-supervised Learning with Ladder Networks
A. Rasmus et al. Conference on Neural Information Processing Systems 2015

[2] Virtual Adversarial Training: A Regularization Method for Supervised and Semi-Supervised Learning
T. Miyato et al. IEEE Transactions on Pattern Analysis and Machine Intelligence 2019

[3] MixMatch: A Holistic Approach to Semi-Supervised Learning
D. Berthelot et al. Conference on Neural Information Processing Systems 2019

[4] FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence
K. Sohn et al. Conference on Neural Information Processing Systems 2020

Link to pre-print

With the funding of Swedish Research Council

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Should we Replace CNNs with Transformers for Medical Images?

Convolutional Neural Networks (CNNs) have reigned for a decade as the de facto approach to automated medical image diagnosis, pushing the state-of-the-art in classification, detection and segmentation tasks. Recently, vision transformers (ViTs) have appeared as a competitive alternative to CNNs, yielding impressive levels of performance in the natural image domain, while possessing several interesting properties that could prove beneficial for medical imaging tasks. In this work, we explore whether it is feasible to switch to transformer-based models for medical image classification as well, or if we should keep working with CNNs - can we trivially replace CNNs with transformers? We consider this question in a series of experiments on several standard medical image benchmark datasets and tasks. Our findings show that, while CNNs perform better if trained from scratch, off-the-shelf vision transformers can perform on par with CNNs when pretrained on ImageNet, both in a supervised and self-supervised setting.

Should we Replace CNNs with Transformers for Medical Images?

Christos Matsoukas[1,2,3], Johan Fredin Haslum[1,2,4], Moein Sorkhei[1,2], Magnus Söderberg[3], Kevin Smith[1,2]

[1] Science for Life Laboratory, Stockholm, Sweden
[2] School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden
[3] Cardiovascular, Renal and Metabolism Safety, Clinical Pharmacology & Safety Sciences, R&D AstraZeneca, Gothenburg, Sweden
[4] Bioscience Cardiovascular, Research and Early Development, Cardiovascular, Renal and Metabolism (CVRM), BioPharmaceuticals R&D AstraZeneca, Gothenburg, Sweden

Convolutional Neural Networks (CNNs) have reigned for a decade as the de facto approach to automated medical image diagnosis, pushing the state-of-the-art in classification, detection and segmentation tasks. Recently, vision transformers (ViTs) have appeared as a competitive alternative to CNNs, yielding impressive levels of performance in the natural image domain, while possessing several interesting properties that could prove beneficial for medical imaging tasks. In this work, we explore whether it is feasible to switch to transformer-based models for medical image classification as well, or if we should keep working with CNNs - *can we trivially replace CNNs with transformers?* We consider this question in a series of experiments on several standard medical image benchmark datasets and tasks. Our findings show that, while CNNs perform better if trained from scratch, off-the-shelf vision transformers can perform on par with CNNs when pretrained on ImageNet, both in a supervised and self-supervised setting.

Methods

We compare two mainstream models for classification:
- ResNet50 [2], as CNN representative.
- DeiT-S [3], for ViTs.

We consider three initialization strategies:
- Randomly initialized weights [4].
- Transfer learning using supervised ImageNet [1] pretrained weights.
- Self-supervised pretraining using DINO [5] on the target dataset, after initialization as in (2).

To asses whether vision transformers are able to produce high quality embeddings for segmentation we consider DeepLabV3 [14] and we simply replace its ResNet50 encoder with DeiT-S.

Results

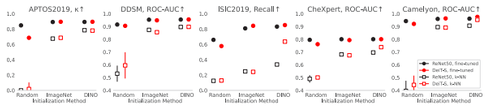

Figure 1: Performance comparison of RESNET50 and DEIT-S, two commonly used CNN-based and ViT-based architectures. The comparison covers several standard medical image classification datasets and different types of initialization including random init, IMAGENET pretraining, and self-supervision using DINO. Performance is measured after fine-tuning on the dataset, as well as using k-NN evaluation without fine-tuning. We report the median over 5 repetitions, error bars represent standard deviation.

References

[1] Deng, Jia, et al. "Imagenet: A large-scale hierarchical image database." 2009 IEEE conference on computer vision and pattern recognition. Ieee. 2009.
[2] He, Kaiming, et al. "Deep residual learning for image recognition." Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.
[3] Touvron, Hugo, et al. "Training data-efficient image transformers & distillation through attention." International Conference on Machine Learning. PMLR. 2021.
[4] He, Kaiming, et al. "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification." Proceedings of the IEEE international conference on computer vision. 2015.
[5] Caron, Mathilde, et al. "Emerging properties in self-supervised vision transformers." arXiv preprint arXiv:2104.14294. 2021.
[6] Kaggle. "Aptos 2019 blindness detection". 2019
[7] Tschandl, Philipp , et al. "The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions." Scientific data 5.1. 2018.
[8] Codella, Noel CF, et al. "Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic)." 2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018). IEEE. 2018.
[9] Combalia, Marc, et al. "BCN20000: Dermoscopic lesions in the wild." arXiv preprint arXiv:1908.02288. 2019.
[10] Sawyer Lee, R., et al. "Curated Breast Imaging Subset of DDSM. The Cancer Imaging Archive.". 2016.
[11] Clark, Kenneth, et al. "The Cancer Imaging Archive (TCIA): maintaining and operating a public information repository." Journal of digital imaging 26.6. 2.013
[12] Codella, Noel, et al. "Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic)." arXiv preprint arXiv:1902.03368 (2019).
[13] Matsoukas, Christos, et al. "Adding seemingly uninformative labels helps in low data regimes." International Conference on Machine Learning. PMLR. 2020.
[14] Chen, Liang-Chieh, et al. "Rethinking atrous convolution for semantic image segmentation." arXiv preprint arXiv:1706.05587 (2017).
[15] Bejnordi, Babak Ehteshami, et al. "Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer." Jama 318.22 (2017): 2199-2210.
[16] Veeling, Bastiaan S., et al. "Rotation equivariant cnns for digital pathology." International Conference on Medical image computing and computer-assisted intervention. Springer, Cham, 2018.
[17] Irvin, Jeremy, et al. "Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison." Proceedings of the AAAI conference on artificial intelligence. Vol. 33. No. 01. 2019.

Model	ISIC2018, IoU ↑	CSAW-S, IoU ↑
DEEPLAB3-RESNET50	0.802 ± 0.012	0.320 ± 0.008
DEEPLAB3-DEIT-S	0.845 ± 0.014	0.322 ± 0.028

Table 1: *Medical image segmentation with* DEEPLAB3 *comparing CNN vs. ViT encoders.*

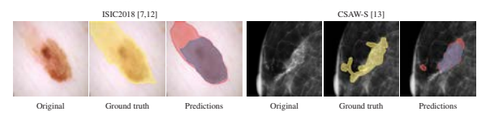

Figure 2: Medical image segmentation results comparing DEEPLAB3-RESNET50 (blue), DEEPLAB3-DEIT-S (red). Ground truth mask appears in yellow. Note that the ViT segmentations tend to do a better job of segmenting distant regions.

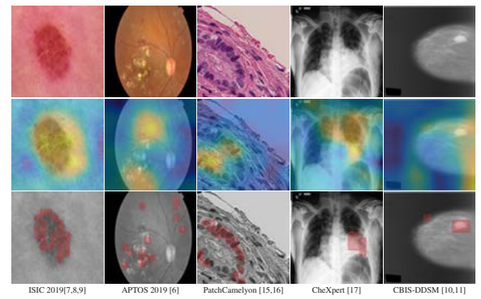

Figure 3: Comparing saliency for RESNET50 (2^{nd} row) and DEIT-S (3^{rd} row) on medical classification. Each column contains the original, a Grad-CAM visualization visualisation for ResNet50 and the top-50% attention map of the CLS token of DEIT-S.

Conclusions

- ViTs can reach the same level of performance as CNNs in small medical datasets, but require transfer learning in order to do so.
- ViTs can outperform CNNs using SSL pre-training when working with limited number of samples, but only marginally.
- ViTs offer built in high-resolution saliency maps that can be used to better understand the model's decisions.

Prototype-Based Off-Policy Evaluation

Before applying a new decision-making policy in safety-critical domains, e.g., in healthcare, we need a reliable estimate of the policy's value. While sampling from this target policy is not possible, we have access to samples from an unknown behavior policy which represents current practice. The problem of evaluating a target policy using data gathered under a behavior policy is known as off-policy evaluation (OPE). Importance sampling (IS) is often used to perform OPE but can provide uncertain value estimates when there are significant differences between the policies. To better diagnose potential problems, we propose estimating the unknown behavior policy for IS using prototype learning. We apply this approach in the evaluation of policies for sepsis treatment, demonstrating that the learned prototypes give a condensed summary of differences between the policies.

Prototype-Based Off-Policy Evaluation

Anton Matsson, Chalmers University of Technology
CHALMERS UNIVERSITY OF TECHNOLOGY
Computer Science and Engineering
Main Advisor: Fredrik Johansson

Abstract

Before applying a new decision-making policy in safety-critical domains, e.g., in healthcare, we need a reliable estimate of the policy's value. While sampling from this *target policy* is not possible, we have access to samples from an unknown *behavior policy* which represents current practice. The problem of evaluating a target policy π using data gathered under a behavior policy μ is known as *off-policy evaluation* (OPE). Importance sampling (IS) is often used to perform OPE but can provide uncertain value estimates when there are significant differences between the policies. To better diagnose potential problems, we propose estimating the unknown behavior policy for IS using prototype learning. We apply this approach in the evaluation of policies for sepsis treatment, demonstrating that the learned prototypes give a condensed summary of differences between the policies.

Background

Importance sampling. Given an observational dataset of trajectories H (sequences of contexts X and actions A) and outcomes R, the standard IS estimator weights the outcomes by the density ratio of the target policy π and the behavior policy μ:

$$\hat{V}_{\text{IS}}(\pi) = \frac{1}{N}\sum_{i=1}^{N}\prod_{t=0}^{T}\frac{p_\pi\left(A_t=a_t^{(i)}\middle|H_t=h_t^{(i)}\right)}{\hat{p}_\mu\left(A_t=a_t^{(i)}\middle|H_t=h_t^{(i)}\right)}r^{(i)}.$$

What can go wrong? The figure shows an example of naïve OPE of two target policies for sepsis management: the AI Clinician [1] and a zero-drug policy. Weighted IS (WIS) was used with different models $\hat{p}_\mu$ of the unknown behavior policy. The zero-drug policy seems to be superior to the behavior policy, followed by physicians in data. But never treating patients with sepsis goes against intuition – can we trust these estimates?

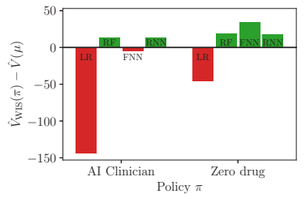

Human evaluation. For $\hat{V}_{\text{IS}}(\pi)$ to be unbiased, overlap must be satisfied. That is, for all t, $p_\pi(A_t|H_t) > 0 \Rightarrow p_\mu(A_t|H_t) > 0$. Because the extent of overlap is unknown when μ is unknown, assessing the quality of $\hat{V}_{\text{IS}}(\pi)$ relies on evaluation by a domain expert. We identify three key questions in such an evaluation:

A. Which observations contribute to the IS estimate?
B. In which situations is overlap violated?
C. If $\hat{V}(\pi) > \hat{V}(\mu)$, what gives π the edge?

References

1. Komorowski, M. et al. The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine*, 24(11): 1716–1720, 2018.
2. Matsson, A. and Johansson, F. D. Case-based off-policy policy evaluation using prototype learning. *arXiv preprint arXiv:2111:11113*, 2021.

Results

Using prototypes. When μ is estimated with black-box models, it can be difficult to detect problems with the IS estimate. Instead, we propose performing OPE using a prototype-based estimate of μ [2]. The overall idea is to compute the probability $p_\mu(A|H)$ by comparing the history H to a small set of prototype cases, which are

- representative trajectories from the training data
- automatically selected by the learning algorithm
- interpretable by a domain expert.

Answering A & B. By evaluating μ and π for each of the prototypes, we get a condensed summary of differences between the policies; see the figure for an example. We can identify groups of patients for which the ratio $p_\pi(A|H)/\hat{p}_\mu(A|H) \gg 1$ for certain actions (see prototype 7) and spot violations of overlap (see prototype 3). Note that the zero-drug policy always predicts action (0, 0) with probability 1.

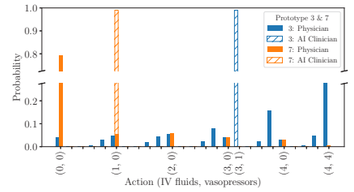

Answering C. The prototypes allows for computing prototype-based contributions to the overall value $V(\pi)$ at each time step. In this way, we can see for which groups of patients a certain policy is most beneficial. The figure shows an example for $t = 2$. Here, $V_t(\pi|J_t = j)$ is the value of π for prototype j at time t and $p_\pi(J_t = j)$ is the probability of being assigned to prototype j at time t.

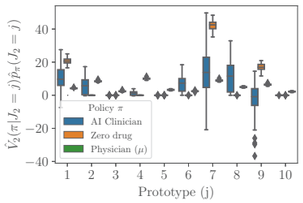

Neural HMMs are all you need (for high-quality attention-free TTS)

Speech synthesis is an application of Generative modelling, where the output is generally conditioned on the input text. This is also referred as Text-To-Speech or TTS Systems. Current experiments on the state of the art speech synthesis systems takes days to realise if the energy hungry GPUs are chunking numbers properly to generate speech and can break down into gibberish. How can we make our current speech / audio experimentation iterations better and save time, effort and energy, without compromising with the quality of synthesised speech? We propose an autoregressive TTS system with a combination of Hidden Markov Models and Deep Neural Networks giving us a smaller size, comparable naturalness, faster iterations, control over speaking rate.

AI MLX | Page 68 B
Mehta, Shivam
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

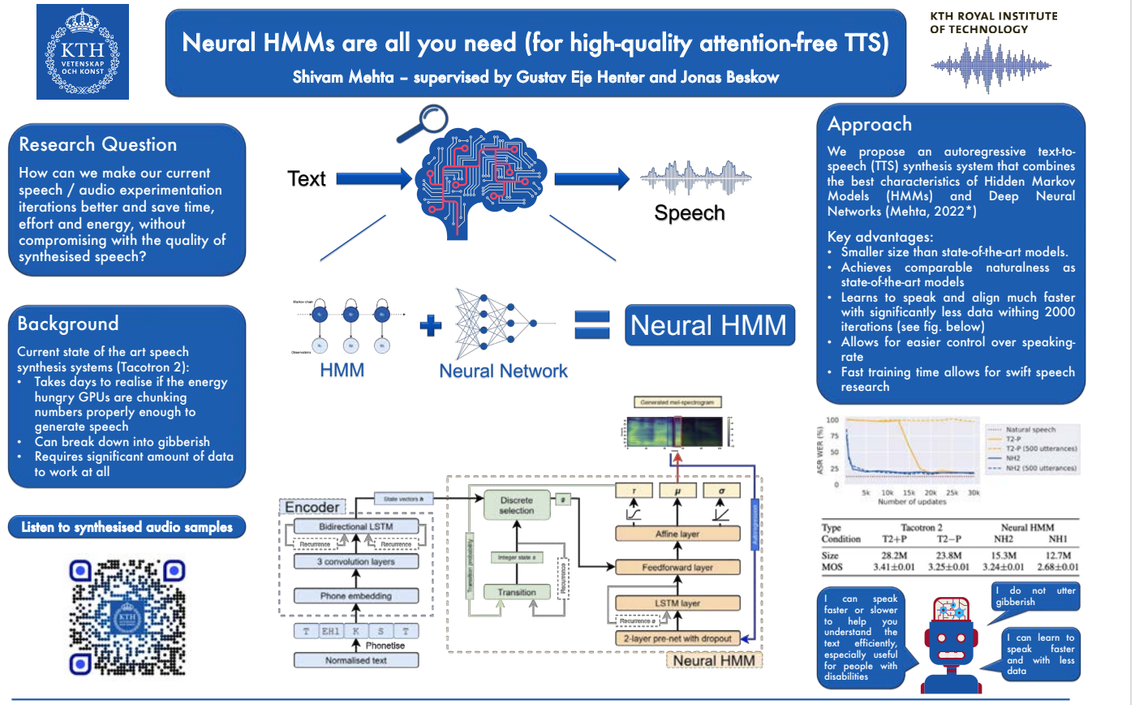

The research is supported by Wallenberg AI, Autonomous Systems and Software Program (WASP) at Department of Speech, Music and Hearing (TMH)

* Neural HMMs are all you need (for high-quality attention-free TTS) - **Shivam Mehta**, Éva Székely, Jonas Beskow, and Gustav Eje Henter – Submitted to ICASSP 2022.

Embed Me If You Can: A Geometric Perceptron

Solving geometric tasks involving point clouds by using machine learning is a challenging problem. Standard feed-forward neural networks combine linear or, if the bias parameter is included, affine layers and activation functions. Their geometric modeling is limited, which motivated the prior work introducing the multilayer hypersphere perceptron (MLHP). Its constituent part, i.e., the hypersphere neuron, is obtained by applying a conformal embedding of Euclidean space. By virtue of Clifford algebra, it can be implemented as the Cartesian dot product of inputs and weights. If the embedding is applied in a manner consistent with the dimensionality of the input space geometry, the decision surfaces of the model units become combinations of hyperspheres and make the decision-making process geometrically interpretable for humans. Our extension of the MLHP model, the multilayer geometric perceptron (MLGP), and its respective layer units, i.e., geometric neurons, are consistent with the 3D geometry and provide a geometric handle of the learned coefficients. In particular, the geometric neuron activations are isometric in 3D, which is necessary for rotation and translation equivariance. When classifying the 3D Tetris shapes, we quantitatively show that our model requires no activation function in the hidden layers other than the embedding to outperform the vanilla multilayer perceptron. In the presence of noise in the data, our model is also superior to the MLHP.

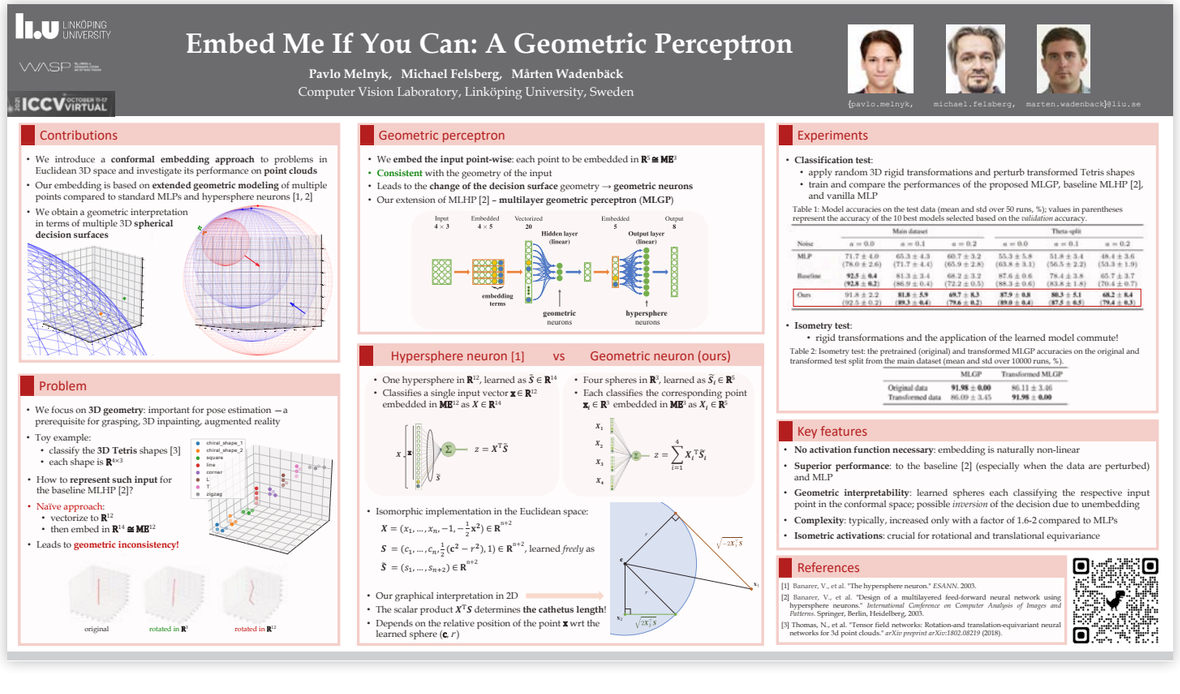

A Language to Bridge the Ethical Gap Between Humans and Machines

AI has raised concerns on whether its use follows ethical principles. We present a new formal language to model ethical requirements. It is statically typed, object-oriented, purely functional, and can be used to create efficient prototypes.

AI MLX | Page 70 B
Mendez, Julian Alfredo
Umeå University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A new object-oriented purely **functional language** formalizes **ethical requirements** for **AI systems**.

A Language to Bridge the Ethical Gap Between Humans and Machines

Julian Alfredo Mendez — julian.mendez@umu.se
Umeå University

Introduction

AI has raised concerns on whether its use follows ethical principles. Our research questions are:

- Can we design a **readable language** to express **ethical requirements** to monitor AI systems?
- Can the specifications be efficiently **prototyped**?

Material and Methods

Why a new language?

The language needs to be not only **expressive** to model ethical requirements, but also **easier** and **clearer** than other similar languages, to **agree on ethical requirements**.

Language properties to model ethical requirements:

the language should ...	we make it ...
be easy to understand	**of simple syntax**
be applicable to many domains	**expressive**
prevent side effect errors	**purely functional**
have powerful design tools	**object-oriented**
detect errors in compile time	**statically typed**
be efficient and easy to integrate	**JVM**

Results and Discussion

We designed a new language that provides:

- **Classes** (to model entities)
- **Class extensions** (Open-Close Principle)
- **Constants and functions** (inside classes)
- **Standard arithmetic operations**
- **Standard basic types**
- **Lambda expressions**
- **Class constructors** (to instantiate objects)
- **Pattern matching** (for class deconstruction)

The descriptions are translated to Scala code, and then to Java Virtual Machine (JVM) bytecode.

Example

```
class RequirementMonitor = {

  has pricing_agent: PricingAgent

  get_price (customer: Customer, flight: Flight, date_in_days: Int): Int =
    pricing_agent.get_price (customer, flight, date_in_days)
}

class Report2 (compliant: Boolean, old_price: Int, new_price: Int)

class Requirement2Monitor
  extends RequirementMonitor = {

  acceptable_yearly_increase = 1.25

  get_report (customer: Customer, flight: Flight, date_in_days: Int): Report2 =
    let
      old_price = get_price (customer, flight, get_a_year_before (date_in_days) )
      new_price = get_price (customer, flight, date_in_days)
    in Report2 (new_price <= old_price * acceptable_yearly_increase,
        old_price, new_price)

  get_a_year_before (date_in_days: Int): Int =
    date_in_days - 365
}
```

State of the Art

We compared available programming languages:

language	version	A	B	C	D	E	F
Agda	2.6.2	Yes	Yes	No	Yes	No	10^9
Clojure	1.10.3	Yes	No	No	No	Yes	10^9
Coq	8.13.2	Yes	Yes	No	Yes	No	10^7
Haskell	8.6.5	Yes	Yes	No	Yes	No	10^8
Idris (2)	0.4.0	Yes	Yes	No	Yes	No	10^{10}
OCaml	4.08.1	Yes	No	Yes	Yes	No	10^9
Prolog	7.6.4	No	No	No	No	No	10^8
Python	3.8.10	No	No	Yes	No	No	10^8
Scala	3.0.2	Yes	No	Yes	Yes	Yes	10^{10}
our language	0.12.0	Yes	Yes	Yes	Yes	Yes	10^{10}

References: **A.** dominantly functional **B.** purely functional (no imperative features) **C.** object-oriented **D.** statically typed **E.** JVM integration **F.** repetitions in 30 s

Conclusion and Future Work

This language is very expressive and can model ethical requirements. We plan to include verification of pieces of code using Coq. Its **readability**, its **expressiveness**, and its **self-consistency** help to bridge the ethical gap between humans and machines.

umu.se/staff/julian-mendez

UMEÅ UNIVERSITY

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Machine learning for improved decision making based on historical data

Historical Data: Simulators make unique benchmarks for causal effect estimation since they do not rely on unverifiable assumptions or the ability to intervene on real-world systems, but are often too simple to capture important aspects of real applications. We propose a simulator of Alzheimer's disease aimed at modeling intricacies of healthcare data while enabling benchmarking of causal effect and policy estimators. We fit the system to the Alzheimer's Disease Neuroimaging Initiative (ADNI) dataset and ground hand-crafted components in results from comparative treatment trials and observational treatment patterns. The simulator includes parameters which alter the nature and difficulty of the causal inference tasks, such as latent variables, effect heterogeneity, length of observed history, behavior policy and sample size. We use the simulator to compare estimators of average and conditional treatment effects.

Decision Making: In personalized medicine, we can assume that there exists a patient latent state $Z\in{R^d}$ which we aim to learn from historical data, and once we learn it, we can recommend personalized near-optimal treatments. In searching for treatments, we usually are constrained in the number of trials that we can perform. We formulate a problem of near-optimal treatment search in a latent bandits setting where: $A\in \{1, ..., k\}, \mu_Z \in R^k$, and budget $T = \beta K$, where β is reasonably small. We apply pure exploration theory to investigate how to present a fixed confidence for near-optimal treatments with a budget in the order $T = \beta K$, where β is reasonably small. We run experiments and simulations with realistic Alzheimer's data from the ADCB environnment.

Machine learning for improved decision making in healthcare with historical data

Newton Mwai Kinyanjui, Fredrik D. Johansson,
Computer Science and Engineering Department, Chalmers

Motivation & Research goals

Healthcare organizations are eager to improve decision making using machine learning applied to records of past decisions and outcomes. Electronic healthcare records are constantly updated with decisions on tests, treatments, procedures and drug prescriptions. If used appropriately, machine learning has the potential to use this data to personalize and improve medicine. Key challenges are a) access to highly realistic observational simulation data in healthcare, and b) ensuring that machine learning systems do not pick up on associations that are not causally related to the results of decisions, to ensure robust decisions when the systems are applied to new problems or new domains.

Decision Making

Personalized medicine with observational data and latent bandits*

In personalized medicine, we can assume that there exists a patient latent state $Z \in \{R^d\}$ which we aim to learn from historical data, and once we learn it, we can recommend personalized near-optimal treatments. As an example of a latent state, it is believed that there are multiple subtypes of Alzheimer's disease which could respond differently to the same treatment.

In searching for treatments, we are usually constrained in the number of trials that we can perform. We formulate a problem of near-optimal treatment search in a latent bandit setting where: treatments (arms) A ∈ {1, ..., k} and budget T = βK, where β is reasonably small.

We apply pure exploration theory to investigate how to present a fixed confidence for near-optimal treatments with a budget in the order T = βK, where β is reasonably small. We run experiments and simulations with realistic Alzheimer's data from the Alzheimer's Disease Causal estimation Benchmark (ADCB) environment.

Goal: Pure exploration with a very small budget

Identify an ϵ – optimal arm (treatment)

$$P(R_{AT} \leq R_A - \epsilon) \leq \delta$$

Subject to T = βK, where β is reasonably small

How does latent state information help us identify such an arm?

** Project ongoing*

Historical Data

Alzheimer's Disease Causal estimation Benchmark (ADCB)

Evaluating learned decision-making policies and observational estimates of causal effects is challenging, as it relies on strong assumptions and access to large samples, hence the research community often turns to simulators for benchmarking which often lack realism.

We propose a new benchmark for evaluating estimators of causal effects that combines the strengths of data-driven simulators with those of hand-crafted components by fitting a longitudinal causal model of patient variables to real data and providing tunable parameters for the generated data.

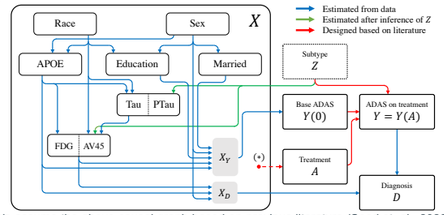

We assume the above causal graph based on previous literature (Sood et. al., 2020) and domain practitioner input. For each continuous(or discrete) attribute, a regression(or classification) model is fit with respect to its parents. We fit sequences autoregressively.

Tunable parameters

μ – Behavior policy ∈ {Random, Diagnosis-based, Hernandez-based (Hernandez et al.(2010))}

ϵ – Overlap parameter ∈ [0, 1]

γ – Treatment Effect Heterogeneity ≥ 0

N – Number of Samples ≥ 0

T – Sample trajectory length (history length) ∈ {0, 12, 24, …, 120}

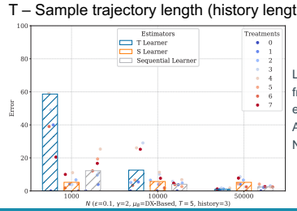

Left: Results of using generated data from ADCB in comparing causal effect estimators based on Conditional Average Treatment Effects Error for Number of patient samples, N

References

1. Kinyanjui, Newton Mwai, and Fredrik D. Johansson. "ADCB: An Alzheimer's disease benchmark for evaluating observational estimators of causal effects." ML4H 2021 - extended abstract *arXiv:2111.06811* (2021).
2. Meemansa Sood, et al. "Realistic simulation of virtual multi-scale, multi-modal patient trajectories using bayesian networks and sparse auto-encoders". Scientific reports, 10(1):1–14, 2020.
3. Hernandez, Santiago, et al. "Pharmacological treatment of Alzheimer's disease: effect of race and demographic variables." Journal of Alzheimer's Disease 19.2 (2010): 665-672.

Norlund, Tobias
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Transferring Knowledge from Vision to Language: How to Achieve it and how to Measure it?

Large language models are known to suffer from the hallucination problem in that they are prone to output statements that are false or inconsistent, indicating a lack of knowledge. A proposed solution to this is to provide the model with additional data modalities that complements the knowledge obtained through text. We investigate the use of visual data to complement the knowledge of large language models by proposing a method for evaluating visual knowledge transfer to text for uni- or multimodal language models. The method is based on two steps, 1) a novel task querying for knowledge of memory colors, i.e. typical colors of well-known objects, and 2) filtering of model training data to clearly separate knowledge contributions. Additionally, we introduce a model architecture that involves a visual imagination step and evaluate it with our proposed method. We find that our method can successfully be used to measure visual knowledge transfer capabilities in models and that our novel model architecture shows promising results for leveraging multimodal knowledge in a unimodal setting.

Norlund, Tobias
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Transferring Knowledge from Vision to Language: How to Achieve it and how to Measure it?

Tobias Norlund*, Lovisa Hagström*, Richard Johansson
Chalmers University of Technology

Motivation & Summary

Despite the ability of language models to learn and hold large quantities of structural knowledge [1], LMs are are also known to suffer from the **hallucination problem** in that they are prone to output statements that are false or inconsistent, indicating significant knowledge gaps [2]. A hypothesis is that this knowledge in many cases simply is missing in the large text corpora typically used for training the models, due to e.g reporting bias [3]. In such cases certain types of knowledge might also be more readily available in a different data modality.

In this work, we investigate **visual knowledge transfer**, i.e. to which extent language models can incorporate and *textually* express knowledge originating from a visual modality.

We investigate this by constructing a novel cloze-style task testing knowledge of **memory colors** for common objects (such as blood is red, a lemon is yellow etc). We also build a large vision-and-language dataset used for self-supervised training, and by careful filtering we make sure the color information is only available through the images and not the text.

Finally, we compare two strategies for how to most effectively query the trained multimodal language model for this visual knowledge.

We show that language models are able to texually express knowledge obtained from a visual modality, as a result from multimodal self-supervised training.

References
[1] Fabio Petroni et al. Language Models as Knowledge Bases? EMNLP (2019).
[2] Robert Logan et al. 2019. Barack's wife Hillary: Using knowledge graphs for fact-aware language modeling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy. Association for Computational Linguistics.
[3] Jonathan Gordon and Benjamin Van Durme. 2013. Reporting bias and knowledge acquisition. In *Proceedings of the 2013 workshop on Automated knowledge base construction*.
[4] Alec Radford et al. 2021. Learning transferable visual models from natural language supervision.

Multimodal Self-supervised Training: CLIP-BERT

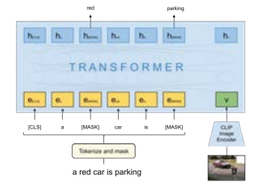

We propose a novel visual-and-language model denoted **CLIP-BERT**, where the image encoder of the pre-trained CLIP [4] model is used to represent the image before appended to the input of a BERT-base model. We train this model using MLM on our vision-and-language dataset, and seek to evaluate how this affects performance on the memory colors task.

Task: Memory colors

With the help of 11 human annotators, we have created a dataset of 109 common objects and their memory colors, with high annotator agreement.

* Equal contribution

Querying strategies

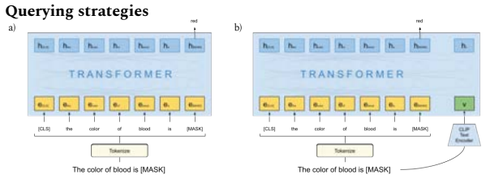

We compare two strategies for querying the language models for this visual knowledge, through [MASK] token prediction.
a) Implicit: The visual knowledge is retrieved from the trained parameters of the language model.

b) Explicit: We use the text encoder of the CLIP model to "imagine" a visual representation of the query text and append to the transformer input. This way the visual knowledge can partly be retrieved from this input, as well as from the trained model parameters.

Results

- ①: The original (text-only) BERT-base performs poorly on this task, close to majority baseline
- ① vs ⑥: Continue MLM training of BERT-base on the text part of our multimodal dataset improves performance slightly despite filtering
- ⑥ vs ⑧: Adding images to training improves performance significantly, showing effective visual knowledge transfer!
- ③ vs ④ and ⑦ vs ⑧: The explicit querying strategy performs better than the implicit

Training	Model	Accuracy	
	Random baseline	0.091 ± 0.026	
	Majority baseline	0.229 ± 0.000	
	Human baseline	0.937 ± 0.051	
None	BERT-base	0.252 ± 0.102	①
Unfiltered	BERT-base	0.724 ± 0.112	②
	CLIP-BERT		
	implicit	0.744 ± 0.080	③
	explicit	0.870 ± 0.086	④
	images	0.876 ± 0.063	⑤
Filtered	BERT-base	0.460 ± 0.083	⑥
	CLIP-BERT		
	implicit	0.541 ± 0.060	⑦
	explicit	0.733 ± 0.098	⑧
	images	0.785 ± 0.055	⑨

Recorded Future

Olmin, Amanda
Linköpings universitet

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Robustness and reliability when training with noisy labels

Algorithms developed for the purpose of handling label noise in supervised training, are commonly evaluated in terms of accuracy. However, robustness in accuracy is not sufficient in applications where reliable uncertainty estimates are critical. For an input-dependent noise model, we investigate the effect of label noise on strictly proper loss functions as well as the set of robust loss functions characterised by noise-insensitive, asymptotic risk minimisers. We find that not only robust, but also strictly proper loss functions offer asymptotic robustness in accuracy. However, neither guarantee that the final model is calibrated. Moreover, both strictly proper and robust loss functions are susceptible to overfitting in practice.

AI MLX | Page 73 B

Olmin, Amanda
Linköpings universitet

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Robustness and reliability when training with noisy labels

Amanda Olmin, Linköping University
Department of Computer and Information Science
Supervisors: Fredrik Lindsten (LiU) and Lennart Svensson (Chalmers)

LINKÖPINGS UNIVERSITET

Motivation

Supervised training of deep learning models is highly dependent upon labelled data. Curating a (possibly large) annotated training data set is costly and time-consuming and the risk of incorporating label noise is imminent. Label noise can hurt model performance by: (i) shifting the asymptotic risk minima towards the conditional distribution over noisy, instead of clean, labels, and (ii) increasing the risk of overfitting (Zhang et al., 2017). While algorithms robust to label noise are commonly evaluated in terms of accuracy, this is not enough if reliable uncertainty estimates are critical. We establish this idea and investigate the effect of label noise on model performance through a critical analysis of robust loss functions.

Simple non-uniform label noise

For noisy label $\tilde{Y} \in \mathcal{Y}$, true label $Y \in \mathcal{Y}$ and input variable $X \in \mathcal{X}$, simple non-uniform label noise (Ghosh et al., 2017) is defined by

$$\mathbb{P}(\tilde{Y} = \tilde{y} \mid Y = y, X = x) = \begin{cases} 1 - \omega(x), & \text{if } \tilde{y} = y \\ \frac{\omega(x)}{K-1}, & \text{otherwise} \end{cases}$$

with $\mathcal{Y} = \{1, \dots, K\}$ and noise parameter $0 \leq \omega(x) < \frac{K-1}{K}$.

Robust loss functions

A loss function ℓ is robust (Ghosh et al., 2017) to label noise if for all asymptotic minimisers f^* of the clean risk, $\mathcal{R}_\ell$, it holds that

$$\tilde{\mathcal{R}}_\ell(f^*) \leq \tilde{\mathcal{R}}_\ell(f), \quad \forall f \in \mathcal{F}$$

where $\tilde{\mathcal{R}}_\ell$ is the risk under the noisy data distribution.

Strictly proper loss functions recover the true conditional $f^*(x) = \mathbb{P}(Y \mid X = x)$ if labels are clean, but are not robust.

Symmetric loss functions are robust under simple non-uniform label noise (Ghosh et al., 2017). They satisfy

$$\sum_{k=1}^{K} \ell(q, k) = C, \quad \forall x \in \mathcal{X},\ \forall q \in \Delta^{K-1}$$

for some constant C.

Training dynamics

We imagine two phases of training a model $f : \mathcal{X} \to \Delta^{K-1}$.

1. The training trajectory "aims" towards the true risk minimiser f^* (or $\tilde{f}^*$ if labels are noisy) and passes "close" to it.
2. The training trajectory diverges and the model overfits to the data.

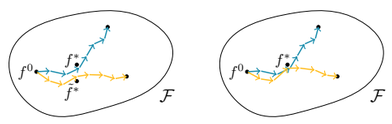

(a) Strictly proper loss, $\tilde{f}^* \neq f^*$

(b) Robust loss, $\tilde{f}^* = f^*$

Robustness does not imply reliability

Under simple non-uniform label noise, we find that strictly proper and symmetric loss functions

- are robust to label noise in accuracy,
- do not result in calibrated models (uncertainty estimates unreliable),

The robustness condition is not sufficient if reliable uncertainty estimates are critical.

- are both susceptible to overfitting in practice.

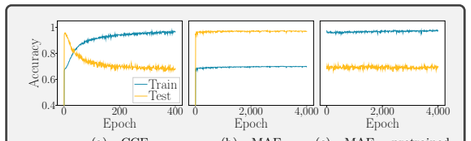

(a) CCE (b) MAE (c) MAE, pretrained

Figure: Experiments on MNIST with simple non-uniform label noise, $\omega(x) = 0.3\ \forall x$. The empirical robustness of Mean Absolute Error (b) has been attributed to its theoretical robustness (Ghosh et al., 2017). However, the theory does not account for overfitting. In practice, the overfitted model (c) achieves a smaller loss (d).

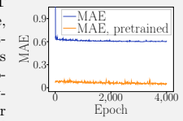

(d) MAE loss

Calibration-based strictly properness

Let $\mathcal{F}_\mathcal{C} \subseteq \mathcal{F}$ be the set of calibrated models. The loss function ℓ, with asymptotic risk minimisers f^*, is calibration-based strictly proper if

$$f^* \in \mathcal{F}_\mathcal{C}, \quad \forall f^* \in \mathcal{F},$$

for all $\mathbb{P}(Y \mid X)$ and for all input distributions μ_X.

A loss function that is both robust and calibration-based strictly proper will preserve accuracy and ensure reliability.

References

Ghosh, A., Kumar, H., and Sastry, P. S. (2017). Robust loss functions under label noise for deep neural networks. In *AAAI Conference on Artificial Intelligence*.

Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2017). Understanding deep learning requires rethinking generalization. In *ICLR*.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Deep Gaussian Markov Random Fields for General Graphs

Machine learning methods on graphs have proven useful in many applications due to their ability to handle generally structured data. The framework of Gaussian Markov Random Fields (GMRFs) provides a principled way to define Gaussian models on graphs by utilizing the sparsity structure. We propose a highly scalable method for defining GMRFs on general graphs based on the layer structure of Deep GMRFs. The parameters of the resulting model can be trained efficiently using variational inference and existing software for Graph Neural Networks. For a Gaussian likelihood exact Bayesian inference is possible for predictions. The usefulness of the model and the multi-layer structure is verified by experiments on a number of synthetic and real world datasets.

Deep Gaussian Markov Random Fields for General Graphs[a]

Joel Oskarsson, joel.oskarsson@liu.se, Linköping University
Dept. of Computer and Information Science, Division of Statistics and Machine Learning
Supervisors: Fredrik Lindsten (LiU), Per Sidén (LiU, Arriver) and Jose M. Peña (LiU)

[a]Ongoing research

LINKÖPING UNIVERSITY

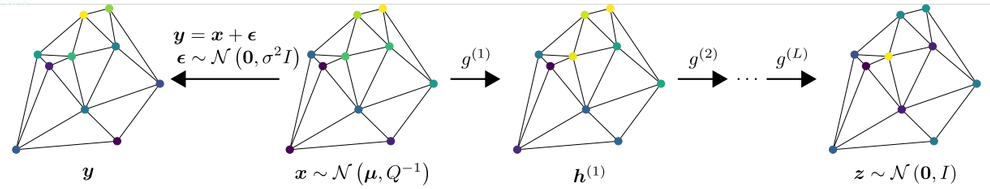

Background: Gaussian Markov Random Fields

Gaussian Markov Random Fields (GMRFs)

- Graphical models where the nodes $\boldsymbol{x}$ are jointly Gaussian [1]
- Conjugate prior to Gaussian likelihood $p(\boldsymbol{y}|\boldsymbol{x}) = \mathcal{N}(\boldsymbol{y}|\boldsymbol{x}, \sigma^2 I)$
  - Posterior $p(\boldsymbol{x}|\boldsymbol{y})$ analytically tractable, also a GMRF!

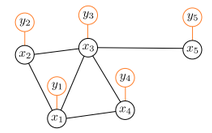

Deep GMRFs (DGMRFs) [2]

- GMRF $\boldsymbol{x}$ defined implicitly by linear map g

$$\boldsymbol{z} = g(\boldsymbol{z}) = G\boldsymbol{x} + \boldsymbol{b}, \quad \boldsymbol{z} \sim \mathcal{N}(\mathbf{0}, I) \tag{1}$$

- g defined as a combination of L simpler layers $g = g^{(L)} \circ g^{(L-1)} \circ \cdots \circ g^{(1)}$
- Originally restricted to lattice-graphs (image-structured data)

Method: Graph DGMRF

- Consider graph $\mathcal{G}$ with adjacency matrix A and degree matrix $D = \mathrm{diag}([d_1, d_2, \ldots, d_N]^\intercal)$
- Define a layer construct $\boldsymbol{h}^{(l)} = g^{(l)}(\boldsymbol{h}^{(l-1)}) = G^{(l)}\boldsymbol{h}^{(l-1)} + \boldsymbol{b}^{(l)}$

$$G^{(l)} = \alpha_l D^{\gamma_l} + \beta_l D^{\gamma_l - 1} A \tag{2}$$
$$\boldsymbol{b}^{(l)} = b_l \mathbf{1} \tag{3}$$

  - b_l, α_l, β_l and γ_l are trainable parameters

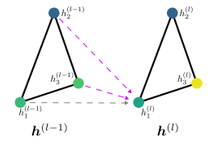

- Closely corresponds to Graph Neural Network (GNN) formulation. Existing frameworks can be utilized for automatic differentiation and GPU acceleration.
- Variational training, maximizing the Evidence Lower Bound (ELBO)
- Training requires efficient evaluation of $\log\left|\det\left(G^{(l)}\right)\right|$. We develop two solutions:
  - Exact method based on pre-computing the eigenvalues of $D^{-1}A$
  - Scalable approximate method based on reformulating log-determinant as a power series

References

[1] H. Rue and L. Held. *Gaussian Markov random fields: theory and applications*. Number 104 in Monographs on statistics and applied probability. Chapman & Hall/CRC, 2005.

[2] P. Sidén and F. Lindsten. Deep gaussian markov random fields. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.

Experiments

- Experiments on inpainting problems where we train on 50% of nodes and evaluate on remaining 50%

Synthetic GMRF Dataset

- Sample from known GMRF created as random mix of multiple DGMRFs
- 1-5 layer DGMRFs trained and evaluated

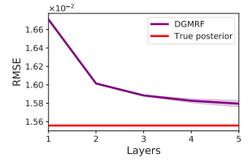

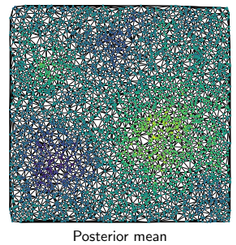

Posterior mean

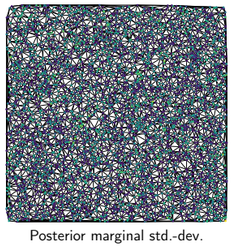

Posterior marginal std.-dev.

Wikipedia Graph

- 11 631 Wikipedia pages as nodes
- Edges are links between pages
- Target $\boldsymbol{y}$ is log average monthly traffic
- CRPS, probabilistic metric taking uncertainty into account

	RMSE	CRPS
Baselines		
Graph GP	2.169	1.251
DGP (GNN)	1.293	0.769
I-GMRF	1.526	0.939
DGMRF		
1 Layer	1.312	0.704
3 Layer	1.230	0.652
5 Layer	**1.179**	**0.620**

$\Rightarrow$ Additional experiments on weighted graphs, use of node features and scaling to $> 10^5$ nodes

Summary

- Layered Gaussian model for graph-structured data
- Efficient variational training using Graph Neural Network framework
- Exact Bayesian inference gives principled uncertainty estimates

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Patil, Minal Suresh
Umeå univeristet

Towards Explainable Agency in Multi-Agent Systems Using Constraint Programming

Logical reasoning is a fundamental aspect of human behaviour, and this is an important criterion to build human-like reasoning in intelligent autonomous multi-agent systems. So far, the field of knowledge representation and reasoning has employed logic-based symbolic techniques to mimic the challenging task of incorporating human-like reasoning in multi-agent systems. However, the field of machine learning has shown increasing interest to take on this challenge. In this research, we describe a methodology that is based on Constraint Logic Programming that enables autonomous agents to generate explanations and logic-based reasoning from a knowledge base and monitor how explanations advance over time. Whilst this preliminary work addresses key limitations such as scalability and adaptability, we strongly emphasise the need for logic-based reasoning in multi-agent systems for interpretability and transparency in their behaviour.

Patil, Minal Suresh
Umeå univeristet

Towards Explainable Agency in Multi-Agent Systems Using Constraint Programming

Minal Suresh Patil, Umeå universitet
Main Supervisor: Kary Främling

Motivation & Research Goals

Logical reasoning is a fundamental aspect of human behaviour, and this is an important criterion to build human-like reasoning in intelligent autonomous multi-agent systems. So far, the field of knowledge representation and reasoning has employed logic-based symbolic techniques to mimic the challenging task of incorporating human-like reasoning in multi-agent systems. However, the field of machine learning has shown increasing interest to take on this challenge. In this research, we describe a methodology that is based on Constraint Logic Programming that enables autonomous agents to generate explanations and logic-based reasoning from a knowledge base and monitor how explanations advance over time. Whilst this preliminary work addresses key limitations such as scalability and adaptability, we strongly emphasise the need for logic-based reasoning in multi-agent systems for interpretability and transparency in their behaviour.

Method

The optimal linkage between explainability and causation, which is the cornerstone to effective human-agent explainability one of the most important ultimate goals of explainable AI systems. Explainability pertains to a system that has the ability to explain itself to others in a natural language. In other words, a system should be able to communicate the reasoning behind its decisions [1]. Cause and effect is crucial for making ethical decisions.

<u>The logical way:</u>
Logic can assist with technical aspects of the problems. As a consequence, we propose a Constraint Logic Programming (CLP) for machine explainability.. It's an approach that is indeed innately interpretable and it is easy to incorporate domain knowledge.

Results

Particularly compared to traditional techniques, CLP has a significant advantages:

- Because logic programs are expressive, CLP systems can understand intricate relational theories.
- Domain-Knowledge can also be used by CLP systems to learn domain specific constraints.
- Because domain knowledge is used as a type of inductive bias, CLP systems might generalize from a small number of examples.
- CLP systems are designed to promote continuous and transfer learning.

References

1. M.S. Patil. Explainability in Autonomous Pedagogically Structured Scenarios. In Proceedings of Workshop on Explainable Agency in Artificial Intelligence at 36th Association for the Advancement of Artificial Intelligence (AAAI), 2022

Paul, Sudipta
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A Subspace matched framework for Federated Learning

In the setup of Federated Learning[1], and more particularly, when using the FedAvg algorithm, all the weights of the locally trained model get averaged in each round. In this setup, the global model is being sent to the ready local devices and they, later, send back the locally trained model to the central server for the next averaging step. Though this setup claims to ensure the security and privacy of the local data at the devices, some research has shown that this algorithm is still prone to membership attack[2], model reconstruction attack and backdoor attack. Our approach introduces a new federated learning framework T\'opos-FL on the basis of subspace and correlation analysis upon the layers of the machine learning models. It mitigates several drawbacks of FedAvg. In particular, model and data reconstruction attacks, and membership attack. In our approach, a conjugated view of the layers is being transferred to the central server where the update is subject to maximizing the correlation between the global and the local models.

Paul, Sudipta
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A Subspace matched framework for Federated Learning

Sudipta Paul, PhD Student, Umeå University
Department of Computing Science
Supervisors: Prof. Vicenç Torra and Lili Jang (UMU)

UMEÅ UNIVERSITET

Motivation & Research Goals

In the setup of Federated Learning[1], and more particularly, when using the FedAvg algorithm, all the weights of the locally trained model get averaged in each round. In this setup, the global model is being sent to the ready local devices and they, later, send back the locally trained model to the central server for the next averaging step. Though this setup claims to ensure the security and privacy of the local data at the devices, some research has shown that this algorithm is still prone to membership attack[2], model reconstruction attack and backdoor attack. Our approach introduces a new federated learning framework Tópos-FL on the basis of subspace and correlation analysis upon the layers of the machine learning models. It mitigates several drawbacks of FedAvg. In particular, model and data reconstruction attacks, and membership attack. In our approach, a conjugated view of the layers is being transferred to the central server where the update is subject to maximizing the correlation between the global and the local models.

Methods

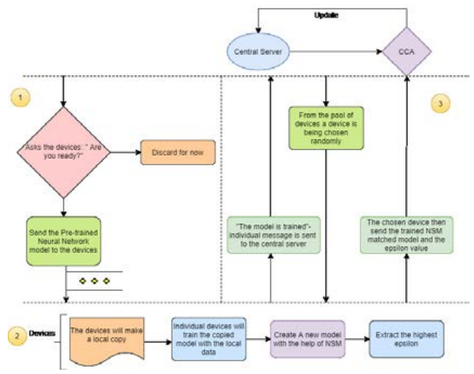

Algorithm 3: Tópos-FL.
The K clients are indexed by $1, 2, 3, 4, \ldots, k, \ldots$
b is the batch size, η is the learning rate, e is the number of epoch, DB_k is the data at device k

At the Server:
initialization w_c;
This is the model at the central server
for *each round* $t = 1,2,3, \ldots$ **do**
  $D_t \leftarrow$ set of ready devices;
  for *each client device* D_k **do**
    Receive $< w_k, \epsilon^* >$ from the client ;
    $w_c \leftarrow w_c + \eta SVCCA(w_c, < w_n, \epsilon^* >)$;
  end
  Send w_c to clients ;
end
At the Clients: *(Run on each client device in parallel)*
$w'_c \leftarrow w_c$ *(Make a copy of the central model)*;
for *each local epoch* $\in e$ **do**
  for *batch size* $\in b$ **do**
    w'_c is being trained by DB_k
  end
end
$< w_n, \epsilon^* > \leftarrow NSM(w_c, w'_c)$;
Send to the server $< w_n, \epsilon^* >$;

Selected Results

Result of the reconstruction attack in the Tópos-FL setup for CIFAR-10

Result of the reconstruction attack in the FedAVG setup for CIFAR-10

Result of the membership attack in the Tópos-FL setup for CIFAR-10

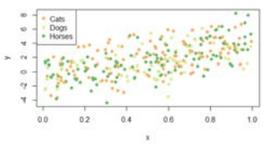

Result of the membership attack in the FedAVG setup for CIFAR-10

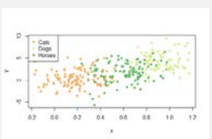

References

[1] Advances and Open Problems in Federated Learning
Peter Kairouz et al.
arXiv preprint arXiv:1912.04977

[2] Membership Inference Attacks Against Machine Learning Models
Shokri, Reza ; Stronati, Marco; Song, Congzheng and Shmatikov, Vitaly
2017 IEEE Symposium on Security and Privacy (SP)

Deep Weighted MMSE Downlink Beamforming

The weighted minimum mean square error (WMMSE) algorithm was proposed to provide a locally optimum solution to the otherwise NP-hard weighted sum rate maximization beamforming problem, but it is still prohibitively complex. With the success of deep unfolding in trading off complexity and performance, we propose to apply deep unfolding to the WMMSE algorithm. With respect to traditional end-to-end learning, deep unfolding incorporates expert knowledge, with the benefits of well-grounded architecture selection, fewer trainable parameters, and better explainability. However, the classical formulation of the WMMSE algorithm given by Shi et al. is not amenable for deep unfolding due to matrix inversions, eigendecompositions, and bisection searches. Therefore, we present an alternative formulation that circumvents these operations. By means of simulations, we show that the deep unfolded WMMSE algorithm performs on par with the original WMMSE algorithm, at a lower computational load.

Deep Weighted MMSE Downlink Beamforming

Lissy Pellaco, Mats Bengtsson, Joakim Jaldén
pellaco@kth.se

https://github.com/lpkg/WMMSE-deep-unfolding/tree/ICASSP2021

Problem Formulation

- MU-MISO interference downlink channel
- Single base station with M transmit antennas
- N single-antenna users
- Linear beamforming

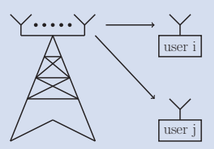

Received signal of user i

$$y_i = \boldsymbol{h}_i^H \boldsymbol{v}_i x_i + \sum_{j=1, j\neq i}^{N} \boldsymbol{h}_i^H \boldsymbol{v}_j x_j + n_i$$

$\boldsymbol{h}_i \in \mathbb{C}^M$: channel, n_i: noise
$\boldsymbol{v}_i \in \mathbb{C}^M$: beamformer vector
x_i: transmitted symbol
$\boldsymbol{V} \triangleq [\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_N]^T \in \mathbb{C}^{N\times M}$

- We address the **NP-hard** problem

$$\max_{\boldsymbol{V}} \quad \sum_{i=1}^{N} \alpha_i \log_2(1+\mathrm{SINR}_i) \tag{1a}$$
$$\text{s.t.} \quad \mathrm{Tr}(\boldsymbol{V}\boldsymbol{V}^H) \leq P \tag{1b}$$

  - $\log_2(1+\mathrm{SINR}_i)$ is the rate of user i
  - α_i is the priority of user i (assumed to be known)
- We want to comply with the **power consumption** and **latency constraints** at the base station

WMMSE algorithm

- It works on an **equivalent reformulation**

$$\min_{\boldsymbol{u},\boldsymbol{w},\boldsymbol{V}} \quad f(\boldsymbol{u},\boldsymbol{w},\boldsymbol{V}) \tag{2a}$$
$$\text{s.t.} \quad \mathrm{Tr}(\boldsymbol{V}\boldsymbol{V}^H) \leq P \tag{2b}$$

- f is jointly nonconvex over $(\boldsymbol{u},\boldsymbol{w},\boldsymbol{V})$
- f is **convex in each optimization variable**

ITERATIVE ALGORITHM PSEUDOCODE

repeat
$\boldsymbol{u} = \mathrm{argmin}_{\boldsymbol{\xi}} f(\boldsymbol{\xi},\boldsymbol{w},\boldsymbol{V})$
$\boldsymbol{w} = \mathrm{argmin}_{\boldsymbol{\xi}} f(\boldsymbol{u},\boldsymbol{\xi},\boldsymbol{V})$
$\boldsymbol{V} = \mathrm{argmin}_{\boldsymbol{\xi}} f(\boldsymbol{u},\boldsymbol{w},\boldsymbol{\xi})$ s.t. $\mathrm{Tr}(\boldsymbol{\xi}\boldsymbol{\xi}^H) \leq P$
until convergence

- Guaranteed to **converge to a local optimum**
- **Relatively high computational complexity**

Deep unfolding

- **Goal**: trade off complexity and performance in presence of computational constraints and latency constraints for iterative algorithms
- **Key idea**: build and train a neural network whose structure is determined by the iterative algorithm
- It **incorporates domain knowledge**

Making the WMMSE unfoldable

- The update equations of $\boldsymbol{u}$ and $\boldsymbol{w}$ can be **easily mapped to neural network layers**
- Conversely, the update of $\boldsymbol{V}$, which is obtained by solving

$$\min_{\boldsymbol{\xi}} \quad f(\boldsymbol{u},\boldsymbol{w},\boldsymbol{\xi}) \tag{3a}$$
$$\text{s.t.} \quad \mathrm{Tr}(\boldsymbol{\xi}\boldsymbol{\xi}^H) \leq P, \tag{3b}$$

  requires **a matrix inversion, an eigendecomposition, and a bisection search** ☹
- We propose to solve (3) with the **projected gradient descent (PGD)** approach ☺
- We **truncate** the sequence of PGD steps to K

UNFOLDABLE WMMSE ALGORITHM PSEUDOCODE

for $l = 1, \ldots, L$
$\boldsymbol{u} = \mathrm{argmin}_{\boldsymbol{\xi}} f(\boldsymbol{\xi},\boldsymbol{w},\boldsymbol{V})$
$\boldsymbol{w} = \mathrm{argmin}_{\boldsymbol{\xi}} f(\boldsymbol{u},\boldsymbol{\xi},\boldsymbol{V})$
for $k = 1, \ldots, K$
$\boldsymbol{V}^{(k)} = \Pi_{\mathcal{C}}\{\boldsymbol{V}^{(k-1)} - \gamma \nabla f(\boldsymbol{V}^{(k-1)})\}$

$$\Pi_{\mathcal{C}}\{\boldsymbol{V}\} = \begin{cases} \boldsymbol{V}, & \text{if } \mathrm{Tr}(\boldsymbol{V}\boldsymbol{V}^H) \leq P \\ \frac{\boldsymbol{V}}{\|\boldsymbol{V}\|}\sqrt{P}, & \text{otherwise.} \end{cases} \tag{4}$$

We prove that the unfoldable WMMSE retains the **same convergence guarantees** of the WMMSE

Deep unfolded WMMSE

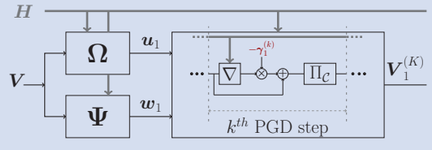

- The step sizes of the PGD ($\boldsymbol{\Gamma}$) are the trainable parameters, where $\boldsymbol{\Gamma} = (\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2, \ldots, \boldsymbol{\gamma}_L)$
- We minimize the **global objective** loss function

$$\mathcal{L}(\boldsymbol{\Gamma}) = -\frac{1}{N_s}\sum_{n=1}^{N_s}\sum_{l=1}^{L} f_{\mathrm{WSR}}(\boldsymbol{H}_n, \boldsymbol{V}_l; \boldsymbol{\Gamma}) \tag{5}$$

where N_s is the size of the training set

Numerical Results

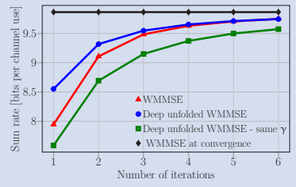

- $\boldsymbol{h}$ is drawn from a Rayleigh distribution
- We initialize $\boldsymbol{V}$ with matched filtering
- We initialize $\gamma_l^{(k)} = 1$ for $k = 1, \ldots, K$ and $l = 1, \ldots, L$
- We set $\alpha_i = 1$ for $i = 1, \ldots, N$
- For each combination of M, L, K, and SNR we train a different network
- We adopt Adam optimizer

Persson, Patrik
Lund University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Parameterization of Ambiguity in Monocular Depth Prediction

Monocular depth estimation is a highly challenging problem that is often addressed with deep neural networks.
While these use recognition of high level image features to predict reasonably looking depth maps the result often has poor metric accuracy.
Moreover, the standard feed forward architecture does not allow modification of the prediction based on cues other than the image.
In this paper we relax the monocular depth estimation task by proposing a network that allows us to complement image features with a set of auxiliary variables. These allow disambiguation when image features are not enough to accurately pinpoint the exact depth map and can be thought of as a low dimensional parameterization of the surfaces that are reasonable monocular predictions.
By searching the parameterization we can combine monocular estimation with traditional photo-consistency or geometry based methods to achieve both visually appealing and metrically accurate surface estimations.
Since we relax the problem we are able to work with smaller networks than current architectures.
In addition we design a self supervised training scheme, eliminating the need for ground truth image-depth-map pairs.
Our experimental evaluation shows that our method generates more accurate depth maps and generalizes better than competing state-of-the-art approaches.

Persson, Patrik
Lund University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Parameterization of Ambiguity in Monocular Depth Prediction

Patrik Perssion
Lund University

Introduction

Dense depth or disparity estimation is a classical problem in computer vision . Traditional methods use stereo (or multi-camera) setups and attempt to match every pixel in the reference image to a corresponding pixel in a neighbouring image using appearance cues. While the accuracy of the recovered depth is often very high for correctly matched pixels, ambiguous texture can degrade the matching and often leads to a noisy depth map. To stabilize the result a popular approach is to add geometric regularization terms such as derivative or curvature penalties. These can be realized as low order potentials in a conditional random field and efficient inference can be performed with move-making or message passing algorithms. While this kind of prior can drastically improve the estimation in ambiguous image regions, they lack any ability to recognize complex geometries and are basically limited to encouraging piece-wise planar or smooth surfaces.

A more recent approach is to use neural networks to directly infer depth or dense matching. An extreme case of this is monocular depth estimation where a neural network is used to estimate depth from a single image [3]. These networks typically require a huge amount of training data and may generalize poorly. In addition, while they achieve meaningful results with plausible object shapes, the resulting depth maps are often inaccurate because of the ambiguous nature of the problem. To resolve these ambiguities, CodeSlam [1] and DeepFactors [2] introduce an image dependent low-dimensional latent scene representation by training a conditional variational auto-encoder. Given an image, a depth basis is predicted and is linearly combined with the latent representation to form the depth. Here both methods are trained supervised using ground truth depth maps, limiting the datasets that can be used for training.

Our approach

Our goal is to use a neural network to extract a low dimensional shape parameterization from a single image that is flexible enough to allow depth fitting using traditional stereo cues such as photo consistency. To achieve this we add an additional input, the latent variable z, to a U-net architecture designed to complement the model with the information that is not directly observed in the image. Our latent variable model is not trained to recreate a depth map through an auto-encoder, as in CodeSlam or DeepFactors, but rather to complement image features with information needed to predict the depth map. A similar approach was introduced in DeepSDF [4] to learn signed distance functions.

Fig. 1: U-Net complemented with the latent vector z

During training and inference, we treat the latent vector z as a parameter to find during optimization instead of predicting it using an encoder. This allows us to train the network and performed inference in essentially the same way.

Training and Inference

During training and inference we use photometric and geometric consistency losses between overlapping views to constrain the depth instead of using ground truth depth maps. This makes the training self-supervised. For each view, we associate a latent vector z and mean depth α, and optimize these jointly for all images in the co-visible set giving us a depth in each view simultaneously. The losses are described below

$$X_i(p) = \pi^{-1}(D_i(p)) \tag{1}$$
$$q_{ij}(p) = \pi\left(T_{ij}X_i(p)\right) \tag{2}$$
$$\mathcal{L}_{\text{photo}_i} = \lambda_p \sum_{j \neq i} \|M_{ij} \odot (I_j \circ q_{ij} - I_i)\|_\delta \tag{3}$$
$$\mathcal{L}_{\text{depth}_i} = \lambda_d \sum_{j \neq i} \|\frac{1}{\alpha_j} M_{ij} \odot (D_j \circ q_{ij} - D_{ij})\|_\delta, \tag{4}$$

where π, π^{-1} are projection and un-projection operations, T_{ij} the relative transformation, M_{ij} a mask that removes invalid projections and $\|\cdot\|_\delta$ the huber norm. The only difference between training and inference is that during inference the network weights are not updated.

Additionally we explicitly handle occlusion by checking the difference between the depth D_j in view j with projected depth D_{ij} originating form D_i, using an adaptive threshold

$$\Delta_{ij}(p) \leq \textbf{Median}(\Delta_{ij}) - \tau \cdot \textbf{MAD}(\Delta_{ij}), \tag{5}$$

where $\Delta_{ij} = D_j - D_{ij}$. A point is classified as occluded if the condition is true.

Results

The Table below shows comparative results between DeepFactors, MegaDepth and Our method, where it can be seen that our method yield significantly better results.

		lower is better			higher is better		
scene	Method	Abs Rel	Sq Rel	RMSE	$\delta < 1.1$	$\delta < 1.25$	$\delta < 1.25^2$
scene0565_00	DeepFactors	0.1517	0.0693	0.3638	46.08%	76.68%	95.67%
	MegaDepth	0.2749	0.2879	0.6672	36.46%	62.34%	83.20%
	Ours	**0.0980**	**0.0492**	**0.3036**	**67.31%**	**88.70%**	**97.83%**
scene0606_02	DeepFactors	0.1736	0.1546	0.5799	44.96%	73.61%	91.01%
	MegaDepth	0.2312	0.1998	0.5959	35.70%	63.30%	87.39%
	Ours	**0.1232**	**0.0804**	**0.3989**	**61.62%**	**83.73%**	**95.32%**
scene0707_00	DeepFactors	0.1669	0.0913	0.3771	43.02%	73.62%	94.70%
	MegaDepth	0.2452	0.2226	0.5443	33.36%	61.90%	86.73%
	Ours	**0.0883**	**0.0365**	**0.2495**	**71.59%**	**91.31%**	**97.94%**
scene0715_00	DeepFactors	0.0959	0.0653	0.4599	64.16%	90.03%	97.74%
	MegaDepth	0.2291	0.4771	0.9298	45.54%	73.12%	89.03%
	Ours	**0.0672**	**0.0471**	**0.3724**	**80.20%**	**94.26%**	**98.20%**
scene0743_00	DeepFactors	0.1537	0.0549	0.23	45.61%	78.26%	95.70%
	MegaDepth	0.2111	0.1236	0.3570	39.86%	70.00%	89.74%
	Ours	**0.0779**	**0.0020**	**0.1709**	**74.52%**	**92.8%**	**98.70%**

One key difference between our method and CodeSlam and DeepFactors is that in our method, the depth is non-linear in z and each component of z has a local impact on the depth map in contrast to the aforementioned. We hypothesise that this is the reason for the significant performance increase since it may allow larger variations in depth and a slightly lower regularizing effect which may allow it to better adapt to new scenes while still regularizing local structure.

Examples

Conclusion

In this work, we have presented a learning approach for monocular depth estimation that takes ambiguities into account by providing a low dimensional parameterization of a family of feasible depth maps. We have shown that optimizing over this representation using photo-consistency losses yields accurate and realistic geometries. Our experimental results indicate, both qualitative and quantitative, that our approach performs better or on par with competing state-of-the-art methods.

Acknowledgements

This work was partially supported by the strategic research projects ELLIIT and eSSENCE, the Swedish Foundation for Strategic Research project, Semantic Mapping and Visual Navigation for Smart Robots (grant no. RIT15-0038) and Wallenberg Artificial Intelligence, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation.

References

[1] M. Bloesch et al. "CodeSLAM - Learning a Compact, Optimisable Representation for Dense Visual SLAM". In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018, pp. 2560–2568. DOI: 10.1109/CVPR.2018.00271.

[2] J Czarnowski et al. "DeepFactors: Real-time probabilistic dense monocular SLAM". In: *IEEE Robotics and Automation Letters* 5 (2020), pp. 721–728. DOI: 10.1109/lra.2020.2965415. URL: http://dx.doi.org/10.1109/lra.2020.2965415.

[3] Zhengqi Li and Noah Snavely. "MegaDepth: Learning Single-View Depth Prediction from Internet Photos". In: *Computer Vision and Pattern Recognition (CVPR)*. 2018.

[4] Jeong Park et al. "DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019, pp. 165–174. DOI: 10.1109/CVPR.2019.00025.

Do kernel and neural embeddings improve optimization and generalization?

We extend the recent results of (Arora et al. 2019) by spectral analysis of the representations corresponding to the kernel and neural embeddings. They showed that in a simple single-layer network, the alignment of the labels to the eigenvectors of the corresponding Gram matrix determines both the convergence of the optimization during training as well as the generalization properties. We generalize their result to the kernel and neural representations and show these extensions improve both optimization and generalization of the basic setup studied in (Arora et al. 2019). In particular, we first extend the setup with the Gaussian kernel and the approximations by random Fourier features as well as with the embeddings produced by two-layer networks trained on different tasks. We then study the use of more sophisticated kernels and embeddings, those designed optimally for deep neural networks and those developed for the classification task of interest given the data and the training labels, independent of any specific classification model.

Do kernel and neural embeddings improve optimization and generalization?

Arman Rahbar, Emilio Jorge, Devdatt Dubhashi, Morteza Haghir Chehreghani
{armanr, emilio.jorge, dubhashi, morteza.chehreghani}@chalmers.se

CHALMERS UNIVERSITY OF TECHNOLOGY
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Description

The authors of [1] has showed that in a simple single-layer network, the alignment of the labels to the eigenvectors of the corresponding Gram matrix determines both the convergence of the optimization during training as well as the generalization properties. We thought about generalizing their results to other representations of the data. Specifically, we worked with representations induced by different Kernels as well as Neural Networks.
In particular, we extend the setup in [1] with Kernels and approximations of Kernel embeddings as well as with the embeddings produced by two-layer neural networks.

Background & Motivation

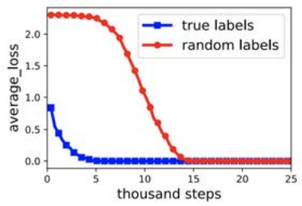

Figure 1: Rethinking Generalization Experiment. Source: ICML slides for [1]
Rethinking Generalization Experiment [3]: Gradient Descent for a neural network reaches nearly 0 training loss for both correct and random labels. **BUT** we see better generalization and faster convergence for correct labels. So we have two fundamental questions to answer:
Optimization. Why do true labels give faster convergence rate than random labels for gradient descent?
Generalization. What property of properly labeled data controls generalization?
In [1], the authors consider a specific simple two-layer network model:

$$f_{W,a}(x) = \frac{1}{\sqrt{m}} \sum_{r=1}^{m} a_r \max(0, w_r^T x_i), \quad (1)$$

with $x \in R^d$, $w_1, ..., w_m \in R^{d \times m}$ and $(a_1, ..., a_m)^T \in R^m$ (where m specifies the number of neurons in the hidden layer, i.e., its width). This network is trained on dataset of data points $\{x_i\}$ and their targets $\{y_i\}$.
They show that both training and generalization are better if the label vector y aligns with the eigenvectors corresponding to the top eigenvalues of $\mathbf{H}^\infty$ (Gram Matrix) where

$$\mathbf{H}^\infty_{i,j} := E_{\mathbf{W} \sim \mathcal{N}(0,\mathcal{I})} \left[\mathbf{x}_i^T \mathbf{x}_j 1[\mathbf{w}^T \mathbf{x}_i \geq 0, \mathbf{w}^T \mathbf{x}_j \geq 0] \right] \quad (2)$$

and

$$\mathbf{H}^\infty = \sum_i \lambda_i \mathbf{v}_i \mathbf{v}_i^T \quad (3)$$

is the orthonormal decomposition of $\mathbf{H}^\infty$.

Our Method

The simple two-layer network can be extended with adding different types of embeddings ϕ at the input layer corresponding to a kernel $\mathcal{K}$:

$$f_{\mathbf{W},\mathbf{a}}(\mathbf{x}) = \frac{1}{\sqrt{m}} \sum_{r=1}^{m} a_r \max(0, \mathbf{w}_r^T \phi(\mathbf{x}_i)). \quad (4)$$

Then the Gram Matrix can be defined in the same way as before. Let its eigenvalues be ordered as $\lambda_0(\mathcal{K}) \geq \lambda_1(\mathcal{K}) \geq \cdots \geq \lambda_{n-1}(\mathcal{K})$ and let $\mathbf{v}_0(\mathcal{K}), \cdots, \mathbf{v}_{n-1}(\mathcal{K})$ be the corresponding eigenvectors. We can also view the representations generated by successive neural network layers as a type of embedding which helps us to understand what happens in deeper networks.
A kernel $\mathcal{K}$ such that the corresponding eigenvectors align well with the labels would be expected to perform well both for training optimization as well as generalization.
For our kernelized network the optimization and generalization are respectively controlled by:

$$\sqrt{\sum_i (1 - \eta \lambda_i(\mathcal{K}))^{2k} (\mathbf{v}(\mathcal{K})_i^T \mathbf{y})^2} \quad (5)$$

and

$$\mathbf{y}^T (\mathbf{H}(\mathcal{K})^\infty)^{-1} \mathbf{y} \quad (6)$$

Our Kernels

We have used the following kernels in our experiments:
1. Gaussian Kernel: The Gaussian kernel is given by $\mathcal{K}(x_i, x_j) := \exp\left(-\gamma \|x_i - x_j\|^2\right)$.
2. Neural Kernel: By adding another layer to the network and training it to convergence we can then use the weights of the first layer as a kernel embedding for training a new network.
2. Arc-cosine Kernel: This kernel mimics the computations in a neural network within an infinite dimensional feature space.
3. Optimized Kernel: We use the method proposed in [2] that suggests an algorithm to learn a new kernel from a group of kernels based on a similarity measure between the kernels, namely centered alignment. The learned kernel is expected to align well with the labels.

How To Find Kernel Embeddings?

We can use existing methods for approximation of kernel embeddings. Specifically, we used *random Fourier features* (RFF) for Gaussian Kernel, and *Nyström method* for other Kernels.

Some of the Results

Below we can see the experimental train and test errors at the different steps of Gradient Descent as well as eigenvector projections for the CIFAR-10 dataset for various Kernels:

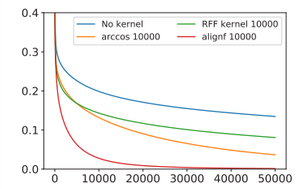

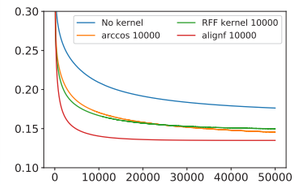

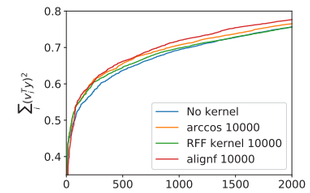

Bibliography

References

[1] Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332, 2019.

[2] Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Algorithms for learning kernels based on centered alignment. *J. Mach. Learn. Res.*, 13:795–828, March 2012.

[3] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. *ArXiv*, abs/1611.03530, 2016.

Saloni Kwatra, Saloni
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A Privacy Preserving Federated Learning Framework with Decision Trees

The aim is to build models for decentralized data following the Federated Learning (FL) approach. In FL, a general approach for learning consists of an iterative process in which (i) a set of agents is selected and they download the current model; (ii) the agents compute an updated model based on their data; (iii) the model updates are sent to the server; and (iv) the server aggregates these models to construct an improved general model. Privacy issues arise concerning the updated model sent to the server and how the models are aggregated. We plan to work on non-numerical (e.g., decision trees) and unsupervised learning models. This study presents an approach for implementing FL with decision trees on horizontally partitioned data.

Saloni Kwatra, Saloni
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A Privacy Preserving Federated Learning Framework with Decision Trees

Saloni Kwatra, PhD at Umeå University
Department of Computing Science
Main Supervisor: Prof. Vicenç Torra
Co-Supervisor: Associate Prof. Lili Jiang

UMEÅ UNIVERSITET

Motivation & Research Goals

The aim is to build models for decentralized data following the Federated Learning (FL) approach. In FL, a general approach for learning consists of an iterative process in which (i) a set of agents is selected and they download the current model; (ii) the agents compute an updated model based on their data; (iii) the model updates are sent to the server; and (iv) the server aggregates these models to construct an improved general model. Privacy issues arise concerning the updated model sent to the server and how the models are aggregated. We plan to work on non-numerical (e.g., decision trees) and unsupervised learning models. This study presents an approach for implementing FL with decision trees on a horizontally partitioned data.

Methods

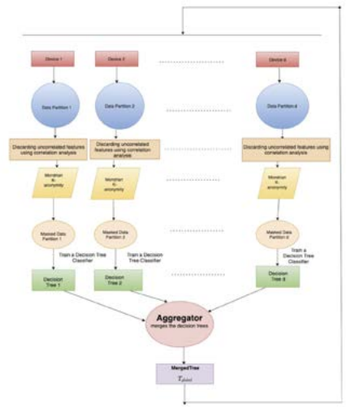

Creating Non-Independent and Identically Distributed Partitions

$$\begin{aligned} \text{Minimize} \quad & p^T Q p + p^T L \\ \text{subject to} \quad & \\ & \textstyle\sum_{i=1}^{d} p_{ij} = d n_j / n \ \text{ for each } \ j = 1, \ldots, l \\ & \textstyle\sum_{j=1}^{l} p_{ij} = 1 \ \text{ for each } \ i = 1, \ldots, d \\ & p_{ij} \geq 0 \ \text{ for each } \ i = 1, \ldots, d \text{ and } j = 1, \ldots, l \\ & p_{ij} = 0 \ \text{ for each } \ p_{ij} \in \mathcal{N} \end{aligned} \tag{1}$$

where, n is number of records, l is number of classes, p_{ij} denotes probability for i^{th} device and j^{th} class, $\sum_{j=1}^{l} p_{ij} = 1$ for all devices $i = 1, \ldots, d$, N denotes the set of probabilities that needs to be set to zero, square matrix $Q = Id$ (i.e., the identity matrix of size $d \cdot l$) and the vector $L = -2\mathrm{A}$ where $\mathrm{A} = (\alpha_{11}, \alpha_{12}, \ldots, \alpha_{dl})$. We consider a quadratic objective function of the form $OF(p; \alpha_{ij}) = (p_{ij} - \alpha_{ij})^2 = (p_{ij}^2 - 2\alpha_{ij} p_{ij} + \alpha_{ij}^2)$, where α_{ij} is a random number taken from a uniform distribution in [0,1].

References

[1] C. Fan, P. Li, Classification acceleration via merging decision trees, in: Pro- ceedings of the 2020 ACM-IMS on Foundations of Data Science Conference 2020, pp. 13–22

[2] LeFevre, K., DeWitt, D. J., & Ramakrishnan, R.(2006, April). Mondrian multidimensional k-anonymity. In 22nd International conference on data engineering (ICDE'06) (pp. 25-25). IEEE.

Selected Results

Firstly, each device protects its raw data using Mondrian, and then trains a decision tree classifier on its protected data. Devices share the root node of their trees with the aggregator. The aggregator merges the trees by choosing the most common split attribute and grows the branches based on the split values of the chosen split attribute. This recursive process stops when all the nodes to be merged are leaf nodes. After the merging operation, the aggregator sends the merged decision tree to the distributed devices. Hence, we build a joint machine learning model based on the data from multiple devices while offering k-anonymity to the participants.

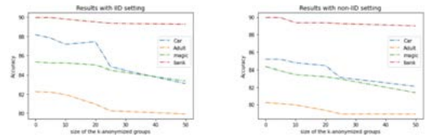

We show the results when the devices use data from the UCI datasets and where the databases follow non-IID data partitions of the whole dataset. We observe a drop of 2% in accuracy in the non-IID setting compared with the IID setting, when k=0 (k=0 means without anonymisation). This shows that the case of non-IID partitioning of data has some effect on the performance of our framework. From our perspective, this behavior is due to the fact that different devices have data with different probability distributions. The classification accuracy remains acceptable in the non-IID case, even when the value of k is as big as 50.
From the privacy point of view, using Mondrian k- anonymity for converting raw dataset into anonymised dataset is a good choice, as it is robust to various kinds of attacks and also it takes into account the multivariate distribution of the data.

Possible Future Improvements of our work

- In our implementation, all devices have the policy of sharing all their nodes with the aggregator. It would be possible that different devices have different policies about sharing (or not) their nodes.
- In our approach, node sharing is based on the decision tree learned with the original data. If partial decision trees are shared by the aggregator, devices can recompute their trees at each iteration.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Sanchez Aimar, Emanuel
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Deep Expert Learning for Long-Tailed Recognition

Many real-world recognition problems present a highly imbalanced or long-tailed label distribution. This setting makes representation learning more challenging and tends to bias classifiers towards head classes, resulting in limited generalization for tail classes. By simultaneously addressing these issues, ensemble learning has shown promising results for long-tailed classification, exhibiting a good trade-off between head and tail performance. The aim of this research is to further study methods for learning and integration of multiple specialized models, called experts, to solve long-tailed visual problems. In addition, we investigate the calibration properties of the proposed expert-based model under long-tailed data regimes.

AI MLX | Page 81 B

Sanchez Aimar, Emanuel
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Deep Expert Learning for Long-Tailed Recognition

Emanuel Sanchez Aimar, Linköping University
Computer Vision Laboratory, ISY
Supervisors: Michael Felsberg, Marco Kuhlmann

LINKÖPING UNIVERSITY

Motivation & Research goals

Many real-world recognition problems present a highly imbalanced or long-tailed label distribution. This setting makes representation learning more challenging and tends to bias classifiers towards head classes, resulting in limited generalization for tail classes [1]. By simultaneously addressing these issues, ensemble learning has shown promising results for long-tailed classification, exhibiting a good trade-off between head and tail performance [2, 3]. The aim of this research is to further study methods for learning and integration of multiple specialized models, called *experts*, to solve long-tailed visual problems. In addition, we investigate the calibration properties of the proposed expert-based model under long-tailed data regimes.

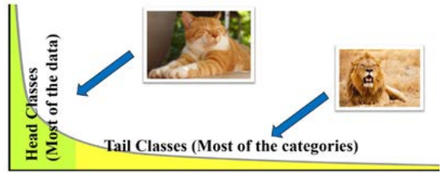

Illustration courtesy of [4].

Classical approaches & drawbacks

Data re-sampling

Loss re-weighting

→ Under-fitting to the head
Over-fitting to the tail
Sub-optimal for feature learning

(Class-dependent) margins

of ensemble learning

How to achieve diversity & specialization?
E.g., by data division

How to combine experts effectively?
E.g., by averaging

Framework

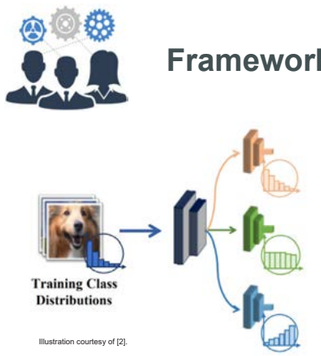

Illustration courtesy of [2].

- Stage 1: **train experts independently**, by simulating different class distributions
- Stage 2: **combine experts in an ensemble** (logit space)

1. *How to improve stage 1?* We propose a new parameterization for the expert losses inspired on [3]
2. *A more flexible fusion strategy for stage 2?* We investigate a learnable gating model
3. Interesting property to look into: *model calibration*

References

1. Kang, Bingyi, et al. "Decoupling Representation and Classifier for Long-Tailed Recognition." *International Conference on Learning Representations* (2019).
2. Zhang, Yifan, et al. "Deep long-tailed learning: A survey." *arXiv preprint arXiv:2110.04596* (2021).
3. Zhang, Yifan, et al. "Test-agnostic long-tailed recognition by test-time aggregating diverse experts with self-supervision." *arXiv preprint arXiv:2107.09249* (2021).
4. Tang, Kaihua, et al. "Long-Tailed Classification by Keeping the Good and Removing the Bad Momentum Causal Effect." *Advances in Neural Information Processing Systems* 33 (2020).

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Machine Learning over Networks: Optimization point of view

One reason for the spectacular success of machine learning models can be the appearance of large datasets. Large datasets cannot be processed on a single machine due to resource limitations or slow training. Moreover, some datasets are private, and data sharing is prohibited. Distributed optimization methods can help use much more sensitive data points to train a machine learning model quickly because the dataset is stored on different nodes, and each node does the computations associated with a collection of data points available to it.

Machine Learning over Networks: From an optimization point of view

Firooz Shahriari-mehr, PhD student at Chalmers University of Technology
Deptartment of Computer Science and Engineering
Supervisors: Ashkan Panahi

CHALMERS UNIVERSITY OF TECHNOLOGY

Motivation

One reason for the spectacular success of machine learning models can be the appearance of large datasets. Large datasets cannot be processed on a single machine due to resource limitations or slow training. Moreover, some datasets are private, and data sharing is prohibited. Distributed optimization methods can help us to use much more sensitive data points to train a machine learning model quickly because the dataset is stored on different nodes, i.e., computational machines, each of them does the computations associated with a collection of data points available to it. In this poster, I will present a novel distributed optimization algorithm for the convex finite-sum minimization problem with explicit convex constraints over strongly-connected directed graphs.

Methods

Finite-sum minimization problem with explicit constraints:

$$\min_{\mathbf{x}\in\mathbb{R}^m} \frac{1}{M}\sum_{v=1}^{M} f_v(\mathbf{x}) \qquad \text{subject to} \quad \mathbf{x}\in\bigcap_{v=1}^{M} S_v$$

The setup that we have considered to solve this problem is a set of M nodes, each of them has access to its own local objective function and constraint set. The nodes communicate over a decentralized network represented by a directed graph.

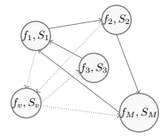

Goal of the network: All nodes converge to a consensus solution which satisfies the following optimality condition:

$$\mathbf{0}\in\sum_{v=1}^{M}\left(\partial I_{S_v}(\mathbf{x}^*)+\nabla f_v(\mathbf{x}^*)\right)$$

Previous proposed algorithms[1] have assumed undirected graph, diminishing step-size, or identical constraints assumption for convergence analysis.

Double Averaging and Gradient Projection[2]**:** Each node v does the following updates in each iteration:

$$\begin{aligned}
\mathbf{z}^v &= \mathbf{x}_k^v - \sum_{u\in\mathcal{N}_v^{in}\cup\{v\}} w_{vu}\mathbf{x}_k^u - \mu\left(\nabla f_v(\mathbf{x}_k^v)-\mathbf{g}_k^v\right)\\
\mathbf{x}_{k+1}^v &= P_{S_v}\left(\mathbf{z}^v\right)\\
\mathbf{g}_{k+1}^v &= \mathbf{g}_k^v + \rho\left[\nabla f_v(\mathbf{x}_k^v) - \mathbf{g}_k^v + \frac{1}{\mu}\left(\mathbf{z}^v-\mathbf{x}_{k+1}^v\right)\right] + \alpha\left(\mathbf{h}_k^v-\mathbf{g}_k^v\right)\\
\mathbf{h}_{k+1}^v &= \mathbf{h}_k^v - \sum_{u\in\mathcal{N}_v^{in}\cup\{v\}} q_{vu}\left(\mathbf{h}_k^u-\mathbf{g}_k^u\right)
\end{aligned}$$

General concepts used in DAGP:

- Weighted averaging: each node receives some information from its neighbors and calculates a weighted average using gossip matrices.
- Projection: each node does a projection onto the local constraint set.
- Constrained gradient tracking: similar to the variance reduction techniques in stochastic optimization, the vectors called g^v store the past values of averaged (sub)gradients and feasible directions.
- Distributed null projection: to reach optimality, $\sum g^v = 0$ (null condition), which is achievable in a distributed way. To this end, we introduce the "distributed projection" of g^vs, called h^v, onto the space $\sum h^v = 0$.

Selected Results

Theorem 1 (Consensus and Optimal Solution)**.** *If the iterates of DAGP algorithm converge, any stopping point is an optimal and consensus solution, i.e.,* $\mathbf{x}^v = \mathbf{x}^*,\ \forall v\in\mathcal{V}$, *and* $\mathbf{x}^*$ *satisfies the sufficient optimality conditions.*

Theorem 2 (Convergence Rates). DAGP recovers the standard rates for convex and smooth objectives under some common assumptions without strong-convexity assumption. We showed feasibility and optimality gaps decay respectively with the rates of $1/K$ and $1/\sqrt{K}$, where K is the total number of iterations.

DAGP has superior empirical results in comparison to the DDPS[1] algorithm in the following constrained problem:

$$\begin{aligned}
f_v(\mathbf{x}) &= \log\left(\cosh(\mathbf{a}_v^T\mathbf{x}-b_v)\right), && v=1,\dots,20\\
S_v &= \{\mathbf{x}\in\mathbb{R}^{10}\mid \mathbf{c}_v^T\mathbf{x}-d_v\le 0\}. && v=1,\dots,20
\end{aligned}$$

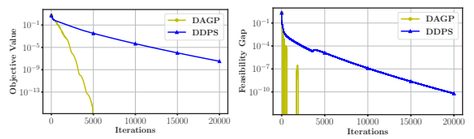

Moreover, **DAGP** shows comparable performance with the state-of-the-art push-pull[3] algorithm in the logistic regression problem for two digits of the MNIST dataset, which is an unconstrained problem.

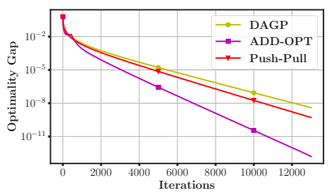

References

[1] Distributed Subgradient Projection Algorithm Over Directed Graphs, C. Xi and U. A. Khan, IEEE Transactions on Automatic Control, vol. 62, Aug. 2017.

[2] Decentralized Constrained Optimization: Double Averaging and Gradient Projection, F. Shahriari-Mehr, D. Bosch, A. Panahi, Conference on Decision and Control, Dec. 2021.

[3] Push-pull gradient methods for distributed optimization in networks, S. Pu, W. Shi, J. Xu, and A. Nedic, IEEE Transactions on Automatic Control, January, 2021.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Sidwall Thygesen, Signe
Linköping university

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Exploring Transition Ensembles using Hierarchical Clustering and Visual Representations

Ensemble analysis is a challenging problem, appearing in many scientific applications. This work introduces a pipeline to explore ensembles by combining automatic and interactive visual analysis, focusing on molecular electronic transition ensembles within the chemistry domain. An electronic transition describes the change in charge distribution between two molecular states. The goal of the pipeline is to characterize and compare the transitions, and how they correlate to physical properties. Each ensemble member is described with a quantitative feature vector, making it possible to utilize hierarchical clustering. A visual summarization for each cluster as well as the whole ensemble is proposed, building on the feature vector representation. Other interactive visual components are used, supporting both exploration of clusters and outlier detection, as well as investigation of correlation. The usefulness of the pipeline is shown by applying it to data from theoretical chemistry.

Sidwall Thygesen, Signe
Linköping university

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Exploring Transition Ensembles using Hierarchical Clustering and Visual Representations

LINKÖPING UNIVERSITY
Signe Sidwall Thygesen, Linköping University
Scientific Visualization, Media and Information Technology

Background & Motivation

Ensemble analysis is a challenging problem, appearing in many scientific applications. This work introduces a pipeline to explore ensembles by combining automatic and interactive visual analysis. It is applied it to molecular transition ensembles from theoretical chemistry, addressing the need to better explore such data. An electronic transition describes the change in charge distribution between two molecular states. The chemists are interested in finding out how much charge is transferred between different subgroups of the molecule. The purpose of the pipeline to characterize and compare transitions, and how they correlate to physical properties.

Pipeline

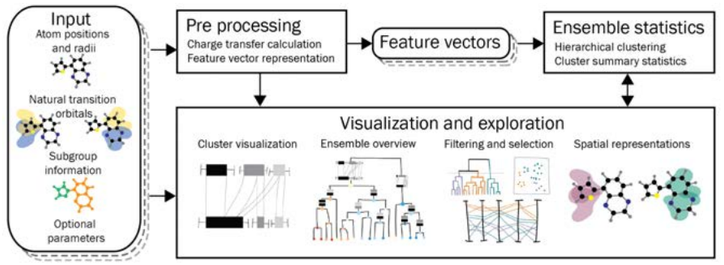

Each ensemble member is quantitatively described with a feature vector containing the charge in each subgroup of the molecule (building on a charge transfer matrix from previous work [1]). This quantitative description gives the possibility to utilize hierarchical clustering and statistical summarization methods. A measure of locality value (ML) is derived, capturing how much charge is transferred in the transition. The result is used in the visual representations. An automated dendrogram, augmented with cluster representations on different level of details gives an overview of the ensemble. Other interactive visual components are used, supporting both exploration of clusters and outlier detection, as well as investigation of correlation. Spatial representations of selected transitions links our abstract visualizations back to a view more familiar to the chemists.

Contributions

- A novel **feature vector** to describe electronic transitions and a quantitative **measure of locality** for distinguishing between transitions of different nature
- A **visual pipeline** for ensemble analysis **combining automatic and explorative** methods
- A **level of detail representation** summarizing and conveying the mean behavior of a cluster
- Introduction of **augmented dendrograms** to provide a hierarchical visual representation of ensemble data

Example

To illustrate the pipeline, it is applied to an electronic transition ensemble for a molecule consisting of two subgroups (G1 and G2).

The **automatically generated augmented dendrogram** gives an overview of the clusters….

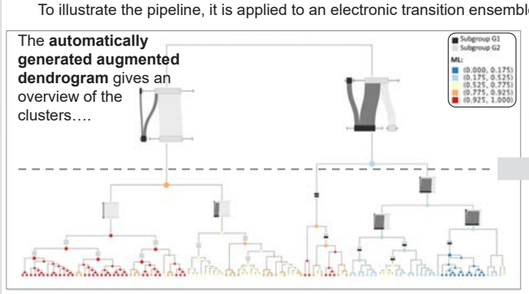

…which can be used to select a good level of detail in the **interactive exploration.**

A **parallel coordinates plot** shows correlation between parameters, and makes it possible to filter and inspect ensemble members.

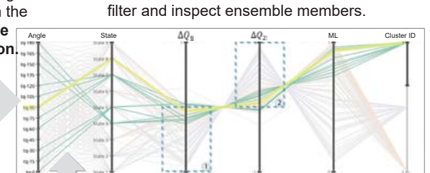

A linked **scatter plot** provides an additional overview and the possibility to select outliers.

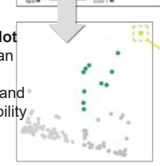

The **spatial representations** for a selected transition helps the chemists to understand the data.

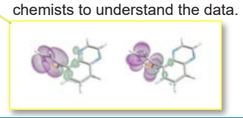

References

1. Bin Masood et al, *Visual Analysis of Electronic Densities and Transitions in Molecules,* Computer Graphics Forum, 2021.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Predicting progression & cognitive decline inamyloid-positive patients with Alzheimer's disease

In Alzheimer's disease, amyloid-β (Aβ) peptides aggregate in the brain forming CSF amyloid levels - a key pathological hallmark of the disease. However, CSF amyloid levels may also be present in cognitively unimpaired elderly individuals. We aim to explain the variance in disease progression among patients with Aβ-pathology. We perform prediction of a) the change of MMSE score using regression models for 2 and 4 years after follow-up and b) the change in diagnostic using classification model for 2 years after follow-up.
We show in our analysis that CSF levels of Aβ are not strong predictors of the rate of cognitive decline in Aβ-positive subjects when adjusting for other variables.

Predicting progression & cognitive decline in amyloid-positive patients with Alzheimer's disease

Hákon Valur Dansson[1], Lena Stempfle[1]*, Hildur Egilsdóttir[1], Alexander Schliep[1], Erik Portelius[2,3], Kaj Blennow[2,3], Henrik Zetterberg[2,3,4,5], and Fredrik D. Johansson[1]

[1]CHALMERS, [2]GU, [3]SAHLGRENSKA UNIVERSITY HOSPITAL, [4]UCL NEURODEGENERATIVE DISEASES, [5]UCL DEMENTIA RESEARCH

Summary

In Alzheimer's disease, amyloid-β (Aβ) peptides aggregate in the brain forming CSF amyloid levels - a key pathological hallmark of the disease. However, CSF amyloid levels may also be present in cognitively unimpaired elderly individuals. We aim to explain the variance in disease progression among patients with Aβ-pathology. We perform prediction of a) the change of MMSE score using regression models for 2 and 4 year after follow-up and b) the change in diagnostic using classification model for 2 year after follow-up. We show in our analysis that CSF levels of Aβ are not strong predictors of the rate of cognitive decline in Aβ-positive subjects when adjusting for other variables.

Determination of amyloid-positive status

Even in individuals with Aβ pathology, there is substantial variation in symptoms, such as cognitive function, for this reason our work focuses on predicting progression in individuals with elevated Aβ CSF levels. Aβ pathology status was determined based on the Aβ42/Aβ40 ratio.

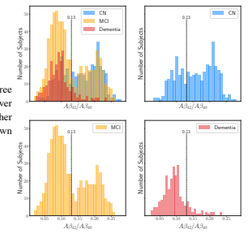

Figure 1: The full cohort was split into three groups: those who had a baseline Aβ ratio lower than 0.13 (Aβ-positive), those who had a higher ratio (Aβ-negative), and those with unknown status.

Potential predictors

Predictive models were built on *two* different sets of features. The first set of features *(all features)* was preselected following [48] and expanded to include key features from the ADNI TadPole competition [49] in addition to a few features that were available for over 90% of the ADNI cohort. This resulted in a set of 37 features including biomarkers tau, ptau and Aβ42 in CSF, and 15 different cognitive tests among others. The second feature set (cognitive tests only) consists only of the 15 cognitive tests.

Derivation and evaluation cohorts

Our experiments are performed on three different cohorts. To compensate for the small number of 1) only Aβ-positive subjects (*Aβ only*), were compared to training cohorts including 2) (*All Subjects*) combined Aβ-positive and Aβ-negative subjects and those without Aβ measurements into one derivation set and 3) (*All Subjects*, *Weighted*), with weighted samples with respect to the All Subjects cohort to mimic a larger sample of Aβ-positive subjects.
We let the latent state $C \in \{0,1\}$ of a Gaussian mixture model (GMM), fit to the Aβ-ratios of the All Subjects cohort, represent Aβ-positivity. The weight w_i was computed as:

$$w_i = \hat{p}(R = r_i | C = 1) / \hat{p}(\mathrm{R} = r_i)$$

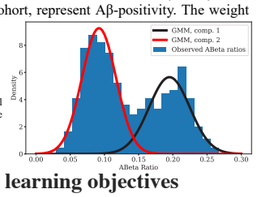

Figure 2: Each subject i was assigned a weight $w_i > 0$ based on the probability that their individual Aβ ratio r_i would be observed for an average hypothetical Aβ-positive subject

Prediction models and learning objectives

We studied the progression of Aβ-positive subjects with respect to two principal outcomes:

Task A: Predicting change in MMSE score for 2 (A1) and 4 (A2) years of follow-up

Change in MMSE score assessed as cognitive function [32], relative to baseline and 2 years. MMSE takes values on a scale from 0 to 30 where a lower score represents worse cognitive function [45].
We considered the performance of linear regression and gradient boosting models that predict the change in MMSE scores measured using the average cross-validated R^2 score and standard deviation.

Task B: Predicting change in diagnosis for 2 years (B1) follow-up

Changes in dementia diagnoses (CN/MCI/AD) were determined by comparing the disease status, indicating whether or not a subject's diagnosis had worsened in 2 year expressed by a binary variable.
For the classification task, a logistic regression model and the same tree-based gradient-boosting approach as for Task A was used and evaluated by cross-validated weighted F1 score.

Results

We perform a set of experiments using ADNI data to predict the change in MMSE after 2 and 4 year of follow-up and the change in diagnosis after 2 years.
The best predictive model of change in cognitive test scores for Aβ-positive subjects at the 2-year follow-up achieved an R^2 score of 0.388 while the best model predicting adverse changes in diagnosis achieved a weighted F1 score of 0.791. When predicting cognitive score change 4 years after baseline, the best model achieved an R^2 score of 0.325 and it was found that fitting models to the extended cohort improved performance. Moreover, using all clinical variables outperformed the best model based only on a suite of cognitive test scores which achieved an R^2 score of 0.228.

Table 1: Performance of the linear and gradient boosting regressions, predicting change in MMSE two and four years after baseline for three different cohort selections. We compare models trained on features a) the all features set from baseline and b) from baseline cognitive scores only.

R^2 (SD)	2-year follow-up	4-year follow-up
All features		
Linear regression, *Aβ Only*	0.372 (0.081)	0.205 (0.227)
Linear regression, *All Subjects*	0.354 (0.083)	0.325 (0.134)
Linear regression, *All Subjects, Weighted*	0.388 (0.073)	0.304 (0.152)
Gradient boosting, *Aβ Only*	0.287 (0.124)	0.156 (0.244)
Gradient boosting, *All Subjects*	0.356 (0.108)	0.252 (0.191)
Gradient boosting, *All Subjects, Weighted*	0.338 (0.950)	0.263 (0.192)
Cognitive tests only		
Linear regression, *Aβ Only*	0.343 (0.087)	0.178 (0.203)
Linear regression, *All Subjects*	0.333 (0.081)	0.228 (0.143)
Linear regression, *All Subjects, Weighted*	0.350 (0.079)	0.225 (0.160)
Gradient boosting, *Aβ Only*	0.272 (0.133)	-0.050 (0.358)
Gradient boosting, *All Subjects*	0.323 (0.118)	0.149 (0.224)
Gradient boosting, *All Subjects, Weighted*	0.293 (0.114)	0.118 (0.227)

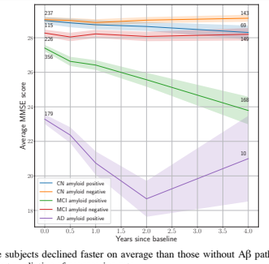

Figure 3: Aβ-positive subjects declined faster on average than those without Aβ pathology, but the specific level of CSF Aβ was not predictive of progression rate.

Conclusion

Our results illustrate high correlation between important predictors which offers future investigation to eventually handle the high missingness in the data. Baseline assessments of cognitive function accounts for the majority of variance explained in the prediction of 2-year decline but is insufficient for achieving optimal results in longer-term predictions.

Acknowledgements

Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu).

References

[32] Galea, M., Woodward, M.: Mini-mental state examination (mmse).Australian Journal of Physiotherapy51(3), 198 (2005)
[45] Dick, J., Guiloff, R., Stewart, A., Blackstock, J., Bielawska, C., Paul,E., Marsden, C.: Mini-mental state examination in neurologicalpatients. Journal of Neurology, Neurosurgery & Psychiatry47(5),496–499 (1984)
[48] Nguyen, M., He, T., An, L., Alexander, D.C., Feng, J., Yeo, B.T.T.:Predicting alzheimer's disease progression using deep recurrent neuralnetworks. NeuroImage222, 117203 (2020).doi:10.1016/j.neuroimage.2020.11720349.
[49] Marinescu, R.V., Oxtoby, N.P., Young, A.L., Bron, E.E., Toga, A.W.,Weiner, M.W., Barkhof, F., Fox, N.C., Eshaghi, A., Toni, T., et al.:The alzheimer's disease prediction of longitudinal evolution (tadpole)challenge: Results after 1 year follow-up. arXiv preprintarXiv:2002.03419 (2020)

Sundqvist, Tobias
Umeå University / TietoEVRY

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Anomaly detection and root-cause analytics in 5G Radio Access Network

Our research objective is to improve the observability and speed up the fault finding process in the 5G Radio Access Network (RAN). We develop and utilize many different machine learning methods to analyze RAN system logs. From the logs we can learn and distinguish between the normal and abnormal behavior and aid the developers in locating the problems.

AI MLX | Page 85 B

Sundqvist, Tobias
Umeå University / TietoEVRY

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Anomaly detection and root-cause analytics in 5G Radio Access Network

We use AI to find faults quicker and more easily

Tobias Sundqvist, Monowar Bhuyan, Erik Elmroth, and Johan Forsman

5G RAN

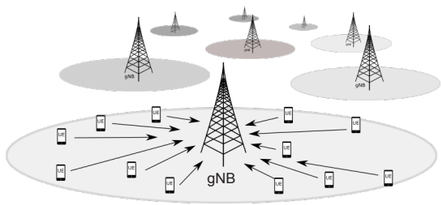

PROBLEM

It will be very challenging to make the 5G Radio Access Network (RAN) reliable due to its complexity:

- Increased network dynamics with different services
- Distributed system of radio units at antenna sites and applications in edge data center cites
- Separation of hardware and software from many vendors with individual configurations
- Virtualized RAN applications share resources with others on multi-purpose hardware
- Analysis of huge metrics from multiple sources
- Difficult to know if the system behaves normally

System logs

AdaBoost ensemble of LSTM

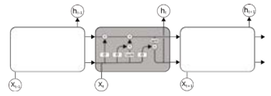

Detect changes in procedures

LSTM models learn the sequential order of procedures in the system log. Anomalies are detected as the order of events changes or abnormal events occur.

Statistical and clustering

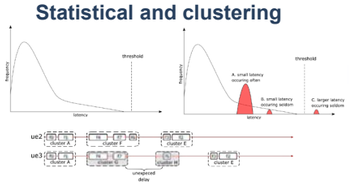

Detect abnormal delays

The methods provide a detailed view of where abnormal delays and large variations occur.

PROOF OF CONCEPT

Continuous Integration

Methods are detecting anomalies in customers data. Speeds up fault localization process.

Train / test

Root-cause detection

To fully understand the root cause of all software bugs in RAN we need to monitor both the function and kernel calls.

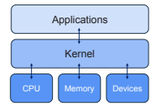

CallGraph

Learns what kind of function and kernel calls that occur between debug printouts in system log.

MultiSpace

Creates a call and mean time vector for each function.

- Find deviation in function call pattern.
- Find largest deviation in time for each function.

Scan me

And learn more about my research
http://www.qvistigt.com/

UMEÅ UNIVERSITET
Department of Computing Science, SE-901 87 Umeå, Sweden
Email: {sundqtob, monowar, elmroth}@cs.umu.se

tieto EVRY

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Tabakovic, Selma
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Co-clustering of Tensor Data Using Sparse Tensor Factorisation

With the ever-increasing amounts of data generated from new sources and scientific methods, e.g. high throughput genome sequencing methods in bioinformatics, powerful tools for exploratory data analysis are required. One such tool is clustering, i.e. grouping together coherent observations in data, which is important for categorising vast amounts of observations into a more manageable format for further analysis. However, this task is subject to new challenges as tensor data, i.e. multidimensional data, has become a frequent occurrence in many applications. For tensor data, a clustering approach called co-clustering has recently attracted research attention. Co-clustering means that the clustering is performed on all the tensor dimensions simultaneously, which enables the detection of joint data expressions that only occur under special circumstances. Here a method for co-clustering of tensor data using a sparse CP decomposition is proposed.

Tabakovic, Selma
Chalmers

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Co-clustering of Tensor Data Using Sparse Tensor Factorisation

Selma Tabakovic,Chalmers University of Technology
Department of Applied Mathematics and Statistics
Supervisors: Rebecka Jörnsten (Chalmers and University of Gothenburg)

CHALMERS UNIVERSITY OF TECHNOLOGY

Motivation & Research Goals

With the ever-increasing amounts of data generated from new sources and scientific methods, e.g. high throughput genome sequencing methods in bioinformatics, powerful tools for exploratory data analysis are required. One such tool is clustering, i.e. grouping together coherent observations in data, which is important for categorising vast amounts of observations into a more manageable format for further analysis. However, this task is subject to new challenges as tensor data, i.e. multidimensional data, has become a frequent occurrence in many applications. For tensor data, a clustering approach called co-clustering has recently attracted research attention. Co-clustering means that the clustering is performed on all the tensor dimensions simultaneously, which enables the detection of joint data expressions that only occur under special circumstances. Here a method for co-clustering of tensor data using a sparse CP decomposition is proposed.

Methods

The *CP decomposition* splits a tensor into a finite sum of rank one tensors, i.e. tensors that can be written as an outer product of N vectors Finding the CP decomposition of a three-way tensor $\mathcal{X} \in \mathbb{R}^{I\times J\times K}$ may be formalised as

$$\min_{\hat{\mathcal{X}}} \|\mathcal{X} - \hat{\mathcal{X}}\| \quad \text{where} \quad \hat{\mathcal{X}} = \sum_{r=1}^{R} \lambda_r \, \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r \,,$$

where R is a positive integer, λ_r is a scaling factor, and $\mathbf{a}_r \in \mathbb{R}^I$, $\mathbf{b}_r \in \mathbb{R}^J$ and $\mathbf{c}_r \in \mathbb{R}^K$ for $r = 1, 2 \dots R$, and are all normalised to length one

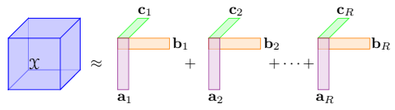

The sparse CP decomposition from [1] regularises the factors by a l_1 norm penalty and calculates them in a sequential manner.

The *mean square residue* (MSR) is used to measure the coherence of a co-cluster $\mathcal{D}$, and is defined as [2]

$$\mathrm{MSR}(\mathcal{D}) = \frac{1}{PQR} \sum_{p\in\mathcal{P}, q\in\mathcal{Q}, r\in\mathcal{R}} (d_{pqr} - d_{p\mathcal{Q}\mathcal{R}} - d_{\mathcal{P}q\mathcal{R}} - d_{\mathcal{P}\mathcal{Q}r} + 2d_{\mathcal{P}\mathcal{Q}\mathcal{R}})^2,$$

where $d_{p\mathcal{Q}\mathcal{R}}$, $d_{\mathcal{P}q\mathcal{R}}$ and $d_{\mathcal{P}\mathcal{Q}r}$ denotes the mean of the pth row, qth column and kth tube, respectively, and $d_{\mathcal{P}\mathcal{Q}\mathcal{R}}$ denotes the mean of $\mathcal{D}$.

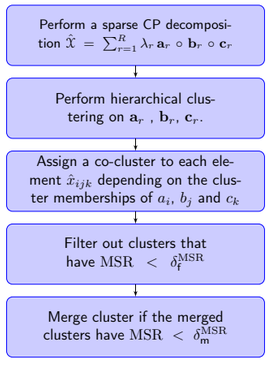

Selected Results

Different kinds of co-clusters were simulated. **Left**: Eight constant co-clusters were simulated. **Right**: The co-clusters found, corresponding to the true co-clusters.

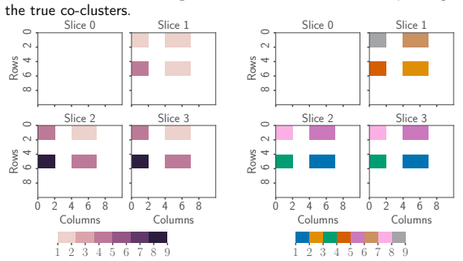

The method was applied on a three-way genomic data set, containing dESeq2 normalised estimates of the fold changes, for different genes, drugs and cell lines.

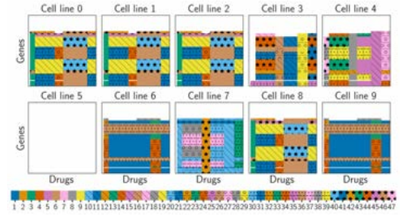

The method has the potential to detect several types of additive coherent co-clusters. Applying it to real genomic data revealed several interesting co-clusters, which provides the biological researches with interesting gene-drug co-clusters to investigate further. Thus the method could be a useful tool for detecting coherent co-clusters in tensor data, and for exploratory data analysis.

References

[1] Sparse higher-order principal components analysis
Allen, Genevera
Artificial Intelligence and Statistics, 2012

[2] δ-TRIMAX: Extracting Triclusters and Analysing Coregulation in Time Series Gene Expression Data
Bhar, Anirban and Haubrock, Martin and Mukhopadhyay, Anirban and Maulik, Ujjwal and Bandyopadhyay, Sanghamitra and Wingender, Edgar
International Workshop on Algorithms in Bioinformatics, 2012

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Developing privacy-aware ML based on ML model spaces

In order to build privacy-aware machine and statistical learning we want to study the relationship between machine learning models and the databases that generate these models. That is, to study the relationship between the space of data and the space of models (hypothesis space). For example, when you have a database DB1 you can build a decision tree (say, DT1). If you have another database DB2 you build another decision tree (say, DT2). It is clear that when you have different databases, you may generate different models. However, we have proven that there are models that appear very often. E.g., it may be that different databases produce DT2. So, some ML models are kind of ”popular” models. We call these models recurrent. We think that these ”recurrent” models are better from a privacy perspective. We plan to study the relationship between the space of data (possible databases, sets of databases) and the space of models (e.g., possible decision trees as in the example above).

AI MLX
Taha, Mariam
Umeå University

Tarle, Magnus
KTH / Hitachi Energy

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning to Control Multiple Power Electronic Converters

The PhD project aims to identify a data-driven control architecture to meet the growing challenges of the electrical power system. In particular, the focus is on optimizing the control of multiple power electronic converters. The main objectives are to achieve greater utilization of the power system capacity, higher power system stability and improved power quality.

AI MLX | Page 88 B

Tarle, Magnus
KTH / Hitachi Energy

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning to Control Multiple Power Electronic Converters

Magnus Tarle (KTH/Hitachi Energy)
Supervisors: Mårten Björkman (KTH), Mats Larsson (Hitachi Energy)
Co-supervisors: Lars Nordström (KTH), Gunnar Ingeström (Hitachi Energy)

Stiftelsen för Strategisk Forskning

Hitachi Energy

Abstract

The PhD project aims to identify a data-driven control architecture to meet the growing challenges of the electrical power system. In particular, the focus is on optimizing the control of multiple power electronic converters. The main objectives are to achieve greater utilization of the power system capacity, higher power system stability and improved power quality.

Motivation & Background

- The electrical power system is evolving and more coordinated, faster and accurate control alternatives are foreseen to be needed

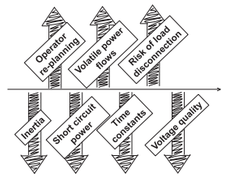

- Using heuristics and engineering, a coordinated control of power electronic converters can increase transfer capability [2]
- Reinforcement learning can potentially self-learn such a control scheme

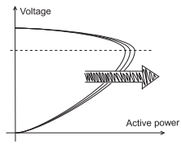

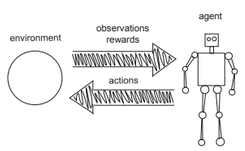

- Control challenges include:
  - Non-convex and sequential-decision-making problem
  - Model and measurement uncertainty
  - Partial observability
  - Variations in topology and dynamics
  - Safety and reliability crucial

Methods & Preliminary Results

- We introduced a world model based reinforcement learning architecture (WMAP) in [1]. Contributions include:
  - Adapting the world model architecture from [3] to the power system domain
  - Increasing safety by adding a shield to make the architecture ask for guidance whenever uncertainty above threshold
  - Performing a case study on the IEEE 14-bus system and benchmarking the performance on a test scenario.

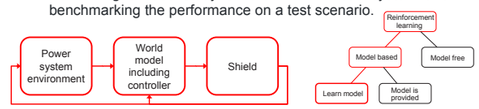

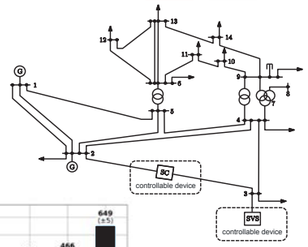

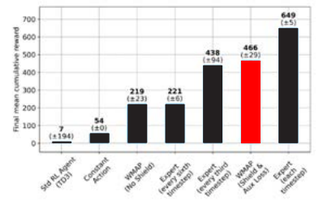

Results showed improved performance compared to model free reinforcement learning algorithm, fixed control setpoints and non-frequent expert setpoint updates.

More results and benchmark variants in the paper.

References

[1] Magnus Tarle, Mårten Björkman, Mats Larsson, Lars Nordström, Gunnar Ingeström, "A World Model Based Reinforcement Learning Architecture for Autonomous Power System Control", *SmartGridComm, 2021. In press.*

[2] M. Larsson, C. Rehtanz, and D. Westermann, "Improvement of Crossborder Trading Capabilities through Wide-area Control of FACTS", *In Bulk Power Syst. Dyn. Control VI, 2004.*

[3] D. Ha and J. Schmidhuber, "Recurrent world models facilitate policy evolution", *Adv. Neural Inf. Process. Syst., vol. 2018-Decem, no. C, pp. 2450–2462, 2018.*

Future Work

Considerations for future work include:

- Continual learning
- Curriculum learning
- Scalability
- Improved safety
- Planning
- Benchmarking against state-of-the-art
- Hardware-in-the-loop implementation and assessment

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

DoubleMatch: improving Semi-Supervised Learning with a Self-Supervised Loss

Following the success of supervised learning, semi-supervised learning (SSL) is now becoming increasingly popular. SSL is the family of methods which in addition to a labeled training set, also use a large set of unlabeled data for training the model. Most of the recent successful SSL methods are based on pseudo-labeling approaches: letting confident model predictions act as training labels. While these methods have shown impressive results on many benchmark datasets, a drawback of this approach is that not all unlabeled data are used during training. We propose a method: DoubleMatch, which combines the pseudo-labeling approach with a self-supervised loss, enabling the model to utilize all unlabeled data through all stages of the training process. We show that this method achieves state of the art accuracies on multiple benchmark datasets while also reducing training times compared to previous SSL methods.

DoubleMatch: improving semi-supervised learning with self-supervision

Erik Wallin, Ind. PhD, Saab AB and Chalmers
Supervisors: Lars Hammarstrand, Fredrik Kahl and Lennart Svensson

CHALMERS UNIVERSITY OF TECHNOLOGY
SAAB

Motivation & Research Goals

Supervised learning has gained a lot of attention in recent years because of remarkable achievements in fields such as image classification, object detection and natural language processing. The great results within supervised learning are typically fueled by huge amounts of labeled data. In practical applications, however, labeled data might be scarce, expensive, or require expert domain knowledge to attain. In contrast, *unlabeled data*, is often much easier to acquire, through e.g. web scraping or unsupervised sensor recordings. **Semi-supervised learning**, using both labeled and unlabeled data for fitting a model, has recently shown impressive results with methods such as FixMatch [1] and UDA [2]. These methods however suffer from not leveraging all unlabeled data during training. We propose a method, DoubleMatch, that takes inspiration from work in *self-supervised learning* to better utilize all unlabeled data. With this method we hope to 1) reduce converge rates of previous methods 2) increase test accuracy on benchmark datasets.

Methods

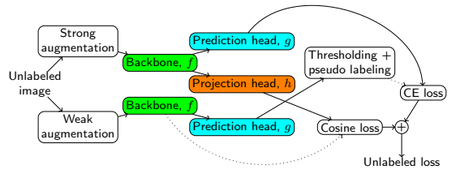

We build upon the FixMatch framework for semi-supervised learning. FixMatch uses the traditional cross-entropy loss for labeled data. For unlabeled data it utilizes both weak, α, and strong, β, data augmentations. A confident prediction on the weakly augmented sample is used as a *pseudo label* for a strong augmentation of that same sample:

$$l_p = \frac{1}{\mu B}\sum_{i=1}^{\mu B} \mathbb{1}\{\max(p(y|\alpha(x_i)) > \tau\} H(\text{argmax}\,(p(y|\alpha(x_i))), p(y|\beta(x_i)))$$

where τ is the confidence threshold, μB is the unlabeled batch size, and x_i are the unlabeled samples.
As we can see, only data with confident predictions on weakly augmented data are used to calculate this loss. We propose adding a feature loss on the output from the penultimate layer of the classification network using the cosine similarity:

$$l_s = -\frac{1}{\mu B}\sum_{i=1}^{\mu B} \frac{h(v_i) z_i}{|h(v_i)||z_i|} = -\frac{1}{\mu B}\sum_{i=1}^{\mu B} \cos(h(v_i), z_i).$$

Here, z_i is the feature vector for weakly augmented sample i and v_i is the feature vector for strongly augmented sample i. h is a trainable linear transformation to allow for different feature representations between weakly and strongly augmented data.

Selected Results

We evaluate our model on benchmark datasets for image classification using different sizes for the labeled training set. We achieve SOTA error rates despite running our method for fewer training iterations than the methods we use as comparisons.

Method	CIFAR10 40 labels	CIFAR10 250 labels	CIFAR10 4000 labels
Π-model	-	54.26±3.97	14.01±0.38
Pseudo-Labeling	-	49.78±0.43	16.09±0.28
Mean Teacher	-	32.32±2.30	9.19±0.19
MixMatch	47.54±11.50	11.05±0.86	6.42±0.10
UDA	29.05±5.93	8.82±1.08	4.88±0.18
ReMixMatch	19.10±9.64	5.44±0.05	4.72±0.13
FixMatch (CTA)	**11.39**±3.35	**5.07**±0.33	**4.31**±0.15
DoubleMatch (ours)	13.59±5.60	5.56±0.42	4.65±0.17

Method	CIFAR100 400 labels	CIFAR100 2500 labels	CIFAR100 10000 labels	STL10 1000 labels
Π-model	-	57.25±0.48	37.88±0.11	26.23±0.82
Pseudo-Labeling	-	57.38±0.46	36.21±0.19	27.99±0.83
Mean Teacher	-	53.91±0.57	35.83±0.24	21.43±2.39
MixMatch	67.61±1.32	39.94±0.37	28.31±0.33	10.41±0.61
UDA	59.28±0.88	33.13±0.22	24.50±0.25	7.66±0.56
ReMixMatch	44.28±2.06	27.43±0.31	23.03±0.56	5.23±0.45
FixMatch (CTA)	49.95±3.01	28.64±0.24	23.18±0.11	5.17±0.63
DoubleMatch (ours)	**41.83**±1.22	**27.07**±0.26	**21.22**±0.17	**4.35**±0.20

To illustrate our increase in training speed the below figure shows test accuracy as a function of training iteration for DoubleMatch and FixMatch during a training run on Cifar100 with 10000 labeled data:

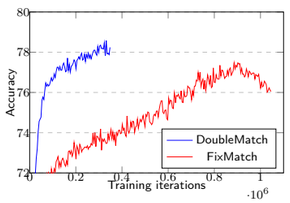

Our method performs well across many datasets. It does however seem to perform worse in the very low-label regime (e.g. Cifar10 with 40 labels). Our hypothesis is that high quality pseudo-labels is more important than self-supervision for the low-label datasets. Another weakness of our method is that the weight for the self-supervised loss needs to be tuned for each dataset.

References

[1] FixMatch: Simplifying Semi-Supervised Learning with Consistency and Confidence
Sohn, K., Berthelot, D., Li, C.L., Zhang, Z., Carlini, N., Cubuk, E.D., Kurakin, A., Zhang, H. and Raffel, C.
NeurIPS 2020

[2] Unsupervised data augmentation for consistency training
Xie, Q., Dai, Z., Hovy, E., Luong, M.T. and Le, Q.V.
NeurIPS 2020

Willemsen, Bram
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Collecting visually-grounded dialogue data with a new vision+language task

We created a new vision+language task to collect visually-grounded dialogue data. The task is framed as a cooperative game in which two players have to come to an agreement on how to rank a set of images given some sorting criterion. The task is designed in such a way that it should lead to naturally-flowing conversations between participants discussing visual information and at the same time enable the study of the grounding and generation of referring expression in the face of distractors.

AI MLX | Page 90 B

Willemsen, Bram
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

There are few visually-grounded dialogue datasets containing symmetric interactions of an unrestricted nature that capture commonly-observed dialogue phenomena and provide the opportunity to study the generation and grounding of referring expressions in the face of distractors, so we created a new vision+language task and decided to collect one

Bram Willemsen / bramw@kth.se
Dmytro Kalpakchi, Gabriel Skantze

"Which cat has the best beard?"

Task: a cooperative image ranking game in which two players have to reach an agreement on how images should be ranked given some sorting criterion

NB: the example dialogue is a highly-simplified artificial example for illustrative purposes only

Robot learning of symbol grounding in multiple contexts through dialog

KTH VETENSKAP OCH KONST

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning Impedance Actions for Safe Reinforcement Learning in Contact-Rich Tasks

1. Extend the action space of RL policy by incorporating variable impedance
2. Our method can be safely deployed on the real robot directly

AI MLX | Page 91 B
Yang, Quantao
Örebro University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning Impedance Actions for Safe Reinforcement Learning in Contact-Rich Tasks

Quantao Yang[1], Alexander Dürr[2], Elin Anna Topp[2], Johannes A. Stork[1], Todor Stoyanov[1]
[1]AMM Lab, Örebro University, Sweden
[2]Dept. of Computer Science, Lund University, Sweden

Contributions

- Extend the action space of RL policy by incorporating ***variable impedance***
- Our method can be safely deployed on the real robot ***directly***

Reinforcement Learning with Skill Priors

- The evidence lower bound (ELBO) to learn a low-dimensional skill latent space[1]:

$$\log p(a) \geq \mathbb{E}_q[\log p(a \mid z) - \beta(\log q(z \mid a) - \log p(z))]$$

- Skill prior model $p_a(z \mid s_t)$ is used to guide Soft Actor-Critic (SAC) policy

Variable Impedance Action Space

- The RL agent predicts the stiffness for contact-rich tasks
- We extend the policy action as the combination of end-effector pose $x \in SE(3)$ in Cartesian space and variable stiffness matrix $K \in \mathbb{R}^{6\times6}$

$$F_a = K(x - x_d) + D(\dot{x} - \dot{x}_d) + M(\ddot{x} - \ddot{x}_d)$$

Experiment Results

Peg-in-hole contact-rich task:

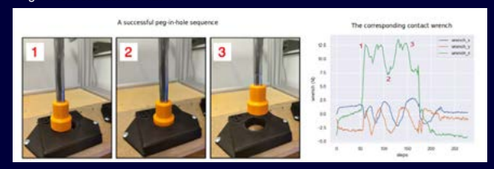

Training results:

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

[1] Karl Pertsch, Youngwoon Lee, and Joseph J. Lim. Accelerating reinforcement learning with learned skill

Zangeneh Kamali, Fereidoon
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Camera Pose Posterior Inference for Visual Localisation

Visual localisation is the problem of estimating the camera pose in an environment from a camera image. Learning-based solutions, such as end-to-end camera pose regression, propose to solve this problem with the help of deep learning. One limitation of such approaches, whether wrapped in a probabilistic formulation or not, is that they assume a uni-modal solution to the pose estimation problem. While this assumption might hold in environments with unique visual features, it falls apart in presence of repetitive structures in the environment, where images from multiple camera poses appear visually similar. In this work we propose to learn inference of the complete pose posterior distribution that is desirable in such scenarios with multi-modal solutions, via variational inference of a simple posterior in a latent space, and learning a map from the latent space to SE(3).

Zangeneh Kamali, Fereidoon
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Camera Pose Posterior Inference for Visual Localisation

Fereidoon Zangeneh, Ind. PhD, Univrses & Kungliga Tekniska högskolan
Division of Robotics, Perception and Learning
Supervisors: Patric Jensfelt, Mårten Björkman (KTH),
Alessandro Pieropan, Amit Dekel (Univrses)

KTH VETENSKAP OCH KONST

UNIVRSES

Motivation & Research Goals

Visual localisation is the problem of estimating the camera pose in an environment from a camera image. Learning-based solutions, such as end-to-end camera pose regression, propose to solve this problem with the help of deep learning. One limitation of such approaches, whether wrapped in a probabilistic formulation or not, is that they assume a uni-modal solution to the pose estimation problem. While this assumption might hold in environments with unique visual features, it falls apart in presence of repetitive structures in the environment, where images from multiple camera poses appear visually similar. In this work we propose to learn inference of the complete pose posterior distribution that is desirable in such scenarios with multi-modal solutions, via variational inference of a simple posterior in a latent space, and learning a map from the latent space to SE(3).

Background

Visual localisation solutions [1] traditionally fall into 2 groups:

- Image-based (i.e. image retrieval in a database),
- Structure-based (i.e. feature-matching in a map).

Recent learning-based methods attempt to improve the visual localisation solutions to achieve invariance towards visual conditions such as seasonal and lighting variations. PoseNet [2] is one of the first methods that attempted to cast visual localisation as an end-to-end absolute pose regression problem.

Problem: PoseNet and its variants all assume a uni-modal (e.g. Gaussian) solution to the visual localisation problem.

Method

Objective: Environments with repetitive structures call for a localisation solution that accommodates multi-modal hypotheses.

This requires a pipeline that for a query image produces the posterior distribution over possible camera poses. To avoid restricting the posterior to a fixed parametric form, we opt for a sampling-based solution that allows the learned distribution to take any form.

We formulate this in a pair of deep networks:

1. First network predicts a posterior distribution for the input query image in latent space with a simple parametric form such as Gaussian.
2. Second network maps the latent space to $SE(3)$.

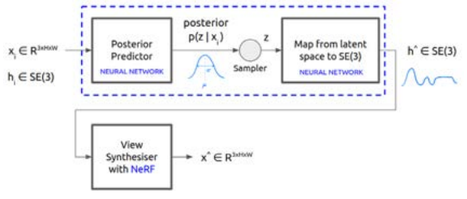

The camera pose posterior distribution for a query image can be simulated by drawing samples from the predicted posterior distribution in the latent space and passing them through the learned map.

Training of the pair of networks is done via 2 supervision signals:

1. KL-divergence between the posterior distribution in the latent space and a prior (e.g. standard Gaussian)
2. photometric error between the query image and the image(s) generated from the camera pose(s) sampled from the posterior distribution

Generative model: A differentiable renderer such as a Neural Radiance Field (NeRF) model [3] pretrained on the scene is used with volume rendering techniques and an appropriate camera model to take $SE(3)$ samples back to image data space.

Challenge: Photometric error as a loss function has a small region of attraction, which requires the samples drawn from the posterior to be close enough the true mode(s) during training.

References

[1] Long-Term Visual Localization Revisited
C. Toft, W. Maddern, A. Torii, L. Hammarstrand, E. Stenborg, D. Safari, M. Okutomi, M. Pollefeys, J. Sivic, T. Pajdla, F. Kahl, T. Sattler
IEEE Transactions on Pattern Analysis and Machine Intelligence, 2020

[2] PoseNet: A Convolutional Network for Real-Time 6-DOF Camera Relocalization
Alex Kendall, Matthew Grimes, Roberto Cipolla
IEEE International Conference on Computer Vision, 2015

[3] NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis
B. Mildenhall, P. Srinivasan, M. Tancik, J. Barron, R. Ramamoorthi, R. Ng
European Conference on Computer Vision, 2020

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning the Pedestrian-Vehicle Interaction for Pedestrian Trajectory Prediction

The prediction of pedestrian behavior can help the drivers and the automated vehicles to make smarter and safer decisions and hence to protect the pedestrians from hazardous situations. The interaction between pedestrians and vehicles is an essential factor that influences pedestrian behavior. In this research, we propose a novel design called the Pedestrian-Vehicle Interaction (PVI) extractor for learning this interaction from data, and implement the proposed PVI extractor on sequential approaches (LSTMs) and non-sequential approaches (CNNs). We use the Waymo Open Dataset consisting of real-world traffic scenes with pedestrians and vehicles. The models using our proposed PVI extractor outperform the state-of-the-art models. The results show that the proposed PVI extractor can capture the interactions between pedestrians and vehicles.

Learning the Pedestrian-Vehicle Interaction for Pedestrian Trajectory Prediction

Chi Zhang, University of Gothenburg | Chalmers
Department of Computer Science and Engineering

* This paper is accepted in 2022 the 8th International Conference on Control Automation and Robotics (ICCAR). 2022.
This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement 860410

Abstract

The prediction of pedestrian behavior can help the drivers and the automated vehicles to make smarter and safer decisions and hence to protect the pedestrians from hazardous situations. The interaction between pedestrians and vehicles is an essential factor that influences pedestrian behavior. In this research, we propose a novel design called the Pedestrian-Vehicle Interaction (PVI) extractor for learning this interaction from data, and implement the proposed PVI extractor on sequential approaches (LSTMs) and non-sequential approaches (CNNs). We use the Waymo Open Dataset consisting of real-world traffic scenes with pedestrians and vehicles. The models using our proposed PVI extractor outperform the state-of-the-art models. The results show that the proposed PVI extractor can capture the interactions between pedestrians and vehicles.

Introduction

Motivation:
Accurately predicting the trajectories of pedestrians in urban traffic scenarios is essential for automated vehicles to prevent hazardous situations. The interaction between pedestrians and vehicles is a key factor that influences the behavior of pedestrians. However, there is limited research on this topic that focuses on pedestrian-vehicle interaction.

Problem Definition:
Use past trajectories of pedestrians and vehicles:

$$X_t^i = (x_t^i, y_t^i), 1 \le t \le T_{obs}, i \in \{1, \dots, N_p\}$$
$$V_t^j = (x_t^j, y_t^j), 1 \le t \le T_{obs}, i \in \{1, \dots, N_v\}$$

to predict future pedestrians' trajectories:

$$\hat{Y}_t^i = (x_t^i, y_t^i), \quad T_{obs} + 1 \le t \le T_{pred}, i \in \{1, \dots, N_p\}$$

where N_p and N_v are the number of pedestrians and vehicles in each frame.

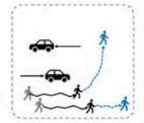

Challenges:
- Multiple factors (features) that influence pedestrian behavior;
- Multiple factors that influence pedestrian-vehicle interaction;
- Difficult to model the interactions between multiple vehicles and multiple pedestrians;
- Use in real-world urban traffic scenarios.

Overall Framework

The features we considered for trajectory prediction:

- Spatial Embedding;
- Social Interaction;
- Pedestrian-Vehicle Interaction.

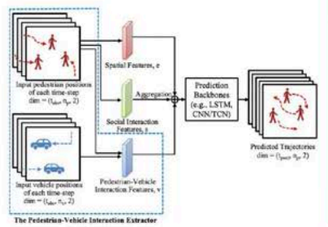

The Pedestrian-Vehicle Interaction Extractor

The PVI Extractor

The Pedestrian-Vehicle Interaction (PVI) extractor:
The factors that we consider in the pedestrian-vehicle interaction:

- Vehicle moving state;
- Relative positions between pedestrians and vehicles;
- Relative velocities between pedestrians and vehicles.

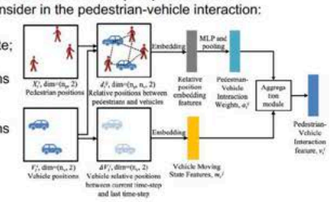

Experiments and Results

Dataset: the urban scenarios in Waymo Open Dataset, consists 450 real-world recording collected in urban traffic scenarios (374 training, 76 test), and each recording spans 20 seconds.

Settings:
- Input: 3.2 sec in the past.
- Output: 4.8 sec in the future.

Evaluations:
- Quantitative evaluation shows the error (in meters) is smaller (better).
- Qualitative evaluation shows the algorithm can handle different scenarios.

TABLE I
THE ADE/FDE METRICS FOR BASELINE METHODS COMPARED TO PROPOSED METHODS ON THE WAYMO OPEN DATASET.

Prediction Method	Model Name	ADE	FDE	Interactions Used in the Model
Sequential	LSTM	0.392	0.844	No interactions.
	Social-LSTM (2016) [4]	0.402	0.840	Social interaction
	Social-GAN (2018) [5]	0.386	0.826	Social interaction
	SI-PVI-LSTM (ours)	**0.372**	**0.796**	Pedestrian-vehicle interaction
	Linear	0.431	0.892	No interactions.
Non-sequential	Social-STGCNN (2020) [6]	0.334	0.530	Social interaction
	Social-IWSTCNN (2021) [16]	0.329	**0.540**	Social interaction
	SI-PVI-Conv (ours)	0.327	0.543	Social interaction; Pedestrian-vehicle interaction

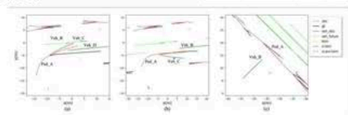

Future Work

- Joint prediction: predict both trajectory and crossing intention of the pedestrian in one framework.
- Hybrid Algorithms: for the pedestrian-vehicle interaction, combine the deep learning-based model with knowledge-based conventional model, and develop a model that is accurate and explainable.

References

1. A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, F.-F. Li, andS. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR), 2016, pp. 961–971.
2. A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 2255–2264.
3. A. Mohamed, K. Qian, M. Elhoseiny, and C. Claudel, "Social-stgcnn: A social spatio-temporal graph convolutional neural network for human trajectory prediction," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 14 424–14 432
4. C. Zhang, C. Berger, and M. Dozza, "Social-iwstcnn: A social interaction-weighted spatio-temporal convolutional neural network for pedestrian trajectory prediction in urban traffic scenarios," in 2021 IEEE Intelligent Vehicles Symposium (IV). IEEE, 2021, pp. 1515–1522.
5. P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caineet al., "Scalability in perception for autonomous driving: Waymo open dataset," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition,2020, pp. 2446–2454.

Other Information

Contact details: chi.zhang@gu.se, zchi@chalmers.se
Supervisor: Associate Professor Christian Berger, GU
Co-Supervisor: Professor Marco Dozza, Chalmers
Our previous work: Social-IWSTCNN, published and presented on IEEE-IV 21 conference. The accessible version (or scan the QR code): https://arxiv.org/pdf/2105.12436.pdf

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Åkerblom, Niklas
Chalmers / Volvo Car Corporation

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Online Learning for Energy Efficient Navigation in Stochastic Transport Networks

Energy-efficient navigation constitutes an important challenge for the electrification of personal transportation, due to the limited battery capacity of electric vehicles. In this project, we employ a Bayesian approach to model the energy consumption at road segments for efficient navigation. In order to learn model parameters, we develop an online learning framework and investigate several exploration strategies, such as Thompson Sampling and Upper Confidence Bound. We also extend the framework to a multi-agent setting, where multiple vehicles adaptively navigate and learn the parameters of the energy model. To establish performance guarantees, we analyze combinatorial Thompson Sampling and derive upper bounds on the expected regret incurred in single-agent and multi-agent settings, through analysis of the algorithm under batched feedback. We demonstrate the performance of our methods via simulation experiments on several real-world city road networks.

Åkerblom, Niklas
Chalmers / Volvo Car Corporation

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Online Learning for Energy Efficient Navigation in Stochastic Transport Networks

CHALMERS UNIVERSITY OF TECHNOLOGY · VOLVO · Niklas Åkerblom · WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM · FFI

Abstract

Energy-efficient navigation constitutes an important challenge for the electrification of personal transportation, due to the limited battery capacity of electric vehicles. In this project, we employ a Bayesian approach to model the energy consumption at road segments for efficient navigation. In order to learn model parameters, we develop an online learning framework and investigate several exploration strategies, such as Thompson Sampling and Upper Confidence Bound. We also extend the framework to a multi-agent setting, where multiple vehicles adaptively navigate and learn the parameters of the energy model. To establish performance guarantees, we analyze combinatorial Thompson Sampling and derive upper bounds on the expected regret incurred in single-agent and multi-agent settings, through analysis of the algorithm under batched feedback. We demonstrate the performance of our methods via simulation experiments on several real-world city road networks.

Background

Widespread adoption of electric vehicle technology is limited by range anxiety. Among other things, public perception is affected by factors like:

- Energy storage capacity limited due to expensive batteries.
- Fast-charging infrastructure still sparse.
- Still comparatively time consuming to charge batteries.

One way to mitigate worries is through better navigation methods. This requires knowledge about the environment!

The problem

- How to utilize a limited set of resources (e.g. autonomous vehicle fleet) to efficiently learn about energy consumption from environment.
- **Energy-efficient navigation problem:** Finding path which minimizes energy consumption through road network graph.
- We consider the energy consumption at different road segments to be stochastic and a priori unknown.
- We want to learn (explore) the parameters of the energy model adaptively while simultaneously solving the navigation problem.

Energy consumption model

- A road network graph may contain millions of edges, hence it is beneficial to utilize prior knowledge for more efficient exploration.
- **Bayesian approach:** Assume prior over expected energy consumption θ_e (e.g. Gaussian conjugate prior) for each edge e:

$$p(\theta_e \mid \mu_{0,e}, \sigma^2_{0,e}) = \mathcal{N}(\theta_e \mid \mu_{0,e}, \sigma^2_{0,e})$$

- Prior uses simplified deterministic energy consumption model based on longitudinal vehicle dynamics.
- Includes knowledge about the vehicle (e.g. mass m, front surface area A) and the road segment (e.g. length l_e, average speed v_e, slope angle α_e):

$$\mu_{0,e} = \frac{mgl_e \sin(\alpha_e) + mgC_r l_e \cos(\alpha_e) + 0.5 C_d A \rho l_e v_e^2}{3600\eta}$$

- Incrementally update posterior with new observations.

Shortest path problem as multi-armed bandit

We consider the stochastic **combinatorial semi-bandit** setting, a multi-armed bandit (MAB) problem where an agent sequentially selects sets of actions (subject to combinatorial constraints) instead of individual actions. The objective is to maximize the long-term expected reward.

- Road network graph $\mathcal{G}(\mathcal{V}, \mathcal{E}, \boldsymbol{\theta}^*)$, where each edge e has expected edge cost $\theta^*_e = -\mathbb{E}[r(e)]$ (we view energy consumption / cost as negative reward $-r(e)$).
- We want to find an optimal action set (edge sequence or path) $\boldsymbol{a}^*$ in the set of all paths $\mathcal{P}$ between two fixed nodes in $\mathcal{V}$ (characterizing the problem instance), such that:

$$\boldsymbol{a}^* = \arg\min_{\boldsymbol{a} \in \mathcal{P}} \sum_{e \in \boldsymbol{a}} \theta^*_e$$

- Bayesian regret (time horizon T, prior π_0 over mean vector $\boldsymbol{\theta}^*$, and expected reward function $f_{\boldsymbol{\theta}}(\boldsymbol{a}) := \sum_{i \in \boldsymbol{a}} \theta_i$):

$$\text{BayesRegret}(T) = \mathbb{E}_{\boldsymbol{\theta}^* \sim \pi_0}\left[\mathbb{E}\left[\sum_{t=1}^{T} \left(f_{\boldsymbol{\theta}^*}(\boldsymbol{a}^*) - f_{\boldsymbol{\theta}^*}(\boldsymbol{a}_t)\right) \;\middle|\; \boldsymbol{\theta}^*\right]\right]$$

Online learning framework

- At time t, use exploration method like Thompson sampling (TS) to compute weights $\boldsymbol{w}_t$.
- Find shortest path $\boldsymbol{a}_t$ w.r.t. weights, e.g. using Dijkstra's algorithm.
- Traverse path $\boldsymbol{a}_t$, observe energy consumption $-r_t$ and update posterior parameters.

Algorithm 1 Online learning for energy-efficient navigation

Require: $\boldsymbol{\mu}_0, \boldsymbol{\sigma}_0$
1: **for** $t \leftarrow 0, 1, \ldots, T$ **do**
2: $\quad \boldsymbol{w}_t \leftarrow$ GetEdgeWeights$(t, \boldsymbol{\mu}_t, \boldsymbol{\sigma}_t)$
3: $\quad \boldsymbol{a}_t \leftarrow$ FindShortestPath$(\boldsymbol{w}_t)$
4: $\quad \boldsymbol{r}_t \leftarrow$ ApplyActionAndObserveReward$(\boldsymbol{a}_t)$
5: $\quad \boldsymbol{\mu}_{t+1}, \boldsymbol{\sigma}_{t+1} \leftarrow$ UpdateParameters$(\boldsymbol{a}_t, \boldsymbol{r}_t, \boldsymbol{\mu}_t, \boldsymbol{\sigma}_t)$

Assign edge weights using bandit algorithms (e.g., TS, BayesUCB):

Algorithm 3 Combinatorial Thompson Sampling

procedure GetEdgeWeights$(t, \boldsymbol{\mu}_t, \boldsymbol{\sigma}_t)$
$\quad$ **for** each edge $e \in \mathcal{E}$ **do**
$\quad\quad \tilde{\theta}_e \leftarrow$ Sample from posterior $\mathcal{N}(\mu_{t,e}, \sigma^2_{t,e})$
$\quad\quad w_{t,e} \leftarrow \mathbb{E}[y_e]$ where $y_e \sim \mathcal{N}^0(\tilde{\theta}_e, \zeta^2_e)$
$\quad$ **return** $\boldsymbol{w}_t$

Algorithm 4 Combinatorial BayesUCB

procedure GetEdgeWeights$(t, \boldsymbol{\mu}_t, \boldsymbol{\sigma}_t)$
$\quad$ **for** each edge $e \in \mathcal{E}$ **do**
$\quad\quad w_{t,e} \leftarrow \max\left(0, Q\left(\frac{1}{t}, \mathcal{N}(\mu_{t,e}, \sigma^2_{t,e})\right)\right)$
$\quad$ **return** $\boldsymbol{w}_t$

Regret analysis

We have derive upper bounds on the Bayesian regret of combinatorial TS for the case where reward feedback is delayed and received in batches of size K:

- Combinatorial TS for batched feedback: $\text{BayesRegret}(T) \leq \tilde{\mathcal{O}}(|\mathcal{E}|K + |\mathcal{E}|\sqrt{T})$.
- Extends to non-delayed case ($K = 1$): $\text{BayesRegret}(T) \leq \tilde{\mathcal{O}}(|\mathcal{E}|\sqrt{T})$.
- Also extends to multi-agent (vehicle fleet of size K) scenario with synchronous sharing of observations, with upper bound on fleet (Bayesian) regret: $\tilde{\mathcal{O}}(|\mathcal{E}|K + |\mathcal{E}|\sqrt{TK})$.

Experiments

- We utilize SUMO traffic simulation data from real-world traffic networks to provide realistic traffic patterns and vehicle speed distributions for simulations.

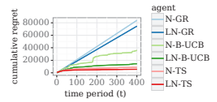

Single-agent algorithm comparison of cumulative regret

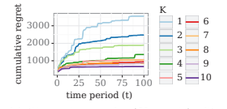

Multi-agent scenarios (K agents) with average cumulative regret per agent in each scenario

- We evaluate exploration through measured cumulative regret and visual inspection.

Exploration with combinatorial TS (explored paths in red)

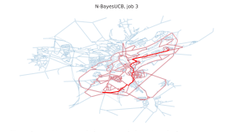

Exploration with combinatorial BayesUCB (explored paths in red

Publications

Åkerblom, N., Chen, Y., and Chehreghani, M. H. (2020).
An online learning framework for energy-efficient navigation of electric vehicles.
In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 2051–2057.

Åkerblom, N., Chen, Y., and Chehreghani, M. H. (2021).
Online learning of energy consumption for navigation of electric vehicles.
arXiv preprint arXiv:2111.02314.

niklas.akerblom@chalmers.se

Poceviciute, Milda
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Unsupervised anomaly detection in digital pathology

Machine learning (ML) algorithms are optimized
for the distribution represented by the training data. For outlier data, they often deliver predictions with equal confidence, even though these should not be trusted. In order to deploy ML-based digital pathology solutions in clinical practice, effective methods for detecting anomalous data are crucial to avoid incorrect decisions in the outlier scenario. We propose a new unsupervised learning approach for anomaly detection in histopathology data based on generative adversarial networks (GANs). Compared to the existing GAN-based methods that have been used in medical imaging, the proposed approach improves significantly on performance for pathology data.
Our results indicate that histopathology imagery is substantially more complex than the data targeted by the previous methods. This complexity requires not only a more advanced GAN architecture but also an appropriate anomaly metric to capture the quality of the reconstructed images.

AI MLX | Page 201 B

Poceviciute, Milda
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Unsupervised anomaly detection in digital pathology

Milda Pocevičiūtė, Gabriel Eilertsen, Claes Lundström
MIT, Linköpings Universitet

LINKÖPINGS UNIVERSITET

Introduction

- Anomaly detection is crucial for safe deployment of digital pathology methods to clinical practice.
- Outlier data is unknown or unavailable at the training time, hence unsupervised detection methods are required.
- GANs learn the distribution of training data, so they are expected to fail to reproduce realistic images of anomalous data
- In our experiments we use healthy patches for training and tumour patches as anomaly data to be detected.

Contributions

•**s2-AnoGAN**: unsupervised StyleGAN2 [1] based anomaly detection method tailored for digital pathology data
•Improved performance compared to two previous methods: f-AnoGAN [2] and pg-AnoGAN [3].

Method

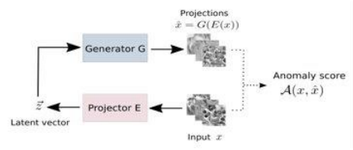

General method for unsupervised anomaly detection with GANs:
•Projector (images -> latent representations)
•Generator (latent representations -> images)
•Anomaly score (comparison of image and its reconstruction)
The generator and gradient descent based projector from StyleGAN2 [1] are used in our method while anomaly score is based on Canny edges [4]: the difference between number of edges in the original images versus in the reconstructed image.
The lower the anomaly score, the better was the reconstruction which implies lower chance the image is an anomaly.
Area under ROC curve (AUC) is used to measure the success of detecting the anomalous images.

References

1 **Karras T et al.** Analyzingand improving the image quality of stylegan. *CVPR 2020*
2 **Schlegl T et al.** f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks. *Medical Image Analysis 2019; 54.*
3 **Berg A et al.** Unsupervised adversarial learning of anomaly detection in the wild. *ECAI 2020*
4 **Canny J**. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence 1986; 6*

Results

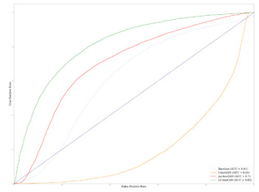

Fig. 1. ROC curves with the corresponding AUC achieved by the tested frameworks: s2-AnoGAN (our), f-AnoGAN, pg-AnoGAN and the baseline (for which only the number of edges is used as anomaly score). Our method had highest performance.

- **s2-AnoGAN** achieved the highest area under ROC curve (AUC) (see Fig. 1)
- As a baseline, we use the count of Canny edges within original images (no GANs' reconstructions involved)
- Poor f-AnoGAN performance is explained by MSE usage in the anomaly metric:
  - Pixel-wise comparison fails as the exact locations of the reconstructed features varies.
- In Fig. 2, we see that our method provides best visual reconstruction for healthy patches. All methods struggle to reconstruct the tumour patches:
  - Implies that more advanced GAN architecture is needed for digital pathology data.
- Table 1 shows AUC scores achieved by different combinations of the GAN + Projector and the anomaly scores

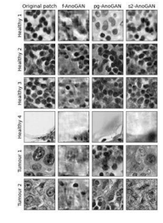

Fig. 2. Examples of test data and their projections by f-AnoGAN, pg-AnoGAN and s2-AnoGAN frameworks.

	f-AnoGAN	pg-AnoGAN	s2-AnoGAN
$\mathcal{A}_{f\text{-}AnoGAN}$	0.21	0.32	**0.37**
$\mathcal{A}_{pg\text{-}AnoGAN}$	**0.74**	0.70	0.71
$\mathcal{A}_{Canny}$	0.75	0.57	**0.82**
$LPIPS$	0.36	0.36	**0.64**
$\mathcal{A}_D$	0.24	0.35	**0.36**
$\mathcal{A}_{MSE}$	0.21	0.32	**0.37**
$\mathcal{A}_{sim}$	0.19	**0.33**	0.27
$\mathcal{A}_{\mathcal{U}}(z)$	**0.75**	0.70	0.71

Table 1. AUC values of anomaly detection when combined with different anomaly metrics.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

WASP WINTER CONFERENCE 2022

Autonomous Systems (AS)

Towards Generalized Robotic Skills and Knowledge Integration

Using skill-based systems to solve robotic tasks in industries has been gaining some popularity. These systems offer easily programmable robotic skills (modular software blocks) that are shareable among robots and rely on integration of planning, learning, sensing and execution. However, one big challenge is the design of a robotic skill (modular software block) that works in general settings. The aim of this research is to develop a framework to solve a variety of industrial tasks using generalized skills with knowledge integration. The framework involves dealing tasks with a lot of parameters and skills, a general way learning new tasks, and knowledge transfer among different hierarchical spaces.

AUTONOMOUS SYSTEMS (AS) | Page 95 B
Ahmad, Faseeh
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Towards Generalized Robotic Skills and Knowledge Integration

Faseeh Ahmad, Lund University
Department of Computer Science, Robotics and Semantic Systems (RSS)
Advisors: Volker Krueger, Elin Anna Topp, Jacek Malek

Motivation & Research Goals

Using skill-based systems to solve robotic tasks in industries has been gaining some popularity. These systems offer easily programmable robotic skills (modular software blocks) that are shareable among robots and rely on integration of planning, learning, sensing and execution. However, one big challenge is the design of a robotic skill (modular software block) that works in general settings. The aim of this research is to develop a framework to solve a variety of industrial tasks using generalized skills with knowledge integration. The framework involves dealing tasks with a lot of parameters and skills, a general way learning new tasks, and knowledge transfer among different hierarchical spaces.

Background

- Learning Parameters in Behavior Trees (BT) for Movement Skills[1]
  - BlackDrops[2]
  - CMAES
  - Behavior Trees
  - Obstacle Avoidance and Peg Insertion
- Learning Tacit Knowledge of Robot Tasks through Planning, Knowledge Integration and Multi-objective Optimization
  - SkiROS[3]
  - Bayesian Optimization (BO)
  - Hypermapper[4]
  - Peg Insertion and Push Task

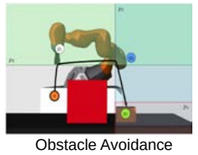

Obstacle Avoidance

Push Task

Focus Points

- Solving tasks with more parameters and skills
- Generalized task specific learning
- Derivative skills
- Knowledge transfer between hierarchical spaces

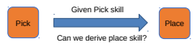

Generalized Task Specific Learning
- Is it robust to changes?
- Can it handle multiple similar scenarios?
- Can it reliably be used to learn new tasks?
- ...

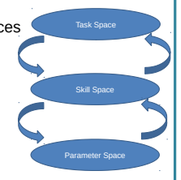

References

1. Mayr, M., Chatzilygeroudis, K., Ahmad, F., Nardi, L. and Krueger, V., 2021. Learning of Parameters in Behavior Trees for Movement Skills. arXiv preprint arXiv:2109.13050.
2. Chatzilygeroudis, K. and Mouret, J.B., 2018, May. Using parameterized black-box priors to scale up model-based policy search for robotics. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 5121-5128). IEEE.
3. Rovida, F., Grossmann, B. and Krüger, V., 2017, September. Extended behavior trees for quick definition of flexible robotic tasks. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 6793-6800). IEEE.
4. Nardi, L., Koeplinger, D. and Olukotun, K., 2019, October. Practical design space exploration. In 2019 IEEE 27th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS) (pp. 347-358). IEEE.
5. https://en.wikipedia.org/wiki/Covariance_matrix

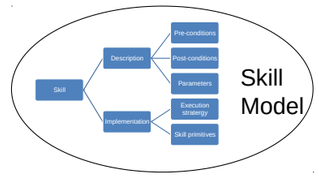

Splitting Tasks

Learning a task with multiple skills and arbitrarily large parameters by maintaining a balance between quality and budget (time)

Case Study: Pick an object and place it at target location

Pick Skill: **Parameters:** $P_1, P_2, \ldots, P_{10}$

Place Skill: **Parameters:** $P_{11}, P_{12}, \ldots, P_{20}$

Solution: Divide the task into sub-tasks and learn the skill independently

Problem Skills can be independent or dependent

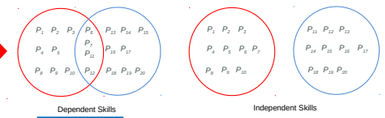

Problem How to determine the dependence of skills?

Use covariance matrix[5] to determine if skills are dependent or independent

$Cov[\mathbf{X},\mathbf{X}] = E[(\mathbf{X} - \mu_x)\ (\mathbf{X} - \mu_x)^T] = E[\mathbf{X}\mathbf{X}^T] - \mu_x\mu_x^T$

where $\mathbf{X}$ = parameter vector and $\mu_x = E[\mathbf{X}]$

Problem How to find the covariance matrix of parameters?

Data Points — Execution of BT with all parameters used

Task may or may not succeed

Truth Table of Execution — Specifies which cases are important

At least one data point per case should exist

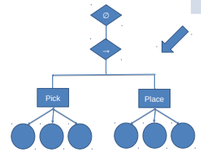

Pick	Place	Reliable data point
Success	Success	Yes
Sucess	Failure	Yes
Failure	Success	No
Failure	Failure	No

P4 Formalization & Verification

Programming Protocol-Independent Packet Processors (P4) is a Domain Specific Language used to program the data plane of networking targets as smart NICs and multi-port switches. The data plane contains two main programmable blocks called parser and control. We analyzed those and built a structural operational semantics (small step) model and executable semantics in order to prove interesting properties related to the processed packets as well as the P4 programs overall. The final outcomes of the project: a P4 interpreter verified down to the binary level, and automation to prove Hoare triple contracts of P4.

AUTONOMOUS SYSTEMS (AS) | Page 96 B
Alshnakat, Anoud
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Anoud Alshnakat, Didrik Lundberg
Roberto Guanciale and Mads Dam

KTH (Theoretical Computer Science Department)

Introduction

Programming Protocol-Independent Packet Processors (P4) is a Domain Specific Language used to program the data plane of networking targets as smart NICs and multi-port switches. The data plane contains two main programmable blocks called parser and control. We analyzed those and built a structural operational semantics (small step) model and executable semantics in order to prove interesting properties related to the processed packets as well as the P4 programs overall. The final outcomes of the project: a P4 interpreter verified down to the binary level, and automation to prove Hoare triple contracts of P4.

Generic P4 Architecture

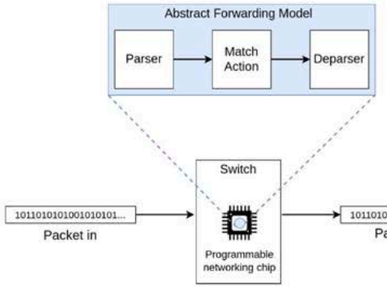

Any packet enters the networking switch supported with a programmable data plane chip (programmed using P4) should be passed in the pipeline to three main stages:

Parser: A finite state machine that maps the bits in the input packet into a type representation. extracts the header fields bits from the packet. It can handle both standard header format as well as custom user-defined header format.

Match Action tables: One or more tables that contain keys and matching kinds, which determine the action to be processed on the packet. This stage requires interacting with the control plain.

Deparser: A control function that assembles the headers back into a well-formed output packet.

References

1. "P4~16~ Language Specification," P4.org, 2017. https://p4.org/p4-spec/docs/P4-16-v1.0.0-spec.html
2. "ott-lang/ott: The Ott tool for writing definitions of programming languages and calculi," GitHub, Jul. 21, 2021. https://github.com/ott-lang/ott

P4 Formalization

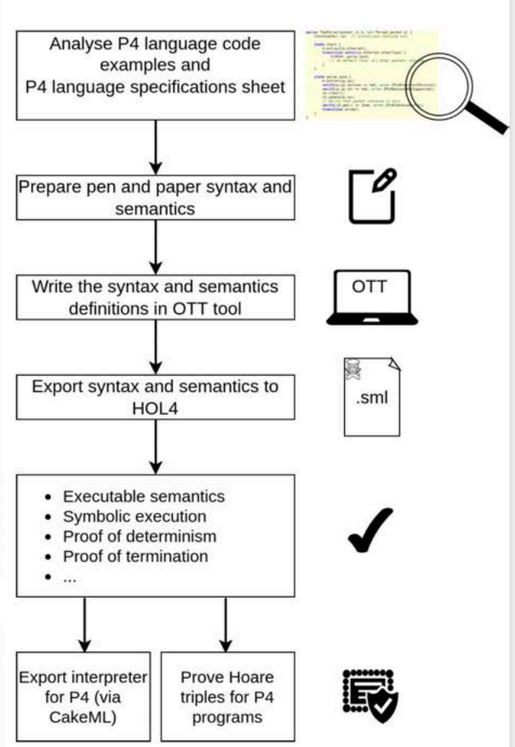

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

GPU Accelerated Sparse Representation of Light Fields

Light field imaging has been an omnipresent research topic during the last decade. With a growing interest, new techniques to capture, sample, and display light fields have been developed. The large amount of data that is produced during the capturing of light fields is a key challenge for acquisition and storage of light fields. We present a GPU-based compression technique based on multidimensional sparse representative. The main goal of the work presented here is to perform light field encoding, which includes an n-mode product, on the GPU and in real time.

AUTONOMOUS SYSTEMS (AS) | Page 97 B
Baravdish, Gabriel
Linköping University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

GPU Accelerated Sparse Representation of Light Fields

VISAPP 2019
Prague 25-27 feb

Gabriel Baravdish, Ehsan Miandji, and Jonas Unger
Department of Science and Technology, Linköping University

Goal

Light field imaging has been an omnipresent research topic during the last decade. With a growing interest, new techniques to capture, sample, and display light fields have been developed. The large amount of data that is produced during the capturing of light fields is a key challange for acquisition and storage of light fields. We present a GPU-based compression technique based on multidimensional sparse representation [3]. The main goal of the work presented here is to perform light field encoding, which includes a tensor-product, on the GPU and in real time. An overview of the method is presented in Figure 1 below.

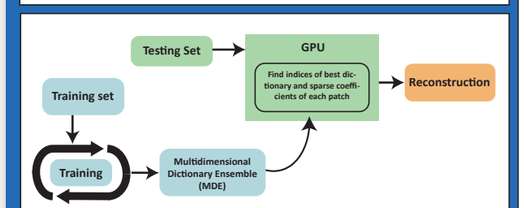

Figure 1: Flowchart of the meted. Firstly we train the MDE. Secondly, as an input to the encoding process, we compute the sparse coefficients and membership indices segmentally in parallel. Lastly we reconstruct the patches.

N-mode product on GPU

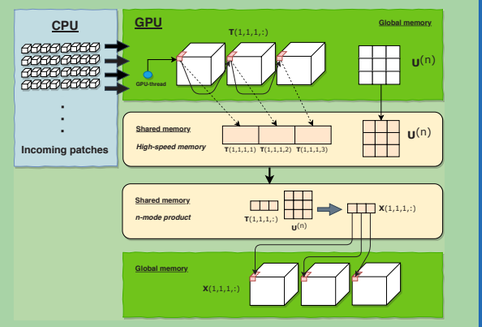

Figure 2: The computation of the tensor-product on the GPU.

We first load N number of patches, represented as tensors, on the GPU. We then let each thread traverse along the n-mode fiber and store the fiber on the fast on-chip memory called shared memory. The projection of a patch on a dictionary, see Equation 2, is computed by a vector-matrix product. This process can be simplified as a massive parallelized dot-product operation. We maximize thread collaboration and minimize bank conflicts by assigning a fiber to each thread. One of the biggest challanges with GPGPU computations is the overhead from memory transactions. We overcome this limitation by processing the patches segmentally in large batches.

Training

Initialiy from a training set of light field data, and as a one-time process, we train a multidimensional dimensional dictionary ensemble (MDE) which is obtained by a learning based approach [1].

Training set

Multidimensional Dictionary Ensemble - MDE

$$\{\mathbf{U}^{(1,k)},\dots,\mathbf{U}^{(n,k)}\}_{k=1}^{K} \longrightarrow$$

An MDE is obtained by minimizing:

$$\min_{\mathbf{U}^{j,k},\boldsymbol{\mathcal{S}}^{(i,k)},\mathbf{M}_{i,k}} \sum_{i=1}^{N_t}\sum_{k=1}^{K}\mathbf{M}_{i,k}\left\|\boldsymbol{\mathcal{L}}^{(i)}-\boldsymbol{\mathcal{S}}^{(i,k)}\times_1\mathbf{U}^{(1,k)}\cdots\times_n\mathbf{U}^{(n,k)}\right\|_F^2$$

subject to

$$(\mathbf{U}^{(j,k)})^T\mathbf{U}^{(j,k)}=\mathbf{I}, \qquad \forall k=1,\dots,K, \quad \forall j=1,\dots,n, \tag{1}$$

$$\left\|\boldsymbol{\mathcal{S}}^{(i,k)}\right\|_0\le\tau_l$$

$$\sum_{k=1}^{K}\mathbf{M}_{ik}=1, \qquad \forall i=1,\dots,N_t$$

GPU Encoding

With the dictionary ensemble, we aim to find the index of the basis where a data point has its most sparse representation with the least error under one of the dictionaries, called the membership index. We use a greedy algorithm [2,3] in order to compute the membership indices segmentally for all data points in parallel on the GPU.

$$[\{\boldsymbol{\mathcal{S}}^{(i)}\}_{i=1}^{N_t},\mathbf{m},\mathbf{z}]=Test(\{\boldsymbol{\mathcal{T}}^{(i)}\}_{i=1}^{N_t},\{\mathbf{U}^{(1,k)},\dots,\mathbf{U}^{(n,k)}\}_{k=1}^{K},\epsilon,\tau_l)$$

We project data points of a testing set denoted as $\boldsymbol{\mathcal{T}}^{(i)}$ onto the dictionaries by parallelizing the tensor product on the GPU, see Equation 2.

Tensor-product $\quad \boldsymbol{\mathcal{X}}^{(i)}=\boldsymbol{\mathcal{T}}^{(i)}\times_1\mathbf{U}^{(1,k)}\cdots\times_n\mathbf{U}^{(n,k)}$ (2)

Finally, we find and store the sparse coefficents and their locations in order to reconstruct the data, see Figure 2.

Results

We maintain the quality of the data from [1] and gain more than a tenfold speedup for the encoding process. The results for the GPU-implementation were achieved with a GTX Tian Xp. We compare our results against a CPU-implementation, see Fig. 3, where we used four Xeon E7-4870 with a total of 40 cores.

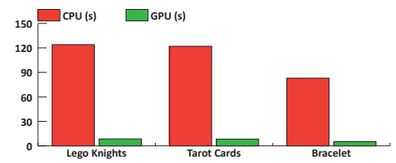

Figure 3: The measured time for the Lego Knights data set with K = 64 dictionaries, sparsity τ = 300 and threshold ε = 5 × 10−5. Further we have K = 64, τ = 390 and ε = 7 × 10−5 for the Tarot Cards. Lastly, we have K = 64, τ = 412 and ε = 5×10−5 for the Bracelet data set.

References

[1] Miandji, Ehsan and Hajisharif, Saghi and Unger, Jonas (2018). A Unified Framework for Compression and Compressed Sensing of Light Fields and Light Field Videos. ACM Transactions on Graphics, ACM, Accepted.

[2] Miandji, E., Kronander, J., and Unger, J. (2015). Compressive image reconstruction in reduced union of subspaces. Comput. Graph. Forum, 34(2):33–44.,

[3] Miandji, E., Kronander, J., and Unger, J. (2013). Learning based compression of surface light fields for realtime rendering of global illumination scenes. In SIGGRAPH Asia 2013 Technical Briefs, SA '13, pages 24:1–24:4, New York, NY, USA. ACM.

Batkovic, Ivo
Chalmers / Zenseact AB

Using MPC to Enable Safe Autonomous Driving

The rapid development of autonomous driving technologies in the past decades has been driven by the objectives of enabling safer and more efficient transportation. However, in order to enable such automated systems to be deployed on a global scale, problems regarding safety must be addressed. In particular, a self-driving vehicle must be able to safely interact with a surrounding environment consisting of other road users, whose intentions cannot be perfectly known. In this poster, we briefly mention how Model Predictive Control (MPC) can be used to ensure safe autonomous driving in uncertain environments.

AUTONOMOUS SYSTEMS (AS)

Batkovic, Ivo
Chalmers / Zenseact AB

Towards Real-world Editable 3D Maps Using Deep Learning

Various mapping frameworks used in robotics allow to build dense 3D world representations. Most of these representations use monolithic data structures such as octrees, point clouds, or meshes. While such representations are well suited for navigation tasks, interaction with such representations is difficult. An important step towards more interactive maps is to infer the 3D shape and pose of objects in the scene from partial observations. With SDFEst we propose an analysis-by-synthesis pipeline for joint pose and shape estimation using signed distance fields. The pipeline combines an initialization network, a generative shape model and a differentiable renderer to enable joint estimation of 7-DoF pose and shape from RGB-D images.

AUTONOMOUS SYSTEMS (AS) | Page 99 B
Bruns, Leonard
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Towards Real-world Editable 3D Maps Using Deep Learning

Leonard Bruns, KTH Royal Institute of Technology
Robotics, Perception and Learning

Motivation & Research Goals

Various mapping frameworks used in robotics allow to build dense 3D world representations. Most of these representations use monolithic data structures such as octrees, point clouds, or meshes. While such representations are well suited for navigation tasks, interaction with such representations is difficult. An important step towards more interactive maps is to infer the 3D shape and pose of objects in the scene from partial observations. With **SDFEst** we propose an analysis-by-synthesis pipeline for joint pose and shape estimation using signed distance fields. The pipeline combines an initialization network, a generative shape model and a differentiable renderer to enable joint estimation of 7-DoF pose and shape from RGB-D images.

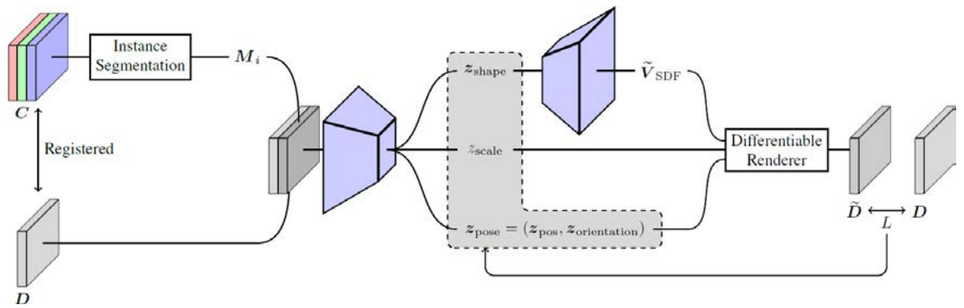

Introduction

We propose an analysis-by-synthesis pipeline for categorical shape and pose estimation. The pipeline consists of three main components: an initialization network, a generative shape model, and a differentiable renderer. The method currently works on a per-category level and only requires a collection of aligned meshes to be trained. No real-world data annotation is required.

Method Overview

Generative Shape Model We train a variational autoencoder (VAE) to compress the shape. Shape is encoded as discretized signed distance fields (SDFs).
Initialization Network We train a neural network which regresses position, orientation, scale and the latent shape of the object from a partial observation. The network is trained in a supervised manner on synthetic data generated from the VAE.
Differentiable Renderer We use a differentiable renderer inspired by SDFDiff [1] to render a posed, discretized SDF, while obtaining gradients for the pose, scale, and SDF.
Iterative Optimization From the initial estimate, we start an iterative optimization procedure by decoding the signed distance field, rendering it in the current pose and formulating a loss comparing the current estimate with the measured depth.

Results

We evaluate our approach on various synthetic and real-world datasets. Below we show qualitative results on the RGB-D Object dataset and the Redwood dataset, which includes handheld object in arbitrary orientations.

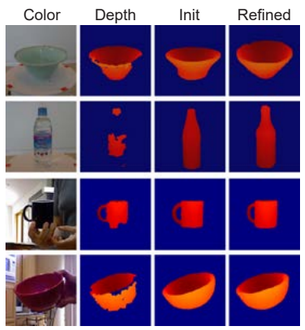

References

[1] Y. Jiang, et al. SDFDiff: Differentiable rendering of signed distance fields for 3d shape optimization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1251-1261, 2020.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Feature Extraction from Transaction Graph

Banks are required to analyse large transaction datasets as a part of the fight against financial crime. Today, this analysis is either performed manually by domain experts or using expensive feature engineering. As part of my PhD, I investigate how vector representations can be learned in an unsupervised way from transaction data. I here present a published method for learning node features from transaction amounts, and concurrent work on extracting vector features from the graph structure.

AUTONOMOUS SYSTEMS [AS] | Page 100 B
Ceylan, Ciwan
KTH / SEB
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Feature Extraction from Transaction Graphs

Ciwan Ceylan, KTH Royal Institute of Technology
Division of Robotics, Perception and Learning, SEB
Main advisor: Danica Kragic

Motivation & Research goals

Each year a vast amount of dirty money is laundered through the financial systems and financial institutions are under pressure to address this. The problem of discovering such schemes nevertheless remains very challenging from a machine learning perspective due to lack of labelled data, large data volumes and data secrecy. My research relates to advanced analytics for detecting financial crime occurring in transaction networks. Specifically, I'm interested in how recent advances in machine learning, network science and graph signal processing can be applied to financial transaction data in order to discovery anomalous customer behaviour, which could indicate illegal activity. A crucial first step is to develop methods which can extract representative vector features from transaction graphs.

Transaction Networks

Row	block_timestamp	from_address	to_address	amount
1	2018-06-01 07:06:17 UTC	0x8e5d083f	0x30727e88	150000.0
2	2018-06-01 07:12:35 UTC	0x30727e88	0x876eabf4	149999.0
3	2018-06-01 07:16:15 UTC	0x876eabf4	0x742d35cc	135130.0
4	2018-06-06 06:53:46 UTC	0x8e5d083f	0x052af704	259457.0
5	2018-06-08 06:33:24 UTC	0x53a46c10	0x6f0881f9	188087.0
6	2018-06-12 16:42:34 UTC	0x052af704	0x36abd00f	150000.0

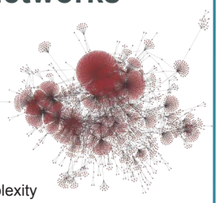

Network structure -- topological complexity

Timestamps -- dynamical complexity

Transferred amount -- data specific complexity

Given transactions observed in a time-window, a transaction network is a directed, weighted and possibly attributed graph, with an additional flow $\boldsymbol{F}$ on the edges:

$$G_T = (\boldsymbol{V}, \boldsymbol{E}, \boldsymbol{W}, A, \boldsymbol{F})$$

For application of anomaly detection algorithms on graph nodes, a common approach is to represent the nodes as D-dimensional embedding vectors [1].

$$f(G_T) = X_T \qquad X_T \in \mathbb{R}^{|V| \times D}$$

The mapping f should be find in an **unsupervised** way.

Network Structure Embeddings

Two types of node embeddings

In the literature, one finds two types of node embeddings.
Adjacency embeddings are similar for nodes close in the network [2, 3].
Role embeddings are similar for nodes with locally similar topology [4, 5].
Most SotA role embeddings are only defined for undirected graphs.

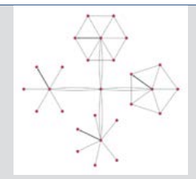

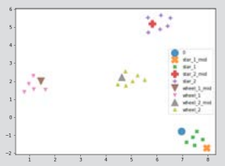

NetMF [3]

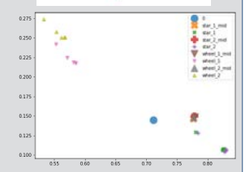

GraphWave [4]

Gated Gradient Model

The network flow $\boldsymbol{F}$ can be captured as node feature vectors. This can be done using either feature engineering (f.e.) or a gradient model (grad) which learns scalar potentials for each node in the network.

The **gated gradient model [6]** extends the gradient model so that vector potentials can be learned instead. This is achieved by introducing a gate function:

$$f^{(ij)} = \boldsymbol{\sigma}\left(\boldsymbol{u}^{(i)}, \boldsymbol{u}^{(j)}\right) \cdot \left(\boldsymbol{z}^{(j)} - \boldsymbol{z}^{(i)}\right)$$

This model is evaluated by its flow prediction performance on a subgraph of the Ethereum transaction graph. The results generalize to bank internal transaction data.

GATED	GRAD	F.E.+DNN2	N2V+DNN2	KUMAR ET.AL.	ZEROS	GRAD (MSE)
9.34	10.39	13.10	11.25	14.19	11.61	396.63

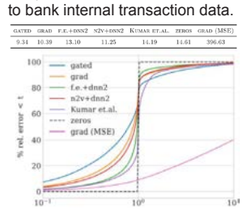

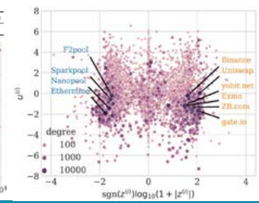

References

1. Akoglu, Leman, Hanghang Tong, and Danai Koutra. "Graph based anomaly detection and description: a survey." *Data mining and knowledge discovery* 29.3 (2015): 626-688.
2. Grover, Aditya, and Jure Leskovec. "node2vec: Scalable feature learning for networks." *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 2016.
3. Qiu, Jiezhong, et al. "Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec." *Proceedings of the eleventh ACM international conference on web search and data mining*. 2018.
4. Donnat, Claire, et al. "Learning structural node embeddings via diffusion wavelets." *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2018.
5. Qiu, Jiezhong, et al. "Gcc: Graph contrastive coding for graph neural network pre-training." *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2020.
6. Ceylan, Ciwan, Salla Franzén, and Florian T. Pokorny. "Learning Node Representations Using Stationary Flow Prediction on Large Payment and Cash Transaction Networks." *International Conference on Machine Learning*. PMLR, 2021.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Decentralized Control of Dynamical Systems under Signal Temporal Logic Specifications

Autonomous systems need often to perform a variety of complex tasks at dynamic environments within certain time intervals . Examples of such tasks could be ” reaching a known area within 5 sec” or ” move with other agents in formation between 5 and 10 sec until the leader agent reaches a safety area”. Each task may often evolve several agents that need to cooperatively design their future actions towards ensuring its satisfaction.

Nevertheless, when agents are working in large environments , communication among them might be difficult , costly or hard to establish. To that end, we propose a decentralized control framework that allows the satisfaction of a global formula with no need of communication. In our work we consider a set of complex tasks expressed as signal temporal logic formulas (STL), the satisfaction of which may depend on several or all agents in the team. As a first step, we decompose the global formula into local formulas whose satisfaction depends on given sub teams of agents using a convex optimization approach. Then, a receding horizon scheme (RHS) is proposed ensuring satisfaction of the local formulas and hence , satisfaction of the global formula . The proposed method is applied in a formation control example.

Decentralized Control of Dynamical Systems under Signal Temporal Logic Specifications

Maria Charitidou, Dimos V. Dimarogonas, KTH
Division of Decision and Control Systems, EECS

Abstract

Autonomous systems need often to perform a variety of complex tasks at dynamic environments within certain time intervals. Examples of such tasks could be ”reaching a known area within 5 sec” or ”move with other agents in formation between 5 and 10 sec until the leader agent reaches a safety area”. Each task may often evolve several agents that need to cooperatively design their future actions towards ensuring its satisfaction. Nevertheless, when agents are working in large environments, communication among them might be difficult, costly or hard to establish. To that end, we propose a decentralized control framework that allows the satisfaction of a global formula with no need of communication. In our work we consider a set of complex tasks expressed as signal temporal logic formulas (STL), the satisfaction of which may depend on several or all agents in the team. As a first step, we decompose the global formula into local formulas whose satisfaction depends on given sub-teams of agents using a convex optimization approach. Then, a receding horizon scheme (RHS) is proposed ensuring satisfaction of the local formulas and hence, satisfaction of the global formula. The proposed method is applied in a formation control example.

1. STL Decomposition

Signal Temporal logic is a specification language defined over continuous time signals. Let $\mu \in \{\bot,\top\}$ be a predicate defined after the evaluation of a continuously differentiable predicate function $h\colon \mathbb{R}^n \to \mathbb{R}$ as follows:

$$\mu = \begin{cases} \top, & h(x) \geq 0 \\ \bot, & h(x) < 0 \end{cases}.$$

In our work, we consider a restricted STL fragment defined as follows:

$$\varphi = G_{[a,b]}\mu \mid F_{[a,b]}\mu,$$
$$\phi = \wedge_{i=1}^{p} \varphi_i,$$

where μ is a predicate and $[a,b] \in \mathbb{R}_{\geq 0}$. We consider a global formula ϕ and a set of disjoint subteams of agents $\mathcal{V}_l, l = 1,\dots,v$ with $\cup_{l=1}^{v} \mathcal{V}_l = \mathcal{V}$. Then, for every $\varphi_i, \mathrm{i} = 1,\dots,\mathrm{p}$, let $V_i \subseteq \{1,\dots,v\}$ denote the indices of the subteams including at least one agent contributing to the satisfaction of φ_i and $\bar{z}_i^l \in \bar{Z}_i^l$ are the states of $\mathcal{V}_l$ satisfying $h_i(x) = h\left(\bar{z}_i^{l_1},\dots,\bar{z}_i^{l_w}\right)$, where $l_1,\dots,l_w \in V_i$ and $\bar{Z}_i^l$ a compact, convex, nonempty set. Then, the local formulas $\bar{\varphi}_i^l$, $l \in V_i$ are defined as follows [1]:

$$\bar{\varphi}_i^l = \mathcal{T}_{\left[a_i^l, b_i^l\right]}\left(h_i^l\left(\bar{z}_i^l;\, \theta_i^l\right) \geq 0\right),$$

where

$$\mathcal{T} = \begin{cases} F, & i \in I_F \\ G, & i \in I_G \end{cases},$$
$$\left[a_i^l, b_i^l\right] = \begin{cases} [a_i, b_i], & i \in I_G \\ [t_i, t_i], & i \in I_F \end{cases},$$

$$h_i^l\left(\bar{z}_i^l;\theta_i^l\right) = r_i^l - \left\|\bar{z}_i^l - c_i^l\right\|_\infty,$$

$t_i \in [a_i, b_i]$ and $\theta_i^l = (r_i^l, c_i^l) \in \mathbb{R}_{\geq 0} \times \bar{Z}_i^l$ are parameters found as the solution to the following optimization problem:

$$\max_{\theta_i^l} \sum_{l \in V_i} r_i^l$$
$$\bar{z}_i^l \in \left\{\xi \in \bar{Z}_i^l\colon \xi(\eta) = r_i^l + c_i^l(\eta) \; or \; \xi(\eta) = -r_i^l + c_i^l(\eta)\right\},$$
$$\bar{z}_i = \left[\bar{z}_i^l\right]_{l \in V_i} \in int\{x \in X\colon h_i(x) \geq 0\},$$
$$\theta_i^l = \left(r_i^l, c_i^l\right) \in \mathbb{R}_{\geq 0} \times \bar{Z}_i^l,$$

where $\xi(\eta)$ is the η-th element of ξ and X is a known, compact, convex, nonempty set. Then, the local formula corresponding to $\mathcal{V}_l$ is defined as:

$$\varphi^l = \wedge_{i \in J_l} \bar{\varphi}_i^l.$$

2. Results

Given each local formula φ^l, we encode the STL constraints using the barrier function:

$$b(z_l, t) = -\ln\left(\sum_{i \in J_l} \exp\left(-b_i\left(\bar{z}_i^l, t\right)\right)\right),$$

where $b_i\left(\bar{z}_i^l, t\right) = -\gamma_i(t) + h_i^l\left(\bar{z}_i^l; \theta_i^l\right)$ and $\gamma_i\colon \mathbb{R}_{\geq 0} \to \mathbb{R}$ are performance functions to be designed according to [2] ensuring satisfaction of all $\bar{\varphi}_i^l, i \in J_l$ with a minimum robustness r. Then, satisfaction of φ^l is ensured, when $b(z_l, t)$ remains nonnegative for every $t \geq 0$. Given the dynamical system:

$$\dot{z}_l = Az_l + Bu_l,$$

we may define the trajectories of agents in $\mathcal{V}_l$ using the receding horizon control scheme proposed in [3].

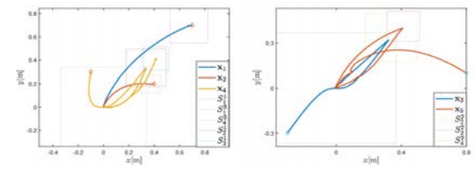

The proposed approach is implemented in a formation control problem of 5 agents. We consider 3 sub-teams $\mathcal{V}_1 = \{1,4\}, \mathcal{V}_2 = \{3,5\}$ and $\mathcal{V}_3 = \{2\}$ and the STL formula ϕ defined as:

$$\phi = G_{[0,2]}(\|x_1 - x_2 - p_x\|^2 \leq 0.1) \wedge G_{[2.5,4]}(\|x_3 - x_4\|^2 \leq 0.2)$$
$$\wedge \; F_{[3,7]}\left(\|x_5 - x_4\|^2_{P_1} \leq 0.2\right) \wedge F_{[8,10]}\|x_5 - x_2\|^2_{P_2} \leq 0.25).$$

The STL formulas are decomposed to local formulas φ^l. The trajectories of the agents, shown in the Figure satisfy the local formulas. This ensures the satisfaction of ϕ with minimum robustness 0.005.

References

1. M. Charitidou, D. V. Dimarogonas, "*Signal Temporal Logic Task Decomposition via Convex Optimization*", IEEE Control Systems Letters, vol. 6, pp. 1238-1243, 2022.
2. L. Lindemann and D. V. Dimarogonas, "*Decentralized Control Barrier Functions for Coupled Multi-Agent Systems under Signal Temporal Logic Tasks*", European Control Conference, 2019, pp. 89-94.
3. M. Charitidou, D. V. Dimarogonas, " *Barrier Function-based Model Predictive Control under Signal Temporal Logic Specifications*", European Control Conference, 2021.

Sybil-Based Attacks on Google Maps or How to Forge the Image of City Life

Participatory sensing (PS) applications collect all sorts of data by many users to maintain up-to-date data on everyday life, contributing to our well-being. Beyond occasional faults, it is often assumed that users are benign, thus strong security is not deployed. Controlling multiple users, an attacker can submit a large volume of forged data to dominate the PS-collected data. The result can be outright manipulation of the sensing process. We showcase the importance of this issue by selecting one of the most popular applications, Google Maps. The attacker model in our system is modest yet effective and efficient that are Sybil-based, leveraging non-existing, fake users. We instantiate automated attacks we term script- and emulator-based. The former submits crafted traffic in volume to manipulate the application data. The latter trades-off attack efficiency for increased versatility to attack other Google Maps features. We complete this investigation with human-based false data injection. This is the motivation of this work: to raise awareness on such a vulnerability and risk and improve the trustworthiness of such a popular application. We responsibly disclosed our findings to Google that acknowledged the issue and granted a reward.

Optimal Coordination of Mixed-Traffic Vehicles

In the context of mixed-traffic, the presence of Human-Driven Vehicles (HDVs) can pose several challenges to vehicles coordination due to uncertain, non-cooperative behaviors. In this work, we present an optimal control-based strategy for handling the HDVs by exploiting and coordinating Connected and Automated Vehicles (CAVs) forming the so-called mixed-platoons. A timeslot-based approach is used to schedule the vehicles that are going to occupy any intersection or roundabouts, with respect to safety requirements and physical limitations. In addition, we conduct a study and analyze the impact of human drivers' uncertainties in vehicles coordination.

AUTONOMOUS SYSTEMS (AS) | Page 103 B
Faris, Muhammad
Chalmers
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Optimal Coordination of Mixed-Traffic Vehicles

Muhammad Faris, Chalmers
Division of Systems and Control
Main supervisor: Paolo Falcone
CHALMERS UNIVERSITY OF TECHNOLOGY

Abstract

In the context of mixed-traffic, the presence of Human-Driven Vehicles (HDVs) can pose several challenges to vehicles coordination due to **uncertain, non-cooperative behaviors**. In this work, we present an **optimal control-based strategy** for handling the HDVs by exploiting and coordinating Connected and Automated Vehicles (CAVs) forming the so-called **mixed-platoons**. A timeslot-based approach is used to schedule the vehicles that are going to occupy any intersection or roundabouts, with respect to safety requirements and physical limitations. In addition, we conduct a study and analyze the impact of human drivers' uncertainties in vehicles coordination.

Mixed-Platoons

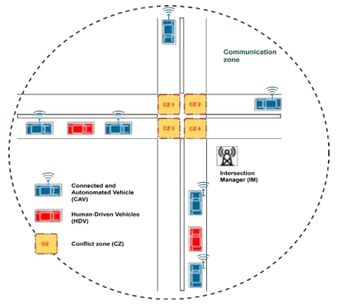

We consider the following

- **High-penetration** market rate of CAVs hence automated unsignalized intersections can be realized
- **Small** numbers of HDVs are squeezed in between **high** numbers of CAVs. CAVs are utilized as traffic **actuators** and **sensors** to regulate the HDVs.
- **Timeslot-based** intersection occupancy and safe distance approaches are applied to **avoid collisions** among vehicles with conflicting trajectories

Control Methods

The vehicles coordination problem under consideration can be formally described in a constrained optimization problem

$$\min_{U,T} \; J(X)$$

$$\text{s.t. } x^+ = f(x,u),$$
$$X \in \mathcal{X}, U \in \mathcal{U}, T \in \mathcal{T}$$

... that can be recast as a receding horizon optimal control problem. Particularly, we propose a bi-level optimization based on **Model Predictive Control (MPC)** [1].

References

[1] J. B. Rawlings, D. Q. Mayne, and M. M. Moritz, *Model predictive control: Theory, Computation, and Design 2nd Edition*, vol. 197. 2019.
[2] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*, vol. 1, no. 4. 2017.

Reachability Analysis

We study the impact of uncertainties from HDVs on the vehicles coordination problem using **reachability tools** [2]

$$\mathrm{K}_t(\mathrm{S},\mathrm{W}) = \mathrm{Pre}(\mathrm{K}_{t-1}(\mathrm{S},\mathrm{W}),\mathrm{W}) \cap \mathrm{X}$$

In particular, we evaluate the initial states feasible set of a free-driving CAV against a preceding mixed-platoon in terms of its ability to maintain safety, e.g., braking, under additive uncertainties of a human driver.

Results

We perform numerical simulations of mixed-traffic vehicles coordination using two different HDV prediction models:
(1) car-following and (2) constant measurements in a small-scale setting. The result is given in the following table

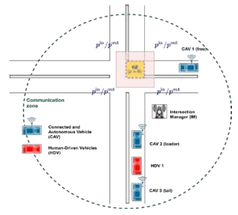

Model	Clearance time
(1) Car-following	2.9 sec
(2) Constant measurements	2.4 sec

Conclusions:

- The algorithm successfully prevents any collision or platoon cut-in.
- Using a model like in (1) might take a much longer time to clear an intersection due to misleading assumptions.

Furthermore, we carry out reachability analysis and present feasible sets of two different cases: (a) one with added uncertainty and (b) the other is not. The results are shown in the figures below

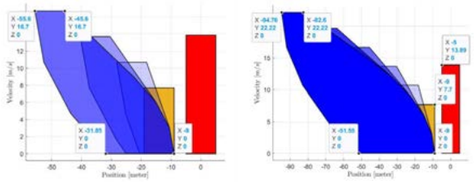

(a) (b)

Conclusions:
Uncertainties shrink the feasible sets of initial states. The higher the uncertainty, the smaller the area.

The fundamental lemma based on second order moments

In this paper, we propose variations of the fundamental lemma that utilize second-order moments such as correlation functions in the time domain and power spectra in the frequency domain.We believe that using a formulation with estimated correlation coefficients is suitable for data compression, and possibly can reduce noise.
Also, the formulations in the frequency domain can enable modeling of a system in a frequency region of interest.

Willems' fundamental lemma based on second-order moments

Mina Ferizbegovic, Håkan Hjalmarsson, Per Mattsson, Thomas B. Schön

KTH VETENSKAP OCH KONST · UPPSALA UNIVERSITET · WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Summary and contributions

We propose:

- a variation of Willems' fundamental lemma based on the correlation functions, which can be useful for large noisy datasets,
- a variation of Willems' fundamental lemma using spectra, which can be useful for modeling in a frequency region of interest.

Willems' fundamental lemma

Willems' fundamental lemma[1]: it is possible to describe all trajectories of a linear deterministic system from a single input-output data trajectory under some assumptions.

We assume:

- collected data (u^d, y^d) of length T
- the system is controllable
- input u^d is persistently exciting of order $L + n_x$:

$$\begin{bmatrix} u_0^d & u_1^d & \cdots & u_{T-L-n_x}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L+n_x-1}^d & u_{L+n_x}^d & \cdots & u_{T-1}^d \end{bmatrix} \text{ full row rank}$$

For any trajectory (u, y) of length L of the system there exists $g \in \mathbb{R}^M$ ($M = T - L$)

$$\begin{bmatrix} u_0 \\ \vdots \\ u_{L-1} \\ \hline y_0 \\ \vdots \\ y_{L-1} \end{bmatrix} = \begin{bmatrix} u_0^d & u_1^d & \cdots & u_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L-1}^d & u_L^d & \cdots & u_{T-1}^d \\ \hline y_0^d & y_1^d & \cdots & y_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ y_{L-1}^d & y_L^d & \cdots & y_{T-1}^d \end{bmatrix} g$$

Data driven simulation

Given: $\underbrace{u_{\text{ini}} = \begin{bmatrix} u_{-L_0} \\ \vdots \\ u_{-1} \end{bmatrix}, y_{\text{ini}} = \begin{bmatrix} y_{-L_0} \\ \vdots \\ y_{-1} \end{bmatrix}}_{\text{estimate } x_0}, \underbrace{u = \begin{bmatrix} u_0 \\ \vdots \\ u_{L'-1} \end{bmatrix}}_{\text{predict output}}$, data (u^d, y^d)

Find: $y = \text{col}\,(y_0 \dots y_{L'-1})$

Classic: Find A, B, C, D. Using $u_{\text{ini}}, y_{\text{ini}} \xrightarrow{A,B,C,D} x_0$, $u \xrightarrow{A,B,C,D,x_0} y$

Assumptions: the system is minimal, u^d is persistently exciting, lag $\leq L_0$

$$\begin{bmatrix} U_p \\ \hline U_f \end{bmatrix} = \begin{bmatrix} u_0^d & u_1^d & \cdots & u_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L_0-1}^d & u_{L_0}^d & \cdots & u_{M+L_0-2}^d \\ \hline u_{L_0}^d & u_{L_0+1}^d & \cdots & u_{M+L_0-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L-1}^d & u_L^d & \cdots & u_{T-1}^d \end{bmatrix} \begin{matrix} Y_p, Y_f \\ \text{similarly as} \\ U_p, U_f \end{matrix}$$

Solution[2]: Compute a solution of $\begin{bmatrix} U_p \\ Y_p \\ U_f \end{bmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \end{bmatrix} \implies y = Y_f g$

Data driven predictive control

- A data-driven simulation used instead of a traditional model in the MPC algorithm.
- We obtain the following optimization problem:

$$\min_{g,u,y} \quad \sum_{k=0}^{L'-1} (y_k - r_{t+k})^\top Q (y_k - r_{t+k}) + u_k^\top R u_k + \lambda \times \text{regularization}$$

$$\text{offline noise} \to \begin{bmatrix} U_p \\ Y_p \\ U_f \\ Y_f \end{bmatrix} g = \begin{bmatrix} u_{\text{ini}} \\ y_{\text{ini}} \\ u \\ y \end{bmatrix} \leftarrow \text{online noise}, \quad u_k \in \mathcal{U}, \quad y_k \in \mathcal{Y}$$

offline data · online data

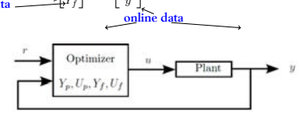

- We optimize g online (receding horizon).
- If T is large, $\dim(g) \approx T \implies$ high dimensional optimization problem. **Data compression!** [3]

Willems' Lemma Using Correlation Functions

- autocorrelation $R_{uu}(\tau)$ of the input u_t^d
- cross-correlation $R_{yu}(\tau)$ between output y_t^d and input u_t^d signals
- the system is controllable and stable
- u^d is persistently exciting of order $L + n_x$

$$\begin{bmatrix} R_{uu}(0) & R_{uu}(1) & \cdots & R_{uu}(L-1) \\ R_{uu}(1) & R_{uu}(0) & \cdots & R_{uu}(L-2) \\ \vdots & \vdots & \ddots & \vdots \\ R_{uu}(L-1) & R_{uu}(L-2) & \cdots & R_{uu}(0) \end{bmatrix} > 0$$

For any trajectory (u, y) of length L of the system there exists $g \in \mathbb{R}^M$

$$\begin{bmatrix} u_0 \\ \vdots \\ u_{L-1} \\ \hline y_0 \\ \vdots \\ y_{L-1} \end{bmatrix} = \begin{bmatrix} R_{uu}(n_x) & R_{uu}(n_x-1) & \cdots & R_{uu}(n_x+1-M) \\ R_{uu}(n_x+1) & R_{uu}(n_x) & \cdots & R_{uu}(n_x+2-M) \\ \vdots & \vdots & \ddots & \vdots \\ R_{uu}(n_x+L-1) & R_{uu}(n_x+L-2) & \cdots & R_{uu}(n_x+L-M) \\ \hline R_{yu}(n_x) & R_{yu}(n_x-1) & \cdots & R_{yu}(n_x+1-M) \\ R_{yu}(n_x+1) & R_{yu}(n_x) & \cdots & R_{yu}(n_x+2-M) \\ \vdots & \vdots & \ddots & \vdots \\ R_{yu}(n_x+L-1) & R_{yu}(n_x+L-2) & \cdots & R_{yu}(n_x+L-M) \end{bmatrix} g.$$

Comparison to the Willems' Lemma

- We have assumed that the input is a stationary process and that the system is stable.
- Data compression (reducing dimension of g)
- Ongoing work in a finite-time domain using the following approximation of correlation coefficients: $R_{xy}(k) \approx \frac{1}{T-k} \sum_{i=0}^{T-k-1} x_{i+k} y_i^\top$
- Using sample correlations reduce the effect of noise in the data compared to using raw data.

Calculating correlation coefficients

Using that $\begin{bmatrix} x_0 & x_1 & \cdots & x_{M-1} \\ \hline u_0^d & u_1^d & \cdots & u_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L-1}^d & u_L^d & \cdots & u_{T-1}^d \end{bmatrix}$ is full row rank, we can obtain

$$\begin{bmatrix} R_{uu}(k) \\ \vdots \\ R_{uu}(k+L-1) \\ \hline R_{yu}(k) \\ \vdots \\ R_{yu}(k+L-1) \end{bmatrix} = \begin{bmatrix} u_0^d & u_1^d & \cdots & u_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L-1}^d & u_L^d & \cdots & u_{T-1}^d \\ \hline y_0^d & y_1^d & \cdots & y_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ y_{L-1}^d & y_L^d & \cdots & y_{T-1}^d \end{bmatrix} g.$$

Willems' Lemma Using Spectra

- spectrum $\Phi_{uu}(\omega)$ of the input u_t^d, at different M frequencies
- and the cross-spectrum $\Phi_{yu}(\omega)$ between output y_t^d and input u_t^d signals, at different M frequencies
- the system is controllable and stable
- full row rank

$$\begin{bmatrix} W_{L+n_x}(\omega_1) \otimes \Phi_{uu}(\omega_1) & \cdots & W_{L+n_x}(\omega_M) \otimes \Phi_{uu}(\omega_M) \end{bmatrix},$$

where $W_L(\omega) = \begin{bmatrix} 1 & e^{j\omega} & \cdots & e^{j\omega(L-1)} \end{bmatrix}^\top$

For any trajectory (u, y) of length L of the system there exists $g \in \mathbb{R}^M$

$$\begin{bmatrix} y_0 \\ \vdots \\ y_{L-1} \\ \hline u_0 \\ \vdots \\ u_{L-1} \end{bmatrix} = \begin{bmatrix} W_L(\omega_1) \otimes \Phi_{yu}(\omega_1) & \cdots & W_L(\omega_M) \otimes \Phi_{yu}(\omega_M) \\ \hline W_L(\omega_1) \otimes \Phi_{uu}(\omega_1) & \cdots & W_L(\omega_M) \otimes \Phi_{uu}(\omega_M) \end{bmatrix} g.$$

Future reading

[1] J. C. Willems, P. Rapisarda, I. Markovsky, and B. De Moor. A note on persistency of excitation. *Systems & Control Letters*, 2005.

[2] I. Markovsky and P. Rapisarda. Data-driven simulation and control. *International Journal of Control*, 2008.

[3] J. Coulson, J. Lygeros, and F. Dörfler. Data-enabled predictive control: In the shallows of the DeePC. In *18th European Control Conference (ECC)*, 2019.

[4] M. Ferizbegovic, H. Hjalmarsson, P. Mattsson, and T. B. Schön. Willems' fundamental lemma based on second-order moments. In *60th IEEE Conference on Decision and Control (CDC)*, 2021.

Forough, Javad
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Anomaly Detection and Countermeasures for Edge Clouds

The accelerated growth of the Internet of Things (IoT) and emerging 5G infrastructure has opened up opportunities to develop intelligent applications that transform data into business and societal value in plenty of application domains such as public services, intelligent transportation, augmented reality, industrial automation, and smart healthcare. The centralized cloud computing model has shown to have inherent problems when it comes to meet certain requirements of bandwidth-hungry or response-time-critical applications at the edge of the network. Thus, centralized clouds cannot provide services with high performance and reliability for such applications. Edge clouds are distributed computing infrastructures comprising edge nodes, fog nodes, and distant clouds, where the massive amount of data moves back-and-forth between the edge and distant cloud datacenters that concern data privacy and security issues. The ultimate goal of this project is to design, develop and deploy decentralized autonomous anomaly detection and countermeasures for ensuring performance and security in edge clouds using emerging machine learning models against unexpected service performance, security flaws and cyber attacks.

AUTONOMOUS SYSTEMS (AS) | Page 105 B

Forough, Javad
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Anomaly Detection and Countermeasures for Edge Clouds

Javad Forough, Umeå University
Department of Computing Science
Supervisor: Erik Elmroth
Co-supervisor: Monowar Bhuyan

UMEÅ UNIVERSITY

Motivation

The accelerated growth of the Internet of Things (IoT) and emerging 5G infrastructure has opened up opportunities to develop intelligent applications that transform data into business and societal value in plenty of application domains such as public services, intelligent transportation, augmented reality, industrial automation, and smart healthcare. The centralized cloud computing model has shown to have inherent problems when it comes to meet certain requirements of bandwidth-hungry or response-time-critical applications at the edge of the network. Thus, centralized clouds cannot provide services with high performance and reliability for such applications. Edge clouds are distributed computing infrastructures comprising edge nodes, fog nodes, and distant clouds, where the massive amount of data moves back-and-forth between the edge and distant cloud datacenters that concern data privacy and security issues. The ultimate goal of this project is to design, develop and deploy decentralized autonomous anomaly detection and countermeasures for ensuring performance and security in edge clouds using emerging machine learning models against unexpected service performance, security flaws and cyber attacks.

Works Done

- **Forough, J., Bhuyan, M., & Elmroth, E. (2021, August). Detection of VSI-DDoS Attacks on the Edge: A Sequential Modeling Approach. In *The 16th International Conference on Availability, Reliability and Security* (pp. 1-10).** (With 20.34% acceptance rate)
- **Forough, J., Bhuyan, M., & Elmroth, E. (2021, October). DELA: A Deep Ensemble Learning Approach for Cross-layer VSI-DDoS Detection on the Edge. *Submitted to The 37th ACM/SIGAPP Symposium On Applied Computing (SAC2022)***

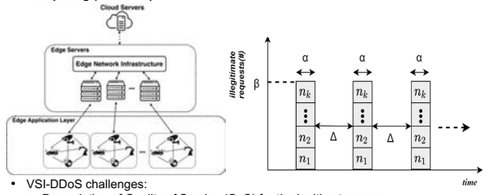

- VSI-DDoS challenges:
  - Degradation of Quality of Service (QoS) for the legitimate users
  - Hard to detect, as the pattern of the monitored metrics remains very similar to the one in non-attack periods
  - Impact is even worse when attack happens on different layers of the system

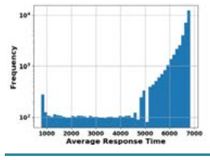

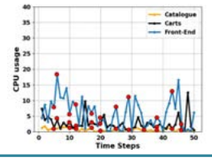

References

1. Shan, H., Wang, Q., & Yan, Q. (2017). Very short intermittent DDoS attacks in an unsaturated system. In International Conference on Security and Privacy in Communication Systems (pp. 45-66). Springer.
2. Park, J., Nyang, D. and Mohaisen, A., (2018). Timing is almost everything: Realistic evaluation of the very short intermittent ddos attacks. In *2018 16th Annual Conference on Privacy, Security and Trust (PST)* (pp. 1-10). IEEE.
3. Sherstinsky, A. (2020). Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network. *Physica D: Nonlinear Phenomena*, *404*, 132306.
4. Luong, M.T., Pham, H. and Manning, C.D., 2015. Effective approaches to attention-based neural machine translation. arXiv preprint arXiv:1508.04025.

Testbed Setup

- Container-based microservice application
- Docker and swarm for service deployment and orchestration
- Locust for normal load generation
- Apache Bench for VSI-DDoS implementation
- Keras for model implementation

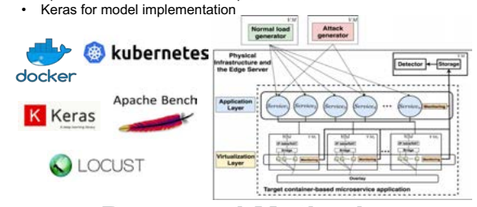

Proposed Methods

- Deep sequential modelling
- Long Short-Term Memory (LSTM) [3]
- Local attention layer [4]
- Ensemble learning

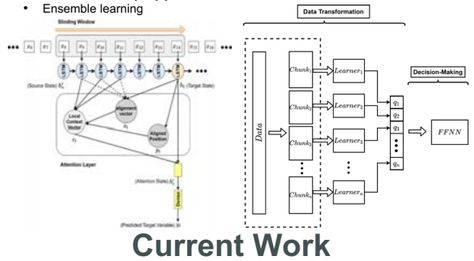

Current Work

- The Quality of Service (QoS) of microservice applications experienced by users on the edge clouds is always exposed to security and performance anomalies such as slow-http, DDoS, spike overload, resource failure, etc. To address such problems of anomalies, we need to first identify the root cause of the problem as a “Security” or “Performance” anomaly, to be able to mitigate the issues in the next step. Hence, in this work, we plan to develop a comprehensive Anomaly Detection System (ADS) which will be able to detect both security and performance anomalies.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Considerations for safety assurance of ADSs

Safety assurance of Automated Driving Systems (ADS) is arguably one of the largest outstanding challenges before large-scale deployment of such systems on public roads. In my research I focus on the aspects of providing, not only effective, but also efficient safety assurance of ADSs. Central to safety assurance is the compilation of a compelling safety (assurance) case that presents evidence-supported arguments for the system's safety fulfilment. I have investigated different ways to approach safety assurance and break down this task by using the Operational Design Domain (ODD), but also by looking at different assurance methods to support the safety case construction and maintenance.

AUTONOMOUS SYSTEMS (AS) | Page 106 B
Gyllenhammar, Magnus
Zenseact
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Considerations for safety assurance of ADSs

Magnus Gyllenhammar, Zenseact, KTH
Mechatronics

Abstract

Safety assurance of Automated Driving Systems (ADS) is arguably one of the largest outstanding challenges before large-scale deployment of such systems on public roads. In my research I focus on the aspects of providing, not only effective, but also efficient safety assurance of ADSs. Central to safety assurance is the compilation of a compelling safety (assurance) case that presents evidence-supported arguments for the system's safety fulfilment. I have investigated different ways to approach safety assurance and break down this task by using the Operational Design Domain (ODD) [1], but also by looking at different assurance methods to support the safety case construction and maintenance [2].

Operational Design Domain (ODD) for safety assurance

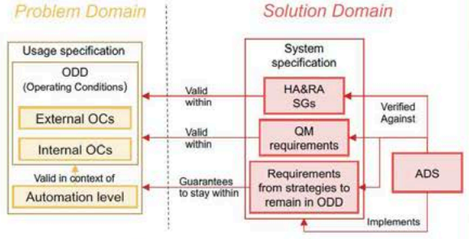

The Operational Design Domain (ODD) of an ADS has been proposed to limit or constrain the development, design, verification and validation activities by limiting the operations of the ADS to this ODD. This thus splits the problem of safety assurance into two parts:

- Designing, developing and providing evidence for a safe ADS within the limits of the ODD, and
- Ensuring that the ADS does not operate outside the ODD.

For the second aspect the following strategies have been proposed [1]:

I. Internal, inherent in ADS feature definition
II. External, checking mission when accepting strategic task
III. External, statistically defined spatial and temporal triggering conditions
IV. External, run-time measurable triggering conditions related to operating conditions

References

1. Gyllenhammar, Magnus, et al. "Towards an operational design domain that supports the safety argumentation of an automated driving system." *10th European Congress on Embedded Real Time Systems (ERTS 2020)*. 2020.
2. Gyllenhammar, Magnus, Carl Bergenhem, and Fredrik Warg. "ADS Safety Assurance–Future Directions." *CARS 2021 6th International Workshop on Critical Automotive Applications: Robustness & Safety*. 2021.
3. Gyllenhammar, Magnus, et al. "Minimal Risk Condition for Safety Assurance of Automated Driving Systems." *CARS 2021 6th International Workshop on Critical Automotive Applications: Robustness & Safety*. 2021.

Data and models for assurance

Throughout the development processes as well as the assurance there is need for using models. These models can either be based on conservative assumptions, such as using the laws of physics for determining possible actions of other traffic participants. On the other end of the spectrum, one could consider using predefined rules, with some presumptions about where and how they are applicable. However, both these extremes will result in a suboptimal system considering both safety and performance. A better approach is letting data support those models.

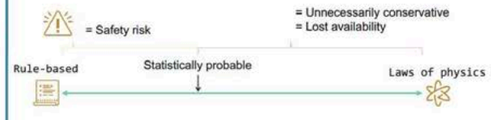

Minimal Risk Condition for safety assurance

To fulfil the ODD, one should ensure that the strategies (I) – (IV) is determined for all types of ODD exits. However, the system itself could also experience deterioration or degradation to its capabilities. In such cases it might not be possible to safely continue fulfilling the mission. If this is the case the system should go into MRC. To contextualise that decision, one could use the following decision hierarchy [3]:

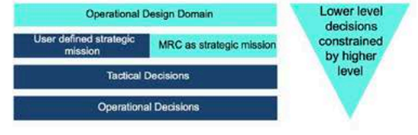

The key element of this is the split of the decisions on the strategic level. The mission of the ADS is either to fulfil the user defined mission or to go into MRC. This is in relation to the ODD at the top or in the case of degradations this ODD is exchanged with the Restricted Operational Domain (ROD).

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Unified task and motion planning

Many robotic applications involve both high-level (discrete) task planning and low-level (continuous) motion planning. Solving the two planning problems separately one after the other often leads to suboptimal solutions, or no feasible solution at all. This doctoral project aims at tightly integrating methods for task planning with methods for optimal-control-based motion planning in order to solve the task and motion planning problems simultaneously. This poster gives a brief background to the problem, presents the overarching research questions of the doctoral project and presents the work that is currently ongoing.

AUTONOMOUS SYSTEMS [AS] | Page 107 B
Hellander, Anja
Linköping University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Unified Task and Motion Planning

Anja Hellander, Linköping University
Department of Electrical Engineering

Abstract

Many robotic applications involve both high-level (discrete) task planning and low-level (continuous) motion planning. Solving the two planning problems separately one after the other often leads to suboptimal solutions, or no feasible solution at all. This doctoral project aims at tightly integrating methods for task planning with methods for optimal-control-based motion planning in order to solve the task and motion planning problems simultaneously. This poster gives a brief background to the problem, presents the overarching research questions of the doctoral project and presents the work that is currently ongoing.

Background

Task and motion planning are naturally interdependent in many robotic applications where robots have to perform both high-level task planning in order to achieve some objective as well as low-level motion planning in order to determine how to perform actual movements. Hierarchical approaches where the task and motion planning are performed separately often give suboptimal or even infeasible solutions. In order to achieve reliable task and motion planning it is therefore necessary to formulate and solve a single integrated planning problem where the task and motion planning problems are solved jointly.

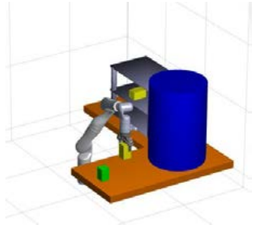

Example of a task and motion planning problem. The manipulator robot must perform task planning in order to determine which objects to pick and place, and motion planning in order to determine how to perform the operations.

Research questions

- How to extend an existing action language (e.g. PDDL) to include specifications of optimal control problems for dynamic systems?
- How to tightly integrate methods behind classical AI planners with methods behind motion planners using graph search and numerical optimal control?
- How to develop efficient heuristics for problems that tightly integrate task and motion planning?

Ongoing work

The current focus is developing a framework for unified task and motion planning that rather than only finding a feasible solution (if one exists) to a task and motion planning can also perform optimization of this solution (at least to some degree).

The ongoing work is therefore focused on:

- How to handle that the motion planning problem has continuous variables whereas the task planning problem is discrete. How should the discrete values be generated? In advance, during the search? Randomly or deterministically?
- The search will require calls to some function in order to determine if feasible motion plans exist or not, which will be expensive. How can the number of calls be reduced?
- How can optimization be integrated into already existing frameworks for task and motion planning?

Ongoing work: Drill planning

- **Setting:** A drill rig must drill a number of holes at given positions. Once a hole has been drilled, the rig cannot pass over it. The drill holes are positioned densely relative the drill rig size.
- **Problem:** Decide in which order the holes are to be drilled, and plan feasible paths between them for the rig to follow.
- **Our approach:** Discretize the configuration space of the drill rig. Graph search (backward) in a state space with state consisting of current position (hole), current (discretized) heading and previously drilled holes. Call to lattice-based motion planner to determine if a feasible path between two states exists.

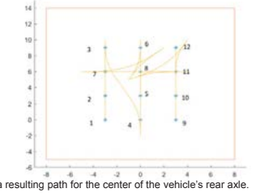

Example of a resulting path for the center of the vehicle's rear axle.

An Adaptive Approach for Task-Driven BCI Calibration

Brain-Computer Interfaces (BCI) use brain signals as inputs and machine learning algorithms to decipher the meaning of these. A BCI system needs to be calibrated before usage, i.e., the machine learning algorithm needs to be trained. The overall goal is to solve a task as fast as possible. The calibration can be terminated with an adaptive approach when the BCI system is good enough to solve the task. Here we present a structure for such a system and show some initial results.

AUTONOMOUS SYSTEMS (AS) | Page 108 B
Heskebeck, Frida
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

An Adaptive Approach for Task-Driven BCI Calibration

FRIDA HESKEBECK AND CAROLINA BERGELING – DEPARTMENT OF AUTOMATIC CONTROL

Background

Brain-Computer Interfaces (BCI) use brain signals as inputs and machine learning algorithms to decipher the meaning of these. A BCI system needs to be calibrated before usage, i.e., the machine learning algorithm needs to be trained. The calibration phase is often tedious for the user. Hence we **focus on improving the calibration** of BCI systems.

Conclusion

Our suggestion for an adaptive Brain-Computer Interface **automatically decides** whether to 1) **listen** to the brain and respond accordingly, using the current machine-learning algorithm to predict the meaning of the brain signals, or 2) **update** the machine learning algorithm to make better predictions in the future.

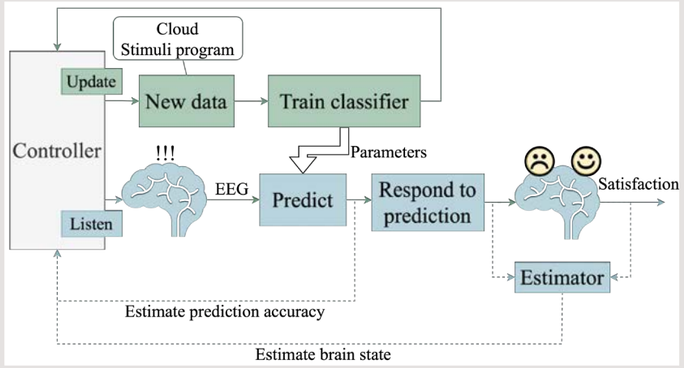

Recent work

- Aim: Fast task solving
  - Hearing aid - Attenuate sound source
- Controller decides action
  - *Update* – retrain machine learning algorithm with new data
  - *Listen* – predict user intent.

- Model learning rates (shown in the graphs).
  - Estimate p and q.
  - Model: First order system, fixed T, individual start and end points.
  - Across sessions and individuals.

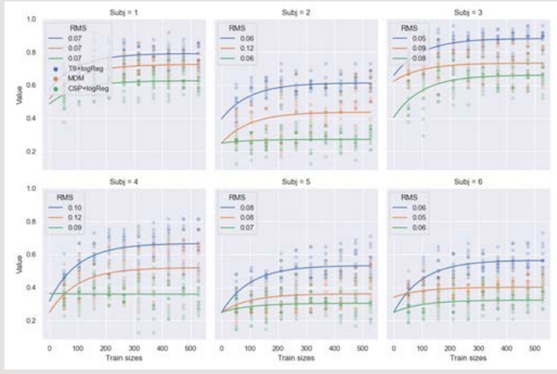

- Markov Decision Process (MDP) to find best action
  - If the inequality below is true, *update* is best action.
  - Parameters: Transition probabilities, rewards, and discount factor.

$$\frac{(1-p)r_0 + pr_1}{1-\gamma(1-p)} < r_2 + \gamma\frac{(1-q)r_3 + qr_4}{1-\gamma(1-q)}$$

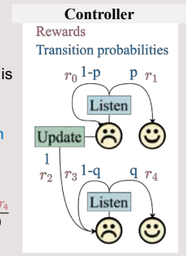

Future

- Resources for *update* action:
  - *Time* to do experiment.
  - *Computer power* to retrain machine learning model.
  - *Money* for accessing the cloud with previous data.
- Controller.
  - Extend Markov Decission Process.
  - Use some “error” based on brain state.
- Estimate prediction accuracy.
  - Without validation data?
- Estimate user’s brain state.
  - Use error potentials – a type of brain response.
  - Estimation with sporadical input from user.

Contact

Frida Heskebeck | frida.heskebeck@control.lth.se | heskebeck.com

ELLIIT WASP

Hynén Ulfsjöö, Carl
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Motion-planning and decision-making under uncertainty for heavy vehicles.

To safely maneuver a heavy vehicle in complex traffic situations, the uncertainty in the prediction of the surrounding vehicles must be considered during planning. In this poster a two-stage approach to this problem is presented that tightly couples a POMDP with scenario-based stochastic MPC, to be able to exploit the properties of both methods.
This is applied to a highway driving situation where the ego vehicle wants to overtake a vehicle in dense traffic, where the prediction of the environment is uncertain and there is uncertainty in how cooperative each driver is. The resulting two-stage motion planner is able to safely plan in this situation and the inclusion of the MPC-step is shown to drastically improve the solution from just using the POMDP.

AUTONOMOUS SYSTEMS (AS) | Page 109 B

Hynén Ulfsjöö, Carl
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Motion-planning and decision-making under uncertainty for heavy vehicles.

Carl Hynén Ulfsjöö, Linköping University
Dept. of Automatic Control, ISY
Supervisor: Prof. Daniel Axehill (LiU)

LINKÖPINGS UNIVERSITET

Introduction

To safely and efficiently maneuver a heavy vehicle in a complex traffic situation, the driver needs to perceive, interpret and predict the motion of multiple surrounding vehicles. Then based on that prediction it must choose an appropriate action that considers the large level of uncertainty in the prediction, without becoming overly conservative.

To realize this in a motion planner it should:

- take the uncertainty in prediction into account
- exploit interactions between the ego and surrounding vehicles
- make joint discrete and continuous decisions.

Method

The developed motion planner is based on a two-stage approach. First a general partially observable Markov decision process (POMDP) is solved, then the solution is used in the second stochastic model predictive control (MPC) step, which improves the solution. This results in a motion planner where the POMDP:

- makes discrete decisions
- handles general uncertainty in perception and prediction
- outputs a coarsely discretized control signal

and the MPC:

- improves the solution locally
- handles unimodal Gaussian uncertainty
- outputs a finely discretized control signal.

Partially observable Markov decision process

The POMDP models a decision process where the noisy system dynamics are known but the underlying state cannot directly be measured. It tries to find the optimal policy (π) that maps a probability distribution over the state-space to an action. The optimization problem that it solves can be written as:

$$
\begin{aligned}
\underset{\pi(\cdot)}{\text{maximize}} \quad & \mathbb{E}\left[\sum_{k=0}^{N} \gamma^k R(x_k, u_k)\right] \\
\text{subject to} \quad & x_{k+1} = f(x_k, u_k) + w_t(x_k, u_k) && \text{(prediction)} \\
& y_k = h(x_k, u_k) + v_k(x_k, u_k) && \text{(observation)} \\
& b_k \sim p(b_k | b_{k-1}, y_k, u_k) && \text{(belief propagation)} \\
& u_k \sim \pi(b_k), x_0 \sim b_0 \\
& x \in \mathcal{X}, u \in \mathcal{U}, y \in \mathcal{O}.
\end{aligned}
$$

The POMDP is solved using the online POMDP solver DESPOT that uses sampling to approximate the uncertainty, which converts the problem to a tree-search problem.
The POMDP solver can handle very general uncertainty and directly consider partial observability. However the resulting tree scales poorly with regards to $|\mathcal{U}|$, which in practice means that the control signal must be coarsely discretized.

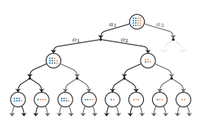

Stochastic model predictive control

The stochastic MPC step is introduced to compensate for the coarse discretization in the solution to the POMDP. A scenario-based stochastic MPC formulation is used, because of the multimodal nature of the prediction of the environment (a surrounding vehicle might or might not yield). This uses discrete scenarios to represent the different modes, and for each mode typical stochastic MPC techniques are used to represent the local uncertainty.
As the MPC step is based on the solution to the POMDP, the solution can be used in several ways to tighten the coupling between them.

- The scenarios in the MPC can be derived from the sampled scenarios in the POMDP, and scenario reduction techniques can be used to only include the most relevant scenarios.
- The solution to the POMDP can directly be used to warm start the optimization solver.
- The POMDP solution can be used to define nonanticipatory constraints in the MPC, which determine when different modes are indistinguishable.

Preliminary results

The motion-planning is applied to a highway driving situation where the ego vehicle wants to overtake a vehicle in dense traffic, where the prediction is uncertain and there is uncertainty in how cooperative each driver is. Despite the uncertainty, the developed motion planner can still find a safe plan that takes the uncertainty directly into account and can perform actions to gather information about the environment.
An example of this can be seen in the first figure to the right where the planner first commands a lane change at time 10 s as it believes that vehicle in the passing lane is cooperative enough. However, as the vehicle does not react cooperatively it postpones the lane change until after the vehicle has passed. In the second figure to the right the result of using the two-stage approach is shown. The blue vehicle is only using the POMDP solution while the green is using the improved solution and is therefore able to perform the lane change much faster.

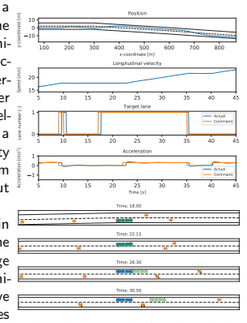

Conclusions and future work

Combining a POMDP with stochastic MPC makes it possible to exploit the best properties of both methods. The method shows promising results in experiments on a typical highway driving situation.
As future work the coupling between the two methods needs to be further investigated, additionally the implementation must be improved to make it real-time capable.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Combining verbal-HRI with Behavior Trees to disambiguate human demonstrations

Fast changing tasks in unpredictable, collaborative environments are typical for medium-small companies, where robotised applications are increasing. Thus, robot programs should be generated in short time with small effort, and the robot able to react dynamically to the environment. To address this, a method exists that combines context awareness and planning to learn Behavior Trees (BTs), a reactive policy representation that is becoming more popular in robotics. The method allows to learn BTs from human demonstration. In those tasks in which the robot is required to fetch items for subsequent manipulation tasks, ambiguities might originate from the presence of identical objects in the scene. To disambiguate the scene, we propose a method that exploits visual data and uses verbal-HRI to request the human intervention, asking questions to understand the target item for the task. We combine this method to the existing BT learning framework to endow the robot with the capability of solving the task in ambiguous scenarios.

AUTONOMOUS SYSTEMS (AS) | Page 110 B
Iovino, Matteo
KTH / ABB Corporate Research
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Combining verbal Human Robot Interaction to solve ambiguities in Behavior Tree execution

Matteo Iovino, Irmak Doğan, Christian Smith, Iolanda Leite, KTH
Robotics Perception and Learning

Motivation & Research goals

Fast changing tasks in unpredictable, collaborative environments are typical for medium-small companies, where robotised applications are increasing. Thus, robot programs should be generated in short time with small effort, and the robot able to react dynamically to the environment. To address this, a method exists that combines context awareness and planning to learn Behavior Trees (BTs) from demonstration [1]. However, situations may arise where the robot is tasked to fetch an item that is present in multiple copies. The robot faces an ambiguous scenario that has to be disambiguated for the task to continue. We propose to combine the existing LfD method for BTs with verbal-HRI that uses visual data to query the scene for the target object and asks questions to the human to disambiguate it [2].

Learn BTs from demonstration

At a high level, our proposed algorithm learns BTs from demonstrations in four steps. Human demonstrations are clustered to infer the context of each action and similarities between them, and then to infer task constraints and goal conditions, which are finally fed to a planner that builds the BT.

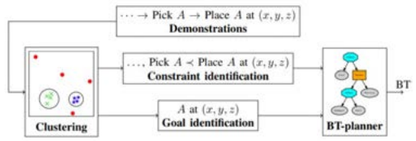

Demonstrations
The teaching method is kinesthetic and there are three actions available: a Pick action will close the robot grippers around the target object and a Place or Drop action will open the grippers, releasing the object. For all actions, the pose of the end-effector is recorded as the target pose for that action.
Behavior Tree Synthesis
The BT is synthesised using the planner proposed in [3], leveraging the idea of backchaining. We run the plan offline because it is preferable to have the full tree available before running it on a real robot.
Goal and Constraints identification
The algorithm infer task constraints by observing the order in which actions appear in the demonstrations and adding each pair of ordered actions to the list of constraint and translated into preconditions that must be fulfilled before executing an action. Conflicting constraints are removed.
Clustering of demonstrated actions
Different actions might also be executed in different reference frames. Thus, equivalent actions across demonstrations have to be identified and their reference frame inferred. If an action possibly belongs to multiple clusters, we can infer the context in which the action is performed.

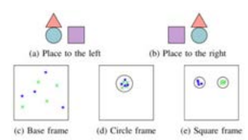

Resolve ambiguities with verbal-HRI

When a human queries the robot to identify an object, situations may arise where the target object is present in multiple copies in the scene. Using RGB data from the camera, the robot uses Grad-CAM [3] to compute the activation regions corresponding to the query and then K-means to cluster them. Then, the DETR Object Detector [4] is used to detect the objects in the scene together with their bounding boxes. The bounding boxes are compared against the clusters to output candidate scene regions. Deep Neural Network techniques are used to parse the natural language sentence to find the target object and to formulate clarification conditions, using other identified items in the scene and referring spatial expression *(left of* **item_x***)*. A conversation is then started to finally disambiguate the object.

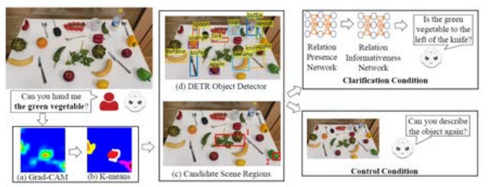

Combined method

We propose to combine the two methods to disambiguate the task during the execution of a BT learned from demonstration. We assume that the scene is not ambiguous during the demonstration and hence the robot is able to successfully grasp the target object for the task. If the task is ambiguous at execution the BT will fail and the disambiguation pipeline is triggered.

References

1. Gustavsson et al. (2021). Combining Context Awareness and Planning to Learn Behavior Trees from Demonstration. *arXiv e-prints*, arXiv-2109.
2. Dogan et al. (2021) Asking follow-up clarifications to resolve ambiguities in human-robot conversation. *Preprint*
3. Selvaraju et al. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision* (pp. 618-626).
4. Kamath et al. (2021). Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 1780-1790).
5. Colledanchise et al. (2019, May). Towards blended reactive planning and acting using behavior trees. In *2019 International Conference on Robotics and Automation (ICRA)* (pp. 8839-8845). IEEE.

Jakobsson, Erik
Linköping University / Epiroc Rock Drills AB

Condition Monitoring for Hydraulic Rockdrills

In this work we investigate data driven methods for classifying patterns in pressure measurements from hydraulic rock drills. By using signatures from induced faults, we aim to handle different configurations and individuals without the need for obtaining training data from all possible configurations. The key is to generate features that capture the difference from a fault in relation to a non-faulty reference. These relative features should also be insensitive to differences from different configurations.

Jakobsson, Erik
Linköping University / Epiroc Rock Drills AB

Condition Monitoring for Hydraulic Rockdrills

Erik Jakobsson
Erik Frisk, Robert Pettersson, Mattias Krysander
erik.jakobsson@epiroc.com

Description

In this project we aim to monitor the internal condition of hydraulic rock drills. This is done using measurements during operation, combined with machine learning schemes to classify different faulty behavior. An important aspect is the high variability between different applications/configurations. The use of Non-Fault reference data is a key technique.

Background & Motivation

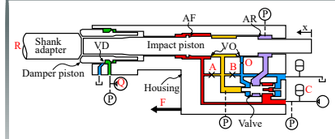

The basic functionality of a rock drill is simple. An impact piston is connected via hydraulic channels to a valve. The interaction between the two components results in a self oscillating mechanism, where impacts between the impact piston and shank adapter generate stress waves used to drill holes in hard rock. Side effects of the oscillations are severe vibrations and pressure pulsations, making the rock drill a difficult application to monitor. The hope is to use a very low number of sensors, possible even positioned away from the machine.

Knowing the current condition of the rock drill is an important step towards autonomous mining, where the information will be used for maintenance planning, logistics and prevention of secondary damage.

Contributions

Pressure Signature Fault Detection A single pressure sensor is used to classify the current condition of a rock drill in a lab setting. An example of such pressure data is shown below, for a single cycle.

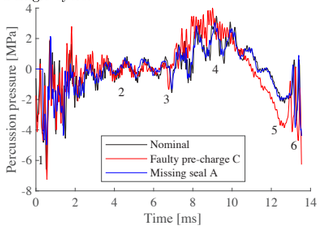

The faulty conditions are seen to slightly differ from the nominal No-fault case. However, *configuration* changes such as different hose lengths can have similar effects and need to be accounted for. We do this by looking at the difference from the No-Fault case, from the same configuration.

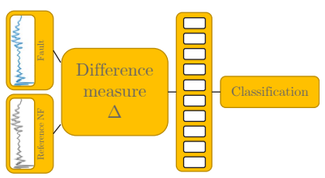

We wish to find a function Δ such that the same fault A in different configurations i, j give the same output,

$$\Delta(NF_i, A_i) = \Delta(NF_j, A_j) \quad \forall i \neq j \tag{1}$$

but different faults A, B give different outputs

$$\Delta(NF_i, A_i) \neq \Delta(NF_i, B_i) \quad \forall A \neq B \tag{2}$$

Dynamic Time Warping Feature Vector So far, the best Δ found consists of a number of difference measures based on Dynamic Time Warping (DTW), as a way to handle differences in stroke duration and pressure wave propagation.

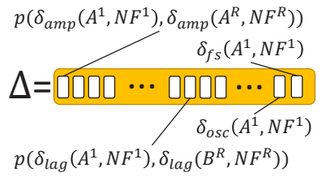

A feature vector is generated for each sample by measuring pairwise difference $p(\delta_1, \delta_{ref})$ and frequency differences to a set of reference samples using various measures δ. SVM classification using such feature vectors give the following accuracy:

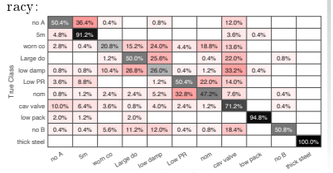

Research Goal & Question

A number of research questions define the area:

1. How can a rock drill be modeled and monitored in order to predict future failure? What data should be collected to maximize the information gathered without creating a too complex product?

2. Can a low number of non-dedicated sensors be used to monitor multiple components, for example different parts of a drilling system?

3. How can No-Fault data be used to give a reference to handle differences between configurations?

An important aspect of the research is to understand how condition monitoring methods can be applied for products with a relatively low volume, high customization, and in a very harsh environment.

Roadmap & Milestones

- Three conference papers accepted [1],[3],[5], two journal papers [2],[4] accepted.
- Licentitate thesis presentation, December 2019.
- Ongoing (Final): Condition monitoring of hydraulic rock drills journal paper.

Publications

[1] Jakobsson et al. "Data driven modeling and estimation of accumulated damage in mining vehicles using on-board sensors" published in Proceedings of *Annual Conference of the Prognostics and Health Management Society, St. Petersburg, Florida, USA, 2017.*

[2] Jakobsson et al. "Fatigue Damage Monitoring and Prognostics for Mining Vehicles using Data Driven Models" published in the *International Journal of Prognostics and Health Management (IJPHM), 2019.*

[3] Jakobsson et al. "Automated Usage Characterization of Mining Vehicles For Life Time Prediction" published in Proccedings of *IFAC World Congress, Berlin, 2020.*

[4] Åstrand et al. "A System for Underground Road Condition Monitoring" published in *International Journal of Mining Science and Technology, 2019.*

[5] Jakobsson et al. "Fault Identification in Hydraulic Rock Drills from Indirect Measurement During Operation" published in Proccedings of *IFAC MMM, Nancy, 2021.*

Contextual Deliberation

Determining the context can help AI systems and agents in interacting with or resembling humans. However proper context dependent reasoning systems do not exist yet. Our aim is to create a framework that can make human like context dependent decisions. The poster gives a simple example that shows how we as humans intuitivily use context determination to make decisions. Then it follows up with an initial conceptual framework that is a start for contextual deliberation in machines.

AUTONOMOUS SYSTEMS (AS) | Page 112 B
Jensen, Maarten
Umeå University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

WHAT IS YOUR CONTEXT?

Contextual Deliberation

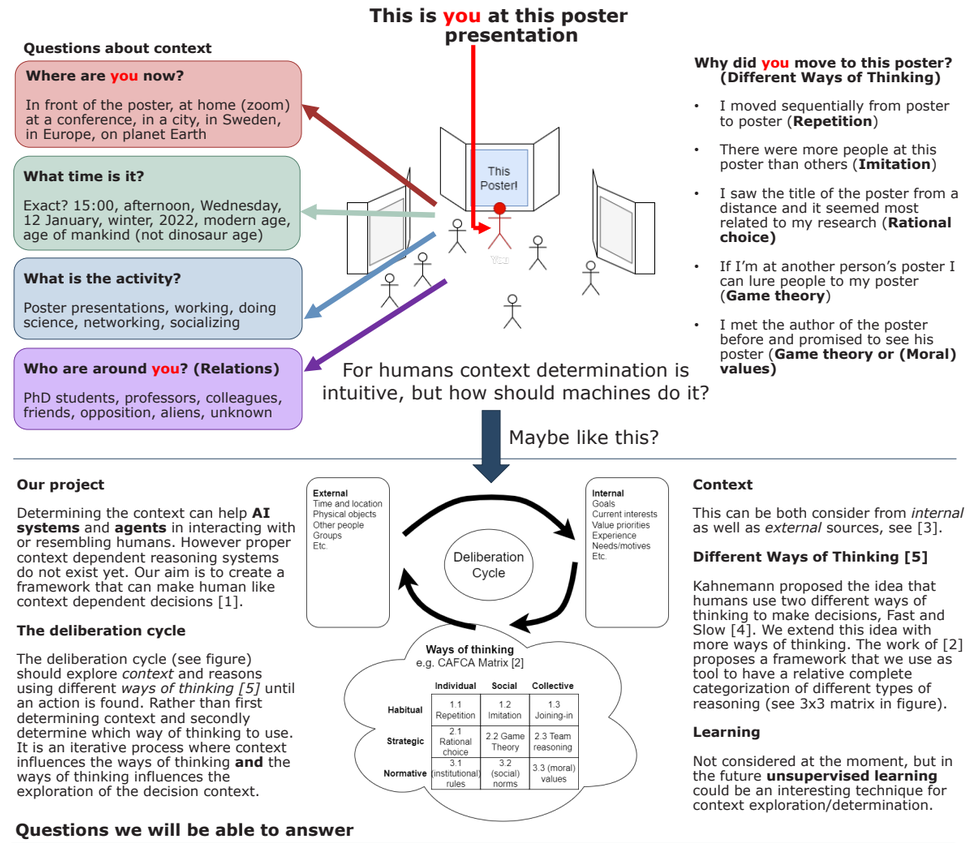

Our project

Determining the context can help **AI systems** and **agents** in interacting with or resembling humans. However proper context dependent reasoning systems do not exist yet. Our aim is to create a framework that can make human like context dependent decisions [1].

The deliberation cycle

The deliberation cycle (see figure) should explore *context* and reasons using different *ways of thinking [5]* until an action is found. Rather than first determining context and secondly determine which way of thinking to use. It is an iterative process where context influences the ways of thinking **and** the ways of thinking influences the exploration of the decision context.

Context

This can be both consider from *internal* as well as *external* sources, see [3].

Different Ways of Thinking [5]

Kahnemann proposed the idea that humans use two different ways of thinking to make decisions, Fast and Slow [4]. We extend this idea with more ways of thinking. The work of [2] proposes a framework that we use as tool to have a relative complete categorization of different types of reasoning (see 3x3 matrix in figure).

Learning

Not considered at the moment, but in the future **unsupervised learning** could be an interesting technique for context exploration/determination.

Questions we will be able to answer

Insight in contextual deliberation can help us with questions like:

1. How do people determine the context?
2. How do people react to a new law?
3. How do people react in a lockdown?
4. How do we program a realistic human-interaction system or robot?

1. [To be published] Jensen, M. Verhagen H., Vanhée L., & Dignum F. (2021) Towards Efficient Context-Sensitive Deliberation
2. Elsenbroich, C., Verhagen, H.: The simplicity of complex agents: a contextual action framework for computational agents. Mind & Society15(1), 131–143 (2016)
3. Zimmermann, A., Lorenz, A., & Oppermann, R. (2007, August). An operational definition of context. Springer, Berlin, Heidelberg.
4. Daniel, K. (2017). Thinking, fast and slow?
5. Minsky, M. (2007). The emotion machine, Simon and Schuster.

WASP Cluster: Agents

UMEÅ UNIVERSITY

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Mit-huset, 901 87 Umeå, ww.umu.se

Author: Maarten Jensen (maartenj@cs.umu.se), Supervisor: Frank Dignum, Loïs Vanhée, Collaboration: Harko Verhagen

Jonnarth, Arvi
Linköping University / Husqvarna

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning to Segment Images Without Mask Labels

Deep learning methods have achieved remarkable results in many computer vision tasks, including semantic segmentation, where the task is to classify each pixel in an image to a predefined set of classes, e.g. person, cat or car. Applications include autonomous driving, video surveillance, and medical image analysis. However, training deep segmentation models requires large datasets of costly human-annotated pixel-wise segmentation masks. In this work, we explore a branch called weakly-supervised semantic segmentation, where the only source of supervision are cheap image-level classification labels. We propose two contributions; importance sampling, and feature similarity loss, for approaching this challenging task, and significantly improve contour accuracy over state-of-the-art methods.

AUTONOMOUS SYSTEMS (AS) | Page 113 B

Jonnarth, Arvi
Linköping University / Husqvarna

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning to Segment Images Without Mask Labels

Arvi Jonnarth, Linköping University & Husqvarna
Department of Electrical Engineering, Computer Vision Laboratory
Supervisors: Michael Felsberg (LiU), Adam Tengblad (Husqvarna)

Abstract

Deep learning methods have achieved remarkable results in many computer vision tasks, including semantic segmentation, where the task is to classify each pixel in an image to a predefined set of classes, e.g. person, cat or car. Applications include autonomous driving, video surveillance, and medical image analysis. However, training deep segmentation models requires large datasets of costly human-annotated pixel-wise segmentation masks. In this work, we explore a branch called weakly-supervised semantic segmentation, where the only source of supervision are cheap image-level classification labels. We propose two contributions; *importance sampling*, and *feature similarity loss*, for approaching this challenging task, and significantly improve contour accuracy over state-of-the-art methods.

Methods

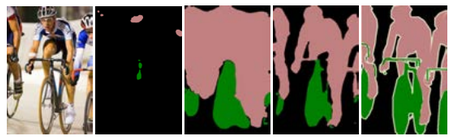

(a) (b) (c) (d) (e)

Figure 1. CAM comparison. (a) Input image; pseudo-masks with (b) max pooling, (c) importance sampling, and (d) importance sampling and feature similarity loss; (e) ground truth.

A fully convolutional neural network is trained in three stages:

1. Training of a multi-label classification network to generate class activation maps (CAMs). Max or average pooling is used to go from pixel-wise to image-level predictions.
2. Training of an AffinityNet [2] to predict pixel affinities.
3. A final segmentation network is supervised by pseudo-masks generated by the CAM and AffinityNet networks.

Contributions

Classification networks are known to (1) mainly focus on discriminative regions, and (2) to produce diffuse CAMs without well-defined prediction contours. We approach both problems with two contributions for improving CAM learning in stage 1.

First, we use **importance sampling** based on K probability mass functions p_k, one per class $k \in \{1, \dots, K\}$, induced by the CAMs $a_\theta \in [0,1]^{W \times H \times K}$ to sample image-level predictions $\tilde{y}_k$.

$$p_k(I, J|x) = \Pr(I = i, J = j|x, k) = \frac{1}{Z_k(a)} a_\theta(x)_{ijk},$$

$$\tilde{y}_k = a_\theta(x)_{ijk}, \quad (i, j) \sim p_k(I, J|x).$$

Second, we formulate a **feature similarity loss** term $\mathcal{L}_{\text{fs}}$ which aims to match the prediction contours with edges in the image.

$$\mathcal{L}_{\text{fs}}(a, x) = -\frac{1}{(HW)^2} \sum_{ij} w_{ij}\, g(a_i, a_j)\, f(\delta(x_i, x_j)),$$

pixel dissimilarity function

Gaussian spatial weight $\frac{1}{2\pi\sigma^2} \exp\left(-\frac{\|p_i - p_j\|_2^2}{2\sigma^2}\right)$

gating function $\frac{1}{2}\|a_i - a_j\|_2^2$

$\tanh\left(\mu + \log\left(\frac{\delta}{1-\delta}\right)\right)$, $\delta = \|x_i - x_j\|_1 / C$

Selected Results

The model is evaluated on the VOC benchmark dataset, with 20 foreground classes. Qualitative results are shown in Figure 2.

In Table 1 we compare our method with state-of-the-art weakly supervised methods in terms of two complementary metrics:

1. Mean intersection over union (mIoU) based on the area of predicted segmentation masks.
2. F-score based on the contours of segmentation predictions.

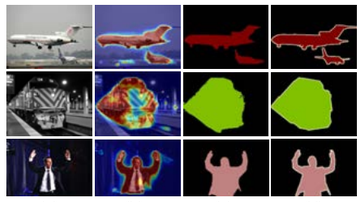

(a) (b) (c) (d)

Figure 2. Qualitative results. (a) Input image, (b) foreground class activations of the CAM network, (c) segmentation predictions of the final model from stage 3, and (d) ground truth.

Table 1. Performance comparison on the VOC validation set.

Method	Area mIoU	Contour F-score
SEAM [1]	64.5	35.7
PMM [3]	68.5	42.1
Ours	66.1	48.6

References

1. Wang et al., *Self-Supervised Equivariant Attention Mechanism for Weakly Supervised Semantic Segmentation*, CVPR, 2020.
2. Ahn et al., *Learning Pixel-Level Semantic Affinity with Image-Level Supervision for Weakly Supervised Semantic Segmentation*, CVPR, 2018.
3. Li et al., *Pseudo-Mask Matters in Weakly-Supervised Semantic Segmentation*, ICCV, 2021.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Kaalen, Stefan
KTH / Scania

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

SMP-tool for quantitative analysis of systems

Systems are growing more and more complex, which makes research and development increasingly relying on model-based development in order to ensure the safety of cyber-physical systems. Stateflow is a tool that supports modeling of systems as finite-state machines and has become the industrial standard practice in among others the automotive industry. However, Stateflow is limited in that it does not explicitly support modeling of stochastic processes, which are essential in model-based safety analysis. In order to overcome this, I have together with my colleagues developed SMP-tool that allow for modeling systems as Stochastic StateFlow (SSF) Models, and analysis of these models by studying the underlying stochastic process on the form of a generalized semi-Markov process.

AUTONOMOUS SYSTEMS (AS) | Page 114 B

Kaalen, Stefan
KTH / Scania

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

SMP-tool for quantitative analysis of systems

Stefan Kaalen, KTH
Mechatronics

KTH VETENSKAP OCH KONST

SCANIA

Abstract

Systems are growing more and more complex, which makes research and development increasingly relying on model-based development in order to ensure the safety of cyber-physical systems. Stateflow is a tool that supports modeling of systems as finite-state machines and has become the industrial standard practice in among others the automotive industry. However, Stateflow is limited in that it does not explicitly support modeling of stochastic processes, which are essential in model-based safety analysis. In order to overcome this, I have together with my colleagues developed SMP-tool that allow for modeling systems as Stochastic StateFlow (SSF) Models, and analysis of these models by studying the underlying stochastic process on the form of a generalized semi-Markov process.

SSF models

SSF models is a stochastic extension of a subset of **Stateflow**. The subset has been chosen according to utility found through numerous case studies joint with a desire to produce a **safe subset** [1]. SSF models can be modeled, although not analyzed directly, in Stateflow. The figure below presents an SSF model in Stateflow of a case study of a subsystem of a gearbox.

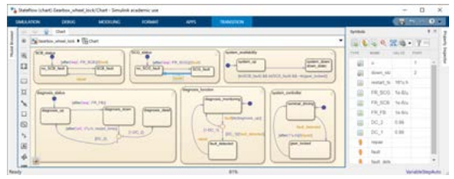

The state marked “down_state” representations the system failure of the wheels of the vehicle locking at their current position caused by an erroneous actuation of the gears in the gearbox. The case study is further explained in [2]

SSF models extends Stateflow both with the option of assigning probability distributions to the waiting time of transitions and with the option to assign discrete probabilistic choices of the destination state of transitions. For full syntax and semantics of SSF models, see [2]

SMP-tool

SMP tool has the ability to perform multiple types of analysis of SSF models modeled in Stateflow. The analyses include transient analyses of the reliability, parameter sensitivity analysis, and steady state analysis. The tool is free and can be downloaded from [3].

The tool has a simulation engine for SSF models and a symbolic/numerical engine for SSF models where the underlying stochastic process is a Hierarchical Semi-Markov Processes (HSMP) [4].

The result from a transient analysis of the case study presented earlier is presented in the following figure.

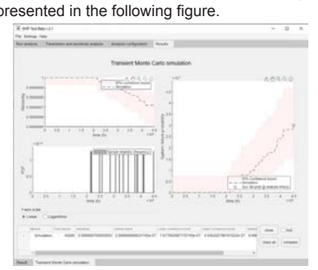

The figure presents how the probability of a system failure develops over time given the parameters specified in the model. By a sensitivity analysis, different parameter configurations, and thereby system specifications, can be found which delivers the same reliability.

References

1. MathWorks Advisory Board. 2020. Control algorithm modeling guidelines using MATLAB, Simulink, and Stateflow. Technical Report. SRI International. https://www.mathworks.com/solutions/mab-guidelines.html. Cited 2021-12-22
2. S. Kaalen, A. Hampus, M. Nyberg, and O. Mattsson. 2022. A stochastic extension of Stateflow. In IPCE ’22 Proceedings of the 13th ACM/SPEC international conference on performance engineering, 2022, ACM (unpublished).
3. kth.se/itm/smp-tool. Cited 2021-12-22
4. Stefan Kaalen, Mattias Nyberg, and Olle Mattsson. 2021. Transient Analysis of Hierarchical Semi-Markov Process Models with Tool Support in Stateflow. In Quantitative Evaluation of Systems: 18th International Conference, Proceedings, Springer Nature , 2021, p. 105-126. Springer nature, 105—126

SMP-tool

We have delivered a stochastic extension of Stateflow for the purpose of evaluation the performance of safety critical systems.

Future work includes extending the symbolic/numerical engine to handle all SSF models. Furthermore, to make SMP-tool tractable for use in the industry, future works includes support for modeling complex systems as SSF models.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Kampik, Timotheus
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Explainable Reasoning and Decision-Making: From Humans to Machines

The line of work that leads to my dissertation studies automated reasoning and its intersection with human decision-making.
Most of the works address one of the following two questions: i) How can principles of human reasoning and decision-making be applied to drawing explainable inferences from knowledge bases with conflicting statements? ii) How can we improve agility and human explainabiltiy of complex 'intelligent' software systems?
The research subject of Question i) is formal argumentation, a graph-based method for non-monotonic reasoning, and the primary method is formal analysis. The main research results are different formal methods to ensure consistency when drawing repeated inferences from changing argumentation graphs, and analyses of the ability of different inference functions to support these approaches; in particular, a novel bridge between formal argumentation as a form of non-monotonic reasoning and economically rational decision-making is built. Regarding Question ii), the research subject is (the engineering of) multi-agent systems, which is studied from engineering and human-computer interaction points of view. The main results are new perspectives on and approaches to deploying agents in multi-agent systems in dynamic, Web-based environments and empirical results on how multi-agent systems can be better explained to human users.

Kampik, Timotheus
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Explainable Reasoning and Decision-Making: From Humans to Machines

Timotheus Kampik, Umeå University
Department of Computing Science
Supervisors: Helena Lindgren and Juan Carlos Nieves

UMEÅ UNIVERSITY

In a Nutshell

The line of work that leads to my dissertation studies automated reasoning and its intersection with human decision-making. Most of the works address one of the following two questions: *i)* How can principles of human reasoning and decision-making be applied to drawing *explainable* inferences from knowledge bases with conflicting statements? *ii)* How can we improve agility and human explainabiltiy of complex "intelligent" software systems? The research subject of Question *i)* is formal argumentation, a graph-based method for non-monotonic reasoning, and the primary method is formal analysis. The main research results are different formal methods to ensure consistency when drawing repeated inferences from changing argumentation graphs, and analyses of the ability of different inference functions to support these approaches; in particular, a novel bridge between formal argumentation as a form of non-monotonic reasoning and economically rational decision-making is built. Regarding Question *ii)*, the research subject is (the engineering of) *multi-agent systems*, which is studied from engineering and human-computer interaction points of view. The main results are new perspectives on and approaches to deploying agents in multi-agent systems in dynamic, Web-based environments and empirical results on how multi-agent systems can be better explained to human users.

Example: 'Formal' Part

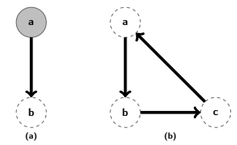

Figure 1: Inconsistent preferences: using many *abstract argumentation* reasoning methods, the left graph implies $\{a\}$ is preferred over $\{\}$, while the right graph implies $\{\}$ is preferred over $\{a\}$.

This semi-formal example illustrates how most inference functions of abstract argumentation, in which conflicts in a set of *arguments* (for example: logical statements, business rules, claims in legal proceedings, *et cetera*) are modeled as a directed graph, violate the *consistent preferences* principle of economic rationality: given any set of choice items A, a rational agent consistently chooses the same items $A^* \subseteq A$, which implies that $\forall A_c \subseteq A$, such that $A_c \neq A^*$, A^* is preferred over A_c. Given a set of options $A' \supseteq A$, the agent must choose A'^* so that $A'^* = A^*$ or $A'^* \not\subseteq A$. This model is too simple to guide real-life decision-making (as has been shown by a range of behavioral economics research). However, it can be used as a sanity check for decision and reasoning algorithms. *E.g.*, the figure above shows that almost all of the well-established inference functions of abstract argumentation are not compliant with properties of economic rationality.

Example: 'Engineering' Part

Figure 2: A multi-agent simulation of a drone delivery scenario, with an explanation message for a human supervisor in the top right corner.

The figure above shows a drone delivery simulation. Each drone can be thought of as an *autonomous agents* that has a partial view of the world; its knowledge may be incomplete and inconsistent with the knowledge that other drones or the "global" operators have. Hence, unexpected situations may occur, such as two drones attempting to pick up the same package, which in turn may result in a drone having to change directions mid-way. To make the overall behavior of the agents *explainable* to human operators, the state of all agents needs to be aggregated and filtered, and it is not exactly clear how to make the right trade-off to achieve explanation granularity that is useful, yet concise. In a human-interaction case study, we have compared different approaches to filtering explanations.

Selected Publications

[I] Ensuring reference independence and cautious monotony in abstract argumentation. Kampik, Nieves & Gabbay. International Journal of Approximate Reasoning. 2022

[II] The quest of parsimonious XAI: A human-agent architecture for explanation formulation. Mualla, Tchappi, Kampik, Najjar, Calvaresi, Abbas-Turki, Galland & Nicolle. Artificial Intelligence. 2022

[III] Governance of autonomous agents on the Web: Challenges and opportunities. Kampik, Mansour, Boissier, Kirrane, Padget, Payne, Singh, Tamma & Zimmerman. ACM Transactions of Internet Technologies. 2022

[IV] Abstract Argumentation and the Rational Man. Kampik & Nieves. Journal of Logic and Computation. 2021

[V] Argumentation-based health information systems: A design methodology. Lindgren, Kampik, Guerrero Rosero, Blusi & Nieves. IEEE Intelligent Systems. 2021

[VI] The burden of persuasion in abstract argumentation. Kampik, Gabbay & Sartor. International Conference on Logic and Argumentation. 2021

[VII] Explanations of non-monotonic inference in admissibility-based abstract argumentation. Kampik & Čyras. International Conference on Logic and Argumentation. 2021

[VIII] The degrees of monotony-dilemma in abstract argumentation. Kampik & Gabbay. European Conference on Symbolic and Quantitative Approaches with Uncertainty. 2021

[IX] Autonomous agents on the edge of things (demonstration). Kampik, Gomez, Ciortea & Mayer. Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems. 2021

[X] A Framework for collaborative and interactive agent-oriented developer operations (demonstration). Amaral, Kampik & Cranefield. Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems. 2020

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Khosravi, Hedieh
Lund university

MM-wave channel sounding for indoor positioning

In this project we investigate the millimeter-wave (mm-wave) channel capability of being utilized for highly accurate indoor positioning purposes such as smart factory, sensing ,etc. Highly accurate radio based positioning relies on the additional information provided by Multi-Path Components (MPCs) which act as Virtual Anchor (VA) points besides the Line of Sight (LoS) and physical anchor. Hence in the first stage of the project we have focused on characterizing the behavior of the MPCs over time, e.g. the number of tractable ones and their life time, in mm-wave channel, by analyzing the real scenario measured data with high resolution in frequency and spatial domains.

AUTONOMOUS SYSTEMS (AS)

Khosravi, Hedieh
Lund university

Kullberg, Anton
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

On Joint State Estimation and Model Learning using Gaussian Process Approximations

State estimation is of interest in essentially every sector of science and engineering. Typically, techniques for state estimation require the specification of a dynamical model of the system in question. It is often possible to derive a partial description of the system dynamics, but depending on the modeling assumptions, this can potentially lead to bad state estimates, due to an insufficient description of the dynamics. This project explores the combination of such a partial dynamical description with a generic black-box structure to allow online learning of parts of the system dynamics. In this way, the model can be improved over time, as more measurements have been obtained, and in extension improve the resulting state estimate. We provide some initial results in this regard.

AUTONOMOUS SYSTEMS (AS) | Page 117 B

Kullberg, Anton
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

On Joint State Estimation and Model Learning using Gaussian Process Approximations

Anton Kullberg, PhD Student, Linköping University
Div. of Automatic Control
Supervisors: Assoc. Prof. Gustaf Hendeby (LiU) and Assoc. Prof. Isaac Skog (LiU)

LINKÖPINGS UNIVERSITET

Motivation & Research Goals

State estimation is of interest in essentially every sector of science and engineering. Typically, techniques for state estimation require the specification of a dynamical model of the system in question. It is often possible to derive a partial description of the system dynamics, but depending on the modeling assumptions, this can potentially lead to bad state estimates, due to an insufficient description of the dynamics. This project explores the combination of such a partial dynamical description with a generic black-box structure to allow online learning of parts of the system dynamics. In this way, the model can be improved over time, as more measurements have been obtained, and in extension improve the resulting state estimate. We provide some initial results in this regard.

Methods

We consider the general discrete-time description of a dynamical system given by

$$\begin{aligned}\mathbf{x}_{k+1} &= \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{g}_f(\mathbf{w}_k + \mathbf{g}(\mathbf{x}_k, \mathbf{u}_k)) \\ \mathbf{y}_k &= \mathbf{h}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{e}_k.\end{aligned}$$

Here, $\mathbf{x}_k$, $\mathbf{u}_k$, $\mathbf{y}_k$ are the state, input and measurement at time k, respectively. Further, $\mathbf{w}_k$, $\mathbf{e}_k$ are mutually independent white noise processes, particularly, $\mathbf{w}_k \sim \mathcal{N}(0, \mathbf{Q})$ and $\mathbf{e}_k \sim \mathcal{N}(0, \mathbf{R})$. The measurement function, $\mathbf{h}$, is assumed known and **parts** of the system dynamics, $\mathbf{f}$ and $\mathbf{g}_f$, are assumed known. Lastly, the function $\mathbf{g}$ captures any system dynamics that are not described by $\mathbf{f}$.
We model the function $\mathbf{g}$ as a basis function expansion

$$\mathbf{g}(\mathbf{x}_k) = \sum_j \phi^j(\mathbf{x}_k)\theta^j,$$

with ϕ^j chosen as radial basis functions, i.e., ϕ^j is a function of $\|\mathbf{x}_k - \xi^j\|$, where ξ^j is the basis function center. The basis functions are placed in the region of the state space in which there is unknown dynamics. Essentially, this constitutes a grid of basis functions, where the extent of the grid determines in what regions unknown dynamics can be learned and the density of the grid determines the fidelity of the learned dynamics. As such, if the state space region of interest is large, the number of parameters θ is large. To facilitate learning of the parameters θ, an augmented state vector is constructed as $\mathbf{x}_k^e = \begin{bmatrix}\mathbf{x}_k^\top & \boldsymbol{\theta}^\top\end{bmatrix}^\top$.
This enables us to estimate the state trajectories and learn (parts of) the model online in a joint fashion using an EKF [1]. Even though this is theoretically computationally efficient, as the number of parameters grows beyond a few thousand, this is not feasible to do in real-time, limiting the model to either a small state space region or low dynamical fidelity [2].
To resolve this issue, we choose ϕ^j such that

$$\phi^j(\|\mathbf{x}_k - \xi^j\|) \equiv 0, \ \|\mathbf{x}_k - \xi^j\| > c^j.$$

As such, each basis function ϕ^j only contributes to the function value in a region close to its center ξ^j, limiting the number of parameters necessary for each function evaluation. With a few modifications to the EKF recursions, this enables real-time online joint state estimation and model learning [2].

Selected Results

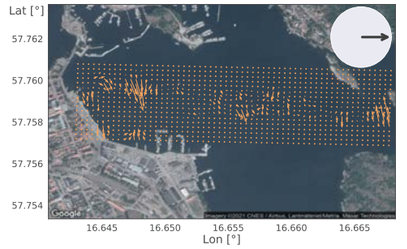

[3] Learned motion dynamics of vessels traveling in a harbor region in Västervik, Sweden. Visualized is the Cartesian acceleration of the vessels traveling out of port. The acceleration was learned from historical data of vessels traveling throughout the port, provided by the Swedish Defense Research Agency (FOI).

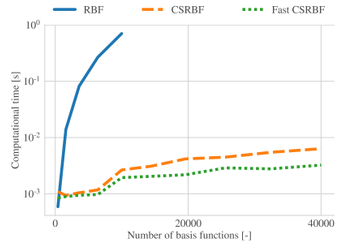

[2] Comparison of the computational time for a single propagation step in the EKF. Blue is for the standard formulation using Gaussian radial basis functions (global support). Orange and green both use compact basis functions but differ slightly in how they identify which basis functions to use. After about 10000 basis functions, the global formulation is no longer runnable at all, whereas the compact formulation easily handles a large number of basis functions.

References

[1] Learning Driver Behaviors Using A Gaussian Process Augmented State-Space Model
A. Kullberg, I. Skog and G. Hendeby
International Conference on Information Fusion (FUSION), July 2020

[2] Online Joint State Inference and Learning of Partially Unknown State-Space Models
A. Kullberg, I. Skog and G. Hendeby
Transactions on Signal Processing 69 2021

[3] Learning Motion Patterns in AIS Data and Detecting Anomalous Vessel Behavior
A. Kullberg, I. Skog and G. Hendeby
International Conference on Information Fusion (FUSION), Nov 2021

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

On Rendezvous in Autonomous Cooperative Landings

We investigate the rendezvous problem for the autonomous cooperative landing of an unmanned aerial vehicle (UAV) on an unmanned surface vehicle (USV). The rendezvous problem is challenging due to several reasons, for example, sudden communication losses or strong disturbances acting on the agents can lead to disastrous consequences. Moreover, even the basic tasks to determine if the rendezvous is possible or not and what strategy to employ when the rendezvous location has to be updated can be complex. Our goal is to create a rendezvous algorithm with an online update of the rendezvous location such that convergence is guaranteed. The preliminary proposed algorithm requires the agents to exchange information only when necessary to maintain the convergence.

AUTONOMOUS SYSTEMS (AS) | Page 118 B
Lapandić, Dženan
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

On Rendezvous in Autonomous Cooperative Landings

Dženan Lapandić, KTH Royal Institute of Technology
Division of Decision and Control Systems
Supervisor: Prof. Bo Wahlberg

Motivation & Research Goals

We investigate the rendezvous problem for the autonomous cooperative landing of an unmanned aerial vehicle (UAV) on an unmanned surface vehicle (USV). The rendezvous problem is challenging due to several reasons [2], for example, sudden communication losses or strong disturbances acting on the agents can lead to disastrous consequences. Moreover, even the basic tasks to determine if the rendezvous is possible or not and what strategy to employ when the rendezvous location has to be updated can be complex. Our goal is to create a rendezvous algorithm with an online update of the rendezvous location such that the convergence is guaranteed. The preliminary proposed algorithm requires the agents to exchange information only when necessary to maintain the feasibility.

Methods

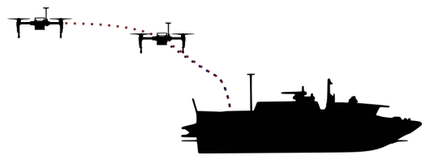

We consider two **heterogeneous** agents with **nonlinear dynamics** and additive disturbances. Each agent solves a corresponding distributed optimal control problem formulated as a **Model Predictive Control** problem penalizing the distance to the rendezvous location θ while satisfying state and input constraints.

The **control objective** is to steer the relevant states of every agent y_i to a rendezvous point $\theta \in \Theta \subseteq \mathbb{R}^p$ in finite time.

- It is assumed that the initial rendezvous location is feasible
- The time planning horizon T is long enough to reach at least one θ in the rendezvous set Θ.
- The agents update and share the rendezvous location only when they are not guaranteed to reach it, i.e. to maintain the feasibility

Based on the deviation of the predicted terminal state output from the rendezvous location $V_o = \|\hat{y}_i(t_k + T; t_k) - \theta(t_k)\|^2$ the agent i updates θ according to the online **update law**

$$\theta(t_{k+1}) = \begin{cases} \theta(t_k) & V_o \leq \varepsilon \\ \theta(t_k) - \eta v_\theta(t_k) & V_o > \varepsilon \end{cases}$$

where η and ε are tuning parameters and $v_\theta(t_k)$ is defined as:

$$v_\theta(t_k) = \frac{\partial V_o}{\partial \theta(t_k)} \left\| \frac{\partial V_o}{\partial \theta(t_k)} \right\|^{-1}.$$

Parameter η is a step size that must be chosen as a small value, in order to avoid overshooting, and it quantifies the correction of θ in the output space.

References

[1] Aperiodic Communication for MPC in Autonomous Cooperative Landing
Dženan Lapandić, Linnea Persson, Dimos V. Dimarogonas, Bo Wahlberg
7th IFAC Conference on NMPC 2021

[2] Model Predictive Control for Autonomous Ship Landing in a Search and Rescue Scenario.
Linnea Persson and Bo Wahlberg,
in AIAA Scitech 2019 Forum, 1169

Selected Results

Feasibility in distributed MPC scenario with a common rendezvous location is challenging to guarantee due to

- communication issues and delays,
- disturbances that may affect one or several agents to be unable to reach the previously agreed rendezvous location,
- update law that may propose a new location which is not feasible for the other agent.

Contributions [1]:

- Distributed rendezvous algorithm with **aperiodic communication** which eliminates unnecessary communication.
- **Time-varying distributed terminal sets** for tracking that depend on the rendezvous point.
- Proof that the proposed algorithm guarantees **recursive feasibility**.

Simulation example: Autonomous cooperative landing

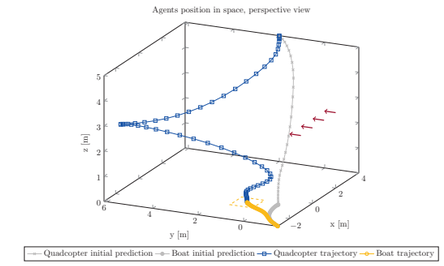

Arrows show wind direction and the yellow dashed polygon represents the boat landing platform.

Agents initiate the landing according to the initial rendezvous location. Due to the strong wind active for $t = [0.5\text{s}, 2.0\text{s}]$, the initial location becomes infeasible and has been updated using the update law to maintain the feasibility.

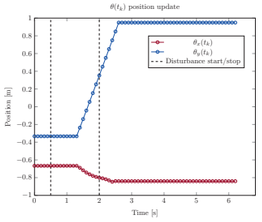

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Sensor Node Calibration in Presence of a Dominant Reflective Plane

In this paper we study the problem of sensor network self-calibration in presence of a single reflective plane. We propose a three-step stratified approach utilizing a rank-1 constraint in the measurements: (i) In the case of time difference of arrival (TDOA) measurements, any offsets in the measurements are solved for. (ii) The heights of the receivers and senders relative to the plane are solved for. (iii) The planar receiver and sender positions are solved for. We evaluate our approach on synthetic and real data.

Sensor Node Calibration in Presence of a Dominant Reflective Plane

Erik Tegler, Martin Larsson, Magnus Oskarsson, Kalle Åström
Centre for Mathematical Sciences, Lund University

Abstract

- In this paper we study the problem of sensor network self-calibration in presence of a single reflective plane.
- We propose a three-step stratified approach utilizing a rank-1 constraint in the measurements:
  1. In the case of time difference of arrival (TDOA) measurements, any offsets in the measurements are solved for.
  2. The heights of the receivers and senders relative to the plane are solved for.
  3. The planar receiver and sender positions are solved for.
- We evaluate our approach on synthetic and real data.

Problem Formulation

Consider the problem of time of arrival (TOA) self-calibration in the presence of a single reflective plane. Then every receiver $R_\wedge$ has a virtual mirror receiver $R_\vee$, and there are two distance measurements $D_\wedge$ and $D_\vee$ to the sender S given by

$$D_\wedge^2 = \|R_\wedge - S\|^2 = d^2 + (g-h)^2,$$
$$D_\vee^2 = \|R_\vee - S\|^2 = d^2 + (g+h)^2.$$

See figure to the right for notation. From these we can derive

$$D_\Delta = \frac{D_\vee^2 - D_\wedge^2}{4} = gh,$$
$$D_\Sigma = \frac{D_\vee^2 + D_\wedge^2}{2} = d^2 + g^2 + h^2.$$

Provided m receivers and n senders we get the rank-1 matrix

$$D_\Delta = \begin{pmatrix} g_1h_1 & \cdots & g_1h_n \\ \vdots & \ddots & \vdots \\ g_mh_1 & \cdots & g_mh_n \end{pmatrix} = \begin{pmatrix} g_1 \\ \vdots \\ g_m \end{pmatrix} \begin{pmatrix} h_1 & \cdots & h_n \end{pmatrix}$$

Offset Estimation

In the case of time difference of arrival (TDOA), additional offsets o_j in the measurements $Z_{\wedge ij}$ and $Z_{\vee ij}$ need to be estimated.

$$D_{\wedge ij} = Z_{\wedge ij} - o_j, \qquad D_{\vee ij} = Z_{\vee ij} - o_j,$$

where $i = 1, \dots, m$ and $j = 1, \dots, n$. D_Δ becomes linear in o_j. The offsets can be found linearly by utilizing the rank constraint on D_Δ.

Height Estimation

The heights g_i and h_j can be retrieved from a robust rank-1 approximation of D_Δ, up to some unknow constant λ, since $D_\Delta = (\lambda \boldsymbol{g})\left(\frac{1}{\lambda}\boldsymbol{h}^T\right)$.

Planar Position Estimation

The planar positions $\boldsymbol{r}_i$ and $\boldsymbol{s}_j$ are retrieved by solving a lower dimensional TOA problem, where the distances d_{ij} depend on λ.

$$d_{ij}^2 = \|\boldsymbol{r}_i - \boldsymbol{s}_j\|^2 = D_{\Sigma ij} - \lambda^2 g_{ij}^2 - h_{ij}^2/\lambda^2$$

We present two new solvers for the minimal case $(m, n) = (3,4)$.

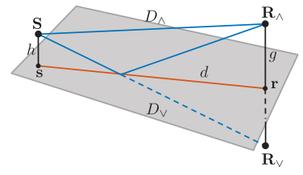

Experiments

Experiments on Synthetic Data

Numerical stability of the proposed solvers compared to the existing $(m, n) = (6,4)$ solver for noise-less data.

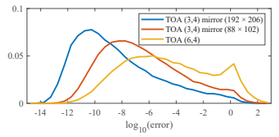

Experiments on Real Data

Estimated planar positions of the receivers and senders compared to ground truth.

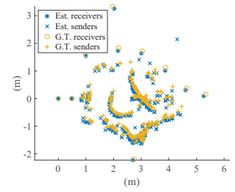

Estimated receiver and sender heights compared to ground truth.

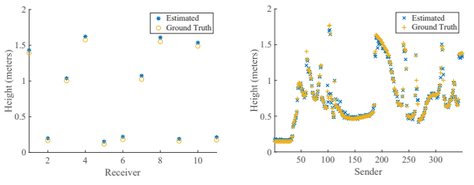

Human-Feedback Shield Synthesis for Perceived Safety in Deep Reinforcement Learning

Despite the successes of deep reinforcement learning (RL), it is still challenging to obtain safe policies. Formal verification approaches ensure safety at all times, but usually overly restrict the agent's behaviors, since they assume adversarial behavior of the environment.Instead of assuming adversarial behavior, we suggest to focus on perceived safety instead, i.e., policies that avoid undesired behaviors while having a desired level of conservativeness. To obtain policies that are perceived as safe, we propose a shield synthesis framework with two distinct loops: (1) an inner loop that trains policies with a set of actions that is constrained by shields whose conservativeness is parameterized, and (2) an outer loop that presents example rollouts of the policy to humans and collects their feedback to update the parameters of the shields in the inner loop.

AUTONOMOUS SYSTEMS (AS) | Page 120 B
Marta, Daniel
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Human-Feedback Shield Synthesis for Perceived Safety in Deep Reinforcement Learning

Daniel Marta, KTH Royal Institute of Technology
RPL: Robotics Perception and Learning

Abstract

Despite the successes of deep reinforcement learning (RL), it is still challenging to obtain safe policies. Formal verification approaches ensure safety at all times, but usually overly restrict the agent's behaviors, since they assume adversarial behavior of the environment.Instead of assuming adversarial behavior, we suggest to focus on perceived safety instead, i.e., policies that avoid undesired behaviors while having a desired level of conservativeness. To obtain policies that are perceived as safe, we propose a shield synthesis framework with two distinct loops: (1) an inner loop that trains policies with a set of actions that is constrained by shields whose conservativeness is parameterized, and (2) an outer loop that presents example rollouts of the policy to humans and collects their feedback to update the parameters of the shields in the inner loop.

Methods

Learning safety constraints from humans

- Inner-loop: takes advantage of self-play by sampling from shield distributions of human feedback [1]
- Outer-loop: updates a shield parameter distribution with human feedback [1]

Human-feedback Shield distribution:

- Computed iteratively from human-feedback datasets in the outer loop.
- Maps high-level human feedback into shield parameter updates.

$$\text{map}(g_j) = \begin{cases} \mu_h - \frac{|\mathcal{H}|\sigma}{2} & \text{if } g_j = \text{very unsafe}, \\ \vdots & \vdots \\ \mu_h & \text{if } g_j = \text{fine}, \\ \vdots & \vdots \\ \mu_h + \frac{|\mathcal{H}|\sigma}{2} & \text{if } g_j = \text{very safe}. \end{cases}$$

$$^u\mu_h = \frac{1}{N_{\text{user}}} \sum_{j=1}^{N_{\text{user}}} \text{map}(g_j),$$

$$^u\sigma_h^2 = \max\left(\frac{1}{N_{\text{user}}} \sum_{j=1}^{N_{\text{user}}} \left(\text{map}(g_j) - (^u\mu_h^2)\right), \sigma_{\min}^2\right)$$

$$\forall i : \text{KL}(^u f_{h_i}, f_{h_i}) \leq \beta$$

$$p(\mathcal{G}|\mu, \sigma^2) = \prod_{j=1}^{|\mathcal{G}|} p(x_j|\mu, \sigma^2), x_j \in \mathcal{G}$$

$$\hat{p}_\theta(\hat{\mu}_\theta|\mathcal{G}, \hat{\sigma}_\theta^2) \propto p(\mathcal{G}|\mu, \sigma^2) p_\theta(\mu|\mu_\theta, \sigma_\theta^2)$$

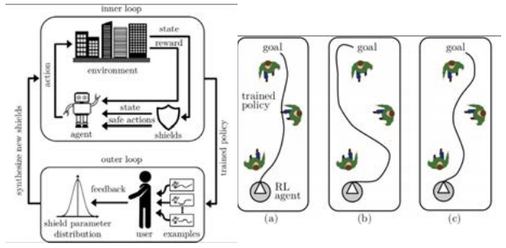

Results

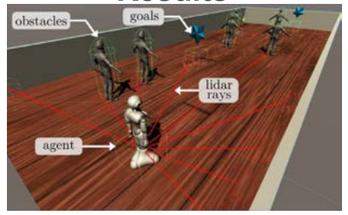

To evaluate perceived safety, we want to estimate how strong the force field should be, i.e., the shield parameter encodes how much the shield considers the full interaction force of the Social Force Model (SFM) We focus our shield synthesis from human-feedback, to address how our approach could have an impact in a robotic scenario. The state space is comprised by the agent's position and velocity, the velocity of other obstacles in the environment, and nine rays of a lidar-like sensor, commonly used in navigation robots. The agent's actions are composed of the accelerations in x- and y-directions, representing the driving force and a third action proportional to the interaction force of the SFM. For each ray, the agent detects either a goal, one of the humans or walls. The rays are one-hot encoded in addition to the distance between the robot and a specific element. In total, there are 9 rays opening in a field of view (FOV) of 200 degrees.

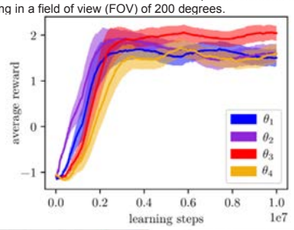

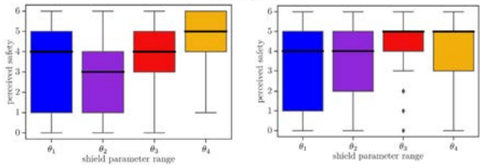

To access the validity of our approach, we ran a study with real humans. The study was run online using Amazon Mechanical Turk (AMT). In total, there were 92 unique participants (59 males, 33 female and none of other gender identities). Their age ranged from 23 to 65 years old, with a median of 34; the majority were in or had completed college education (N=78) and came from the US (N=71). 62 participants reported to have never or only seen robots in media, 14 to have interacted with one robot before, and 2 to do it on a regular basis.

References

1. D. Marta, C. Pek, G. I. Melsión, J. Tumova and I. Leite, "Human-Feedback Shield Synthesis for Perceived Safety in Deep Reinforcement Learning," in IEEE Robotics and Automation Letters, vol. 7, no. 1, pp. 406-413, Jan. 2022, doi: 10.1109/LRA.2021.3128237.

Learning Robot Tasks through Planning, Knowledge Integration and Multi-Objective Optimization

We introduce a framework for integrating planning with targeted learning of scenario-specific parameters. It uses a coarse-to-fine strategy: (1) the user provides a task goal in PDDL, (2) a plan (i.e., a sequence of skills) is generated and the learnable parameters of the skills are automatically identified. An operator then chooses (3) reward functions and hyperparameters for the (4) subsequent learning process. Learning is tightly integrated with a knowledge framework to support planning and to provide priors for learning and using multi-objective optimization, since objectives such as safety and task execution can often affect each other. Our system utilizes extended Behavior Trees for planning, execution. We adopt a multi-objective Bayesian optimization approach to learn the parameters of our tasks statistically efficient. Learning is done entirely in simulation and we use domain randomization techniques to ensure the results work in a real-world scenario.

AUTONOMOUS SYSTEMS (AS) | Page 121 B
Mayr, Matthias
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning Robot Tasks through Planning, Knowledge Integration and Multi-Objective Optimization

Matthias Mayr, Faseeh Ahmad, Konstantinos Chatzilygeroudis, Luigi Nardi and Volker Krueger

Lunds Universitet, Lunds Tekniska Högskola · University of Patras · Stanford University

How to safely learn new robot task with explainable policies?

Motivation

- **Fast adaption** to new tasks
- **Explainable policy representation**: know *what* is performend *when* and *why*
- **Safe learning process**
- **Automated pipeline**: Little user interaction from goal definition to learned policies

Skill-based systems

- Support planning of tasks
- **Reasoning** alone **often leads to complex systems**

How to combine planning and learning to learn the tacit aspects of robot task?

Approach

Skills with Behavior Trees[1]

- Behavior trees as **reactive** and **parametric** policy representation
- **Human-readable** and **editable** as well as **expandable** and **modular**
- Usable for **planning and skill execution**
- **Skills expose parameters** such as:
  - Objects to manipulate
  - Conditions
  - Offsets for motions

Planning and Knowledge Integration

- *SkiROS2* as a **platform** for **skills** and the **world model**
- *PDDL* to formulate the **planning problem**
- Automatically generate planning domain
- The **world model** provides and stores **knowledge** about the given problem
  - supplies **information** about the **learnable parameters**

Policy Optimization

Dynamical system in the form

$$\mathbf{x}_{t+1} = \mathbf{x}_t + M_{sim}\left(\mathbf{x}_t, \mathbf{u}_t, \phi_R\right)$$

with transition dynamics $M_{sim}\left(\mathbf{x}_t, \mathbf{u}_t, \phi_R\right)$ modeled by a simulator.

- **Policy search** with black-box optimization algorithm: *Bayesian optimization*
- **Reward functions** can be select from a **library** by the operator
- **Multiple objectives** can be optimized concurrently

Simulation-supported Learning

- **Less interaction time** with the robot
- **Safe** for **robot and production material**
- **Scales** with **cloud resources**
- Allows **object tracking** without a complex setup

Domain Randomization

- learn more **robust policies** that generalize to reality
- emulate different configurations

Examples for Learned Tasks

Push an Object

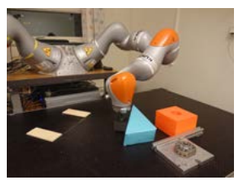

Peg Insertion

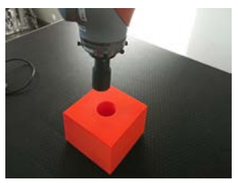

Learning Pipeline

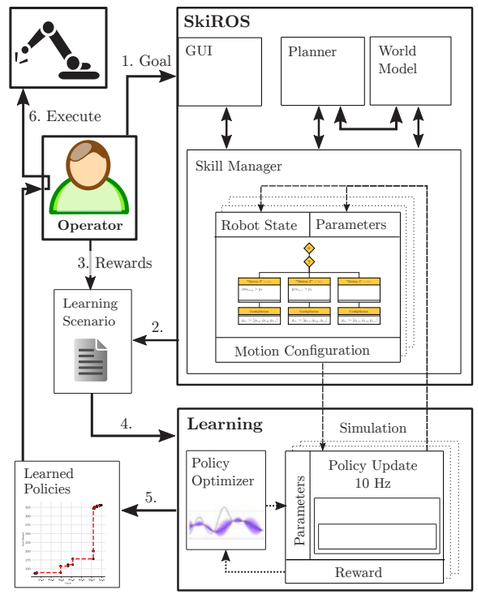

Steps of the Learning Pipeline:

1. Goal definition by the operator
2. Plan generation and parameterization; learnable parameters are identified
3. Operator complements scenario with hyperparameters, rewards and objectives
4. Learning in simulation
5. Pareto front with the best policies
6. Operator selects solutions and executes them on the real system

Future Work

1. **Learning** of the **task structure** including **recovery behaviors**
2. **Multi-fidelity optimization** that can include the real system
3. **Automatic reasoning** about rewards and objectives

References

(1) Rovida, F.; Grossmann, B.; Krueger, V. Extended Behavior Trees for Quick Definition of Flexible Robotic Tasks. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*; 2017; pp 6793–6800.
(2) Mayr, M.; Chatzilygeroudis, K.; Ahmad, F.; Nardi, L.; Krueger, V. Learning of Parameters in Behavior Trees for Movement Skills. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*; 2021.

Moliner, Olivier
Lund University / Sony

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Bootstrapped Representation Learning for Skeleton-Based Action Recognition

In this work, we study self-supervised representation learning for 3D skeleton-based action recognition. We extend Bootstrap Your Own Latent (BYOL) for representation learning on skeleton sequence data and propose a new data augmentation strategy including two asymmetric transformation pipelines. We also introduce a multi-viewpoint sampling method that leverages multiple viewing angles of the same action captured by different cameras. In the semi-supervised setting, we show that the performance can be further improved by knowledge distillation from wider networks, leveraging once more the unlabeled samples.
We conduct extensive experiments on the NTU-60 and NTU-120 datasets to demonstrate the performance of our proposed method. Our method consistently outperforms the current state of the art on both linear evaluation and semi-supervised benchmarks.

AUTONOMOUS SYSTEMS (AS) | Page 122 B

Moliner, Olivier
Lund University / Sony

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Bootstrapped Representation Learning for Skeleton-Based Action Recognition

Olivier Moliner [1,2] Sangxia Huang [2] Kalle Åström [1]

[1]Lund University [2]Sony R&D Center Lund

SONY

LUND UNIVERSITY

Introduction

We study self-supervised representation learning for 3D skeleton-based action recognition.

Motivation Fully-supervised action recognition algorithms require large datasets of 3D skeleton data with accurate annotations, which are time-consuming and costly to prepare.

Goal To learn semantic features from unlabeled 3D skeleton sequence data, making downstream task learning more label-efficient.

Our contribution

- A simple framework for self-supervised representation learning for skeleton-based action recognition based on BYOL.
- A data augmentation strategy for skeleton data based on two distinct transformation pipelines.
- A multi-viewpoint sampling method that makes better use of action sequences captured simultaneously by different cameras.
- We show that our method consistently outperforms the current state of the art on linear evaluation and semi-supervised tasks.

Method

Self-Supervised Skeleton Sequence Representation Learning with BYOL

Two networks, an online network and a target network, encode two augmented views of the same action sequence captured from different viewing angles. The online network is trained to predict the output of the target network, while the target network is updated with an exponential moving average of the online network.

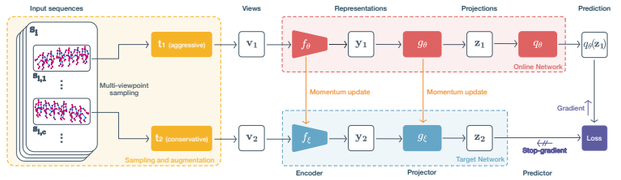

Figure 1: Overview of our proposed method.

Asymmetric Augmentation Pipelines
We tailored an aggressive data augmentation pipeline to learn semantically-relevant features, while the conservative pipeline reduces the distribution shift between self-supervised pre-training and supervised fine-tuning.

Multi-Viewpoint Sampling
We leverage recordings of the same action sequences captured simultaneously from different angles by different cameras to learn representations that are invariant to changes of viewpoint and to different camera properties.

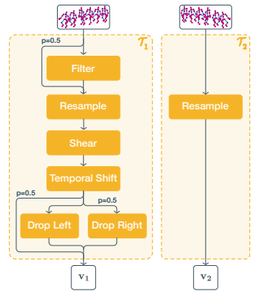

Figure 2: Asymmetric Augmentation Pipelines.

Experimental Results

Method	NTU-60 CS	NTU-60 CV	NTU-120 CSub	NTU-120 CSet
ST-GCN (supervised)	**88.5**	**94.3**	**83.0**	**85.1**
LongT GAN	39.1	48.1	-	-
MS^{2}L	52.6	-	-	-
PCRP	53.9	63.5	-	-
AS-CAL	58.5	64.8	48.6	49.2
Thoker et al.	76.3	85.2	67.1	67.9
3s-CrosSCLR	77.8	83.4	67.9	66.7
Ours	**86.8**	**91.2**	**77.1**	**79.2**

Table 1: Linear evaluation protocol on NTU-60 and NTU-120.

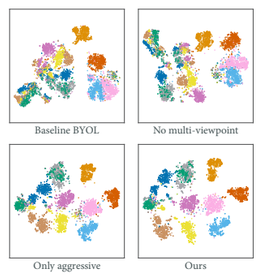

Figure 3:t-SNE projection of learned representations.

Method	Label fraction 1%	5%	10%
ST-GCN (supervised)	19.3	59.1	71.7
SESAR-KT	48.1	55.0	58.2
MS^{2}L	33.1	-	65.2
ASSL	-	57.3	64.3
3s-CrosSCLR	51.1	-	74.4
Thoker et al.	35.7	59.6	65.9
Ours 1×, distilled	**79.4**	83.6	84.6
Ours 2×	79.3	**84.5**	**86.0**

Table 2: Semi-supervised learning on NTU-60 (Cross-Subject).

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Mollevik, Iris
Umeå University / Codemill AB

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Representing Temporal Data in Semantic Graphs

Semantic parsing is the process of taking input data and translating it into a structured representation of its meaning, for example a graph. The input data has traditionally been in the form of text, but in later years semantic parsing of video content has been an active research topic. Semantic parsing of video is useful for automating several otherwise manual tasks. Examples include automatic trailer creation, automatic caption generation, compliance checking, and knowledge extraction. We introduce the research project, which aims to design a suitable formalism for semantic representation of video and multimodal content, as well as developing tools to generate said representations.

AUTONOMOUS SYSTEMS (AS) | Page 123 B

Mollevik, Iris
Umeå University / Codemill AB

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Representing Temporal Data in Semantic Graphs

Iris Mollevik

Research project

- **Semantic parsing** = to translate input data into a representation, suitable for further processing
- The traditional input has been text
- Our focus: Semantic parsing of video and multimodal data

Applications in video industry

- Automatic trailer creation
- Caption generation
- Compliance checking
- Advanced search among video content

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

CODEMILL

UMEÅ UNIVERSITET

Making **computers** understand **video content** can help the video industry automate manual tasks.

For example trailer creation and compliance checking.

It can also facilitate advanced search among video material.

Abstract Meaning Representation (AMR) is a graph format used to capture the meaning of single text sentences.

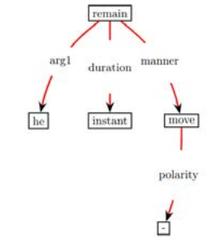

AMR graph for the sentence "He remained motionless for an instant". A very simple example.

AMR has desirable properties, however it was developed for single sentences of text. We are working on extending this formalism to be able to capture longer sequences of text, including temporal information. In the longer term, we would like to extend it to capture video and other data as well.

If the AMR format proves unsuitable, we will instead develop a different format.

This research is still in a very early stage.

Set-Membership Estimation in Shared Situational Awareness for Automated Vehicles in Occluded Scenarios

The objective of this project is to model, formalize, and analyse a shared situational awareness framework for the ego-vehicle and extended vehicles, i.e., connected vehicles and infrastructure. Shared situational awareness is the ability to perceive and comprehend the traffic situation and to predict the intent of vehicles and road users in the surrounding of the ego-vehicle using local and connected sensors. This framework will allow to orchestrate the utilization of shared resources in complex and crowded environments and to define which kind of information each Connected and Autonomous Vehicle (CAV) and the infrastructure should share. Safety-critical application such as these require robust guarantees for the estimation of the road users.

AUTONOMOUS SYSTEMS (AS) | Page 124 B
Narri, Vandana
KTH / Scania AB
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Set-Membership Estimation in Shared Situational Awareness for Automated Vehicles in Occluded Scenarios

Vandana Narri
ATS Research, Scania CV AB : *vandana.narri@scania.com*
Division of Decision and Control Systems, KTH : *narri@kth.se*

Motivation & Research goals

- The objective of this project is to model, formalize, and analyse a shared situational awareness framework for the ego-vehicle and extended vehicles, i.e., connected vehicles and infrastructure.
- Shared situational awareness is the ability to perceive and comprehend the traffic situation and to predict the intent of vehicles and road users in the surrounding of the ego-vehicle using local and connected sensors.
- This framework will allow to orchestrate the utilization of shared resources in complex and crowded environments and to define which kind of information each Connected and Autonomous Vehicle (CAV) and the infrastructure should share.
- Safety-critical application such as these require robust guarantees for the estimation of the road users.

Background

- Local CAV sensors typically provide a limited understanding of the environment due to limited sensor range, blind spots, and occlusions in the environment.
- Vehicle to vehicle (V2V) communication and vehicle to infrastructure (V2I) communication based on 5G or IEEE 802.11p standards, can help gather more information about the environment, and address the shortcomings of CAV sensors.
- CPM (Collaborative Perception Message) service supports sharing information between ITS-Ss (Intelligent Transportation System – Stations) [1].
- The main research areas are connectivity (enabled by V2I and V2V), cooperative driving, situational awareness, set-based estimation and traffic flow optimization.

Problem Formulation

- The problem considered in this research work is formulated around a scenario as shown in Figure 1.
- This scenario consist of two-lane road with a side-walk on each side of the road and a pedestrian crossing.
- The ego-vehicle (blue bus) is traveling from left to right and is approaching the pedestrian crossing. The ego-vehicle is equipped with a sensor having a field of view represented by the blue-shaded circle segment.
- In this scenario, two additional sensors are included. One on the approaching CAV represented by yellow-shaded circle segment and other on the connected road-side sensor units represented by green-shaded circle segment.

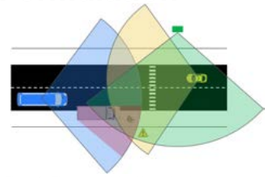

Figure 1 : The ego-vehicle with one local sensor, an additional V2V sensor and an additional V2I sensor.

References

1. Draft ETSI TS 103 324 V0.0.22 Collective Perception Service.
2. Vandana Narri, A. Alanwar, J. Mårtensson, C. Norén, L. Dal Col and K. H. Johansson, "Set-Membership Estimation in Shared Situational Awareness for Automated Vehicles in Occluded Scenarios," 2021 IEEE Intelligent Vehicles Symposium (IV).

Architecture of Share Situational Awareness

- The proposed architecture is presented in Figure 2. It consists of three parts: (i) Local and extended sensor network, (ii) Algorithms for shared situational awareness, and (iii) Decision-making.
- Measurement data from the sensors are collected and fused to perform state estimation.
- Based on these estimates, decisions are made, and actions are planned. In this paper, the focus is on (i) and (ii).

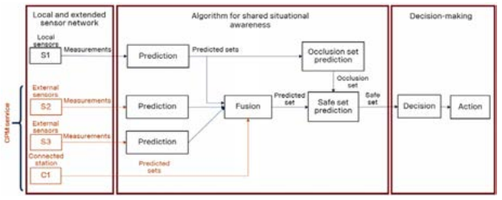

Figure 2 : Proposed architecture for set-based estimation for shared situational awareness.

Discussion

- In this work, set-based approach is considered, which models the noise and disturbance as unknown variables with known bounds.
- One of the most popular set-based approach is set-membership estimator, which is implemented in this project. And in this approach set of states are considered instead of a single state for estimations which will help in providing robust guarantees and safety margins.
- The set are mathematically represented using zonotopes as shown in Figure 3.

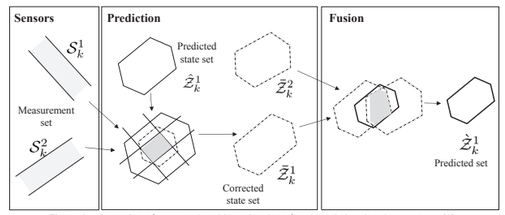

Figure 3 : Illustration of set-membership estimations for shared situational awareness [2].

Nelson, Christian
Lund University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A Multilink Channel Measurement System

Wireless channels in vehicular environments are highly dynamic. To evaluate the propagation channel at any given time, they require measurements of all radio channels between several radios simultaneously. Additionally, vehicles and roadside units are distributed at different locations which means the radios cannot be synced using a shared reference source via cables. To this end, a new channel sounder was developed using software-defined radios. This allows each radio to connect to a host computer and a stable rubidium clock. Another source of reference that can be used – when under the open sky – is the pulse-per-second (PPS) transmitted in the GNSS signals. The system has been implemented in the National Instruments software suit LabVIEW.

AUTONOMOUS SYSTEMS (AS) | Page 125 B

Nelson, Christian
Lund University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A Multilink Channel Measurement System

Christian Nelson, PhD Student, Lund University
Dept. of Electrical Engineering and Information Technology
Supervisors: Prof. Fredrik Tufvesson

LUND UNIVERSITY

Motivation & Research Goals

Wireless channels in vehicular environments are highly dynamic. To evaluate the propagation channel at any given time, they require measurements of all radio channels between several radios simultaneously. Additionally, vehicles and roadside units are distributed at different locations which means the radios cannot be synced using a shared reference source via cables. To this end, a new channel sounder was developed using software-defined radios which allows each radio to connect to a host computer and to a stable rubidium clock. Another source of reference that can be used – when there is a clear view of the sky – is the pulse-per-second (PPS) transmitted in the GNSS signals. The system has been implemented in the National Instruments software suite LabVIEW 2021.

Methods

The radios used have a maximum instantaneous bandwidth of 40 MHz and can be tuned in the range of 1.2 GHz to 6 GHz. For vehicular communication, the frequency of interest is around 5.9 GHz, which is within the operational range. Vehicles and roadside units are distributed in space, so the radios cannot be synced using cables from a shared reference source. Rather, each radio is connected to its host computer for control and a rubidium clock for a stable reference clock. Before each measurement, the rubidium clocks are connected to each other to be synchronized, and can thereafter be separated and maintain coherent for a long enough time to perform needed measurements. If the view of the sky is unobstructed, the internal clocks on the radios can be disciplined using the pulse-per-second (PPS) transmitted in the GPS signal.

[1]The sounding technique used is a correlative type, which means that the transmitted signal inhibits good autocorrelation properties. The signal used is the Zadoff-Chu sequence. It has the good autocorrelation properties that we desire and has a flat frequency response. Each radio access the channel in a predefined order using a time-division multiple access (TDMA) scheme.

Additionally, a RF front-end have been designed and built to control the signal paths and to amplify the transmitted signal.

[2]The system is implemented using LabVIEW, on a Windows 10 industrial-grade computer. Some of the requirements on the computer are that it needs to be portable for field measurements, and it needs to draw power from car batteries. These requirements limit the performance. The computer connects to the radio via an external PCIe interface which allows for transfer rates up to 200 MSamples/s.

Assembled System and Future Work

Below is two figures showing NI LabVIEW 2021 running the host user interface, and the hardware. The system has not yet been tested in the field, but there are measuremnts planned for spring 2022.

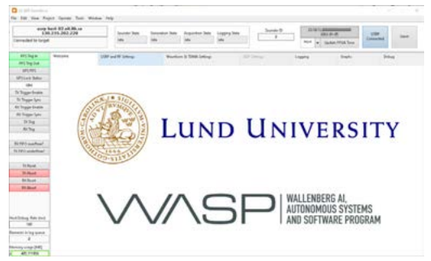

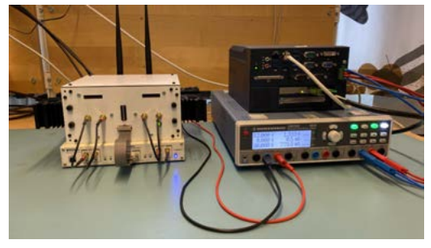

New radios have been acquired, NI-USRP X410 (shown below), with 10x larger instantaneous bandwidth (400 MHz) and 4 RF chains per radio. The work have been initiated to refactor parts of the code to use the new radios. Then radios pose some new challenges regarding data management since it can generate up to 8 GB of data per second.

References

[1] Wireless Communication, 2nd Edition
Andreas F. Molisch
John Wiley & Sons, 2012

[2] National Instruments LabVIEW
https://www.ni.com/labview
Accessed 2022-01-06

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Multi-Map SLAM

An environment that changes in between visits causes problems for long-term positioning. Can SLAM algorithms handle non-static environments by considering multiple hypotheses of landmark positions?

AUTONOMOUS SYSTEMS (AS) | Page 126 B
Nielsen, Kristin
Linköping University | WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Multi-Map SLAM

Kristin Nielsen, Linköping University
Automatic Control
Main supervisor: Gustaf Hendeby

LINKÖPING UNIVERSITY | Epiroc

An environment that changes in between visits causes problems for **long-term positioning**.

Can SLAM algorithms handle non-static environments by considering **multiple hypotheses** of landmark positions?

A robot is travelling along a corridor with **unique markers** visible to the robot attached to the door handles. As the robot moves the doors are **opened or closed**, changing how the robot perceives the environment.

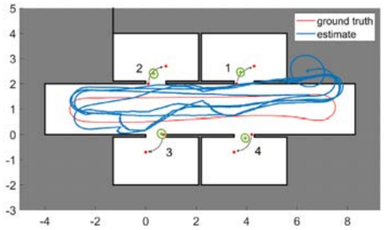

Static landmarks assumed

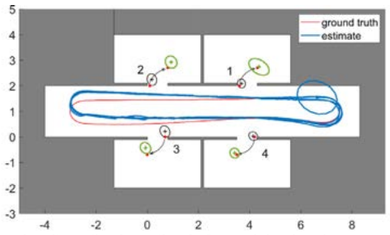

An alternative **hypothesis** for each landmark position. Black circles marks deactivated hypotheses.

A **static world** assumption gives inconsistent estimates of the robot position. By allowing **multiple hypotheses** of landmark positions and statistically decide the most likely hypothesis, the accuracy of the estimate is improved.

Hypotheses

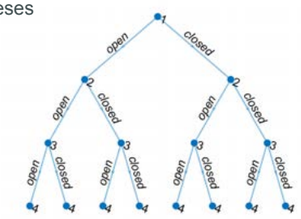

A **hypothesis tree** is maintained where each branch represents one hypothesis.

Multi-Map

In the changing environment of an **underground mine**, different version of the map can be maintained and activated/deactivated online.

Conclusion

A non-static environment that changes in between visits can be represented by a multiple hypothesis map. Decisions and inference in such a map representation has strong connections to theory used in multiple hypothesis tracking.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Developing Tools and Analyzing Methods for Secure Software Update

Secure Over-The-Air (OTA) software upgrade or update is an important maintenance aspect of any network specifically IoT networks. As a result it is important to figure out which security configuration can affect the security or energy consumption of the devices in the network, and which key sharing and key management scheme, or the actual upgrade procedure is more applicable to the network in case of energy efficiency and security.
To handle these issues, in one of our work, we tried to do an actual OTA update in an IoT environment using CoAP and MQTT protocols to see how security can affect the energy consumption of IoT devices. In another work, we have designed RoSym, a robust, secure and pure symmetric based software upgrade solution for IoT networks. Managing and provisioning of symmetric keys is difficult, as a result in another work, we present Flowrider, a novel key provisioning mechanism for cloud networks that unlocks scalable use of symmetric keys and significantly reduces the related computational load on network endpoints with the use of SDN model. Flowrider makes key distribution agnostic of the network topology and communication patterns, of which it does not require any early knowledge.

AUTONOMOUS SYSTEMS (AS) | Page 127 B
Nikbakht Bideh, Pegah
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Developing Tools and Analyzing Methods for Secure Software Update

Pegah Nikbakht Bideh, Martin Hell, Nicolae Paladi
Department of Electrical and Information Technology, Lund University, Sweden

Introduction

Secure Over-The-Air (OTA) software upgrade or update is an important maintenance aspect of any network specifically IoT networks. As a result it is important to figure out which security configuration can affect the security or energy consumption of the devices in the network, and which key sharing and key management scheme, or the actual upgrade procedure is more applicable to the network in case of energy efficiency and security. To handle these issues, in one of our work, we tried to do an actual OTA update in an IoT environment using CoAP and MQTT protocols to see how security can affect the energy consumption of IoT devices. In another work, we have designed RoSym, a robust, secure and pure symmetric based software upgrade solution for IoT networks. Managing and provisioning of symmetric keys is difficult, as a result in another work, we present Flowrider, a novel key provisioning mechanism for cloud networks that unlocks scalable use of symmetric keys and significantly reduces the related computational load on network endpoints with the use of SDN model. Flowrider makes key distribution agnostic of the network topology and communication patterns, of which it does not require any early knowledge.

Background

IoT protocols: CoAP and MQTT are the most common application layer protocols used in IoT environments. CoAP uses UDP at transport layer, while MQTT uses TCP as transport protocol. To add security at transport layer, the natural choice is to use DTLS for CoAP and TLS for MQTT. There are four security modes such as Nosec, PreSharedKey, RawPubicKey and Certificates which can be used for both CoAP and MQTT protocols.
OTA upgrade: In IoT environments, software update and other critical maintenance operations need to be performed over the network and when the device is wireless this often referred to as Over-The-Air (OTA) update/upgrade.
SDN and OpenFlow: SDN emerged in response to the increasing complexity of network deployments, facilitating operation and management of virtualized networks. SDN components are:

- Data plane
- Control plane
- Network functions

Energy Consumption for Securing Lightweight IoT Protocols

The ubiquitous nature of IoT devices often requires them to run on batteries, making energy efficiency a primary concern. The large number of devices make it costly to replace batteries, and it will also make the total energy consumption considerable. At the same time, adding security to the communication will add additional overhead. Thus, it is important to not only develop lightweight security protocols, but also to understand to which extent security affects the energy consumption of the devices. The main contribution of our work is a thorough analysis of CoAP and MQTT and the investigation of their energy footprint in different scenarios:

- **How added security at the transport layer (TLS/DTLS) affects the energy consumption?**
- **How important design choices, such as cipher suite, PKI vs. PSK, and client authentication impact the energy consumption?**

In our experiments we use ESP32 with libcoap, MQTT, and mbed TLS libraries and conduct real-world measurements using Otii.

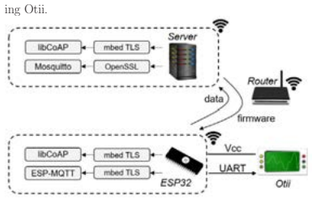

We find that aggregating data to larger packets can significantly reduce the energy consumption. We also find that AES-CCM8 seems slightly more efficient than other modes of operation.

RoSym: Robust Symmetric Key Based IoT Software Upgrade Over-the-Air

Despite the fact that the software upgrade problem have been extensively studied, there is still a need for investigating new design of completely symmetric key-based upgrade solutions with post-quantum resistant. Our new scheme RoSym is such a solution. Furthermore, the upgrade process should not be sensitive to Denial-of-Service (DoS) attacks. In particular, it should not be easy for an attacker to force IoT units to waste extensive power battery resources on false software upgrade requests. These properties should be achieved while simultaneously offering a high level of confidentiality and integrity protection, working well for resource constraint units, and avoiding public key operations during software upgrade.
Our solution is consists of two phases:
Key Establishment and Parameter Setup: During this phase the following information will be tranferred to IoT devices:

1. Two randomly generated symmetric keys, IK_{sw} and K_{sw}.
2. Timing information, T_1 and T_2 determines the validity of update packets and validity of IK_{sw} and K_{sw} keys respectively.
3. A software image one-way hash root value, h_0.
4. The number of packages (n) in the new software distribution.

Upgrade procedure phase: Before the upgrade procedure starts, the DMS package the software update image into suitable size packages.

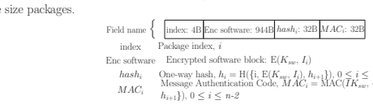

The following corresponding upgrade procedure applies on the IoT side:

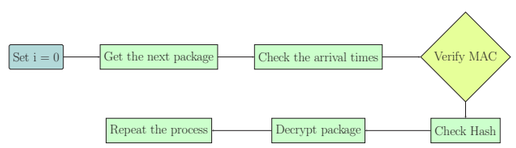

Our security analysis showed that RoSym offers the expected security level.

Flowrider: Fast On-Demand Key Provisioning for Cloud Networks

Generating symmetric keys requires less computational power and has firmware support on many platforms, the use of symmetric keys leads to challenges such as secure key provisioning and key authentication. This introduces the research question:
Can the SDN model be leveraged to conveniently provision symmetric keys and reduce computational resource consumption?
Yes with the use of Flowrider, a novel key provisioning mechanism for network endpoints in SDN deployments that considers the practicalities of cloud systems deployment.

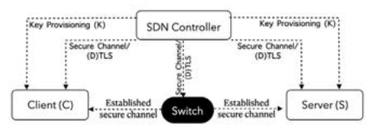

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Nordlöf, Jonas
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Planning for minimal uncertainty

A belief-space planning problem for GNSS-denied areas is studied where the location and number of landmarks available are unknown when performing the planning.
To be able to plan an informative path in this situation, an algorithm using virtual landmarks to position the platform during the planning phase is studied.
The virtual landmarks are selected to capture the expected information available in different regions of the map, based on the beforehand known landmark density.
The approach is tested in a simulated environment, in conjunction with an extended information filter, with successful results.

AUTONOMOUS SYSTEMS (AS) | Page 128 B

Nordlöf, Jonas
Linköping University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Planning for minimal uncertainty

Jonas Nordlöf, Gustaf Hendeby and Daniel Axehill

Summary

Objective: Reduce position uncertainty in landmark-based SLAM by using motion planning.

Challenge: Landmark positions are unknown.

Solution: Introduce approximation using virtual landmarks based on landmark densities.

Outcome: Ability to plan minimal uncertainty path and predict position uncertainty.

Stochastic optimization problem

Future position estimates depend on the actual motion and the obtained measurements of the landmarks positions, which are both unavailable at the planning stage. This leads to a stochastic optimization problem:

$$\begin{aligned} \underset{\pi,I}{\text{minimize}} \quad & \mathbb{E}\big(J(\mathcal{I}_{t|t})\big) && \text{expected performance} \\ \text{subject to} \quad & x_{t+1} = f(x_t, u_t, w_t), && \text{platform state dynamics} \\ & y_t^i = h^i(x_t) + \epsilon_t^i,\ \forall i \in \mathcal{M}_t, && \text{landmark measurements} \\ & \mathcal{I}_{t+1|t+1} = \Lambda(\mathcal{I}_{t|t}, y_t, x_t, u_t), && \text{information model} \\ & u_t = \pi(x_t, \mathcal{I}_{t|t}), && \text{control policy} \\ & \epsilon_t^i \sim \mathcal{N}(0, R_t), \quad w_t \sim \mathcal{N}(0, Q_t) \end{aligned}$$

This problem is generally not solvable.

Deterministic approximation

The stochastic problem can be replaced by a deterministic problem using the following approximations:

Unknown noise realizations: Certainty-equivalent control

Unknown landmark positions: Virtual landmarks based on landmark densities $\rho(\cdot)$

- Each virtual landmark represents a subregion Ω_i.
- Information gained from observing a subregion Ω_i can be calculated as

$$\mathcal{I}_t^i = \int_{\tilde{m}\in\Omega_i} \rho(\tilde{m}) \big(H_t(p_t, \tilde{m})\big)^\top R_t^{-1} H_t(p_t, \tilde{m}) d\tilde{m}. \tag{1}$$

yielding the information update:

$$\mathcal{I}_{t+1|t+1} = \big(F_t \mathcal{I}_{t|t}^{-1} F_t + G_t Q_t G_t\big)^{-1} + \sum_{i\in\mathcal{M}_t} \mathcal{I}_t^i \tag{2}$$

- For range-bearing measurements (1) has closed form solution for circle sectors.
- Approximate Ω_i with circle sectors $\Omega_{i,k}$.

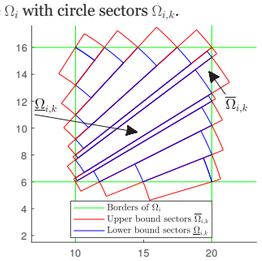

Simulation study

The approach is evaluated using Monte Carlo simulations in a forest environment. Trees are used as landmarks.

- Planned path using the proposed approach (red) and Monte Carlo realizations of path (blue)

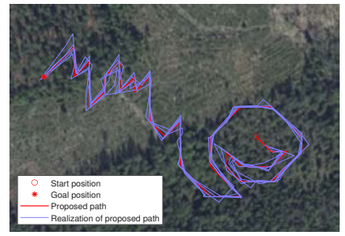

- Estimated performance bounds (blue) and Monte Carlo estimate of performance measure (red)

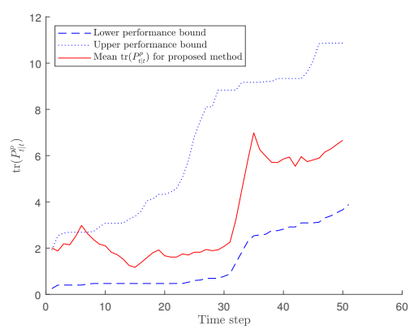

Future work

- Apply the method in a real environment
- Investigate impact of terrain properties on position estimate
- Add known landmarks and visual sensors
- Position correction using receding horizon control

References

J. Nordlöf, G. Hendeby, and D. Axehill, "Belief space planning using landmark density information", in 2020 Proc. IEEE 23rd Int. Conf. Inf. Fusion (FUSION), Rustenburg, South Africa, Jul. 2020, pp. 1–8.

J. Nordlöf, G. Hendeby, and D. Axehill, "Improved Virtual Landmark Approximation for Belief-Space Planning", in 2021 Proc. IEEE 24rd Int. Conf. Inf. Fusion (FUSION), Rustenburg, South Africa, Nov. 2021, pp. 1–8.

Acknowledgment & Collaborations

This work is partially supported by the WASP Affiliated PhD Student Program. The project is being funded by the division of C4ISR within the Swedish Defence Research Agency (FOI). The project follows the Swedish Armed Forces Research and Tech program for the area of autonomous systems and the area of sensors and signature management.

LINKÖPING UNIVERSITY WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM FOI

Nyberg, Truls
KTH / Scania CV AB

Foresee the Unseen: Sequential Reasoning about Hidden Obstacles for Safe Driving

Safe driving requires autonomous vehicles to anticipate unseen objects, such as a cyclist hidden behind a large vehicle, or an object on the road hidden behind a building.
Existing methods are usually not able to consider all possible shapes and orientations of such obstacles.
They also typically do not reason about observations of hidden obstacles over time, leading to conservative anticipations.
We overcome these limitations by (1) modeling possible hidden obstacles as a set of states of a point mass model and (2) sequential reasoning based on reachability analysis and previous observations.
Based on (1), our method is safer, since we anticipate obstacles of arbitrary unknown shapes and orientations.
In addition, (2) increases the available drivable space when planning trajectories for autonomous vehicles, which we demonstrate can give rise to significant reductions in time when traversing various intersection scenarios.

Nyberg, Truls
KTH / Scania CV AB

Foresee the Unseen

Truls Nyberg*, José Manuel Gaspar Sanchez*, Christian Pek, Jana Tumova, Martin Törngren
Affiliated with KTH Royal Institute of Technology and Scania CV AB
*Contributed equally to this research. Article currently under review.

Sequential Reasoning about Hidden Obstacles for Safe Driving

Safe driving requires autonomous vehicles to anticipate **potential hidden traffic participants and objects**. Existing methods typically do not consider **arbitrary shapes** of hidden obstacles and do not reason about **observations over time**. We overcome these limitations by (1) **modeling possible hidden obstacles as a set of states** of a point mass model and (2) **sequential reasoning based on reachability analysis** and previous observations.

The Problem

Autonomous vehicles need to model possible hidden obstacles **conservatively enough**, such that any possible unseen obstacle is represented and considered, regardless of their size or orientation, such as the motorcycle in **Figure 1**.

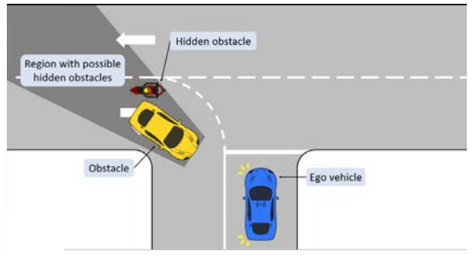

***Figure 1**. A dangerous situation where defensive driving is needed.*

However, modeling possible hidden obstacles **too conservatively** limit autonomous vehicles from finding safe and efficient paths.

Given **past observations** and assumed constraints on driving (e.g., maximum speed and other traffic rules), currently unseen regions can still be concluded free from obstacles, such as the checkered region in **Figure 2**.

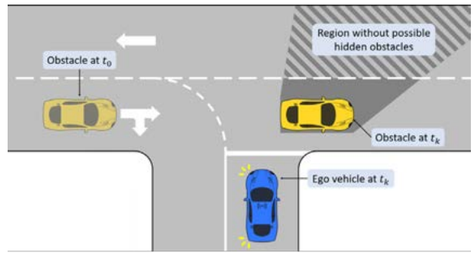

***Figure 2**. A typical situation where autonomous vehicles generally are too conservative. The checkered region was previously seen, and it can be concluded that no object could have reached there.*

The Solution

By **initially** considering the complete unseen region as potentially occupied, our method captures **any** hidden obstacle (**Figure 3a**). By **iteratively** updating which regions possibly can be occupied, we avoid being **too conservative.**

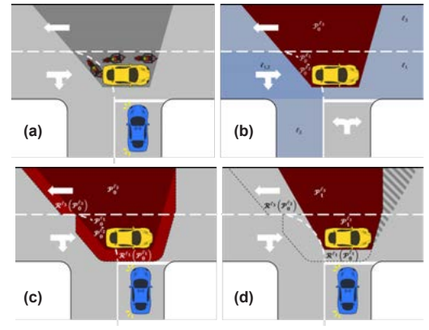

***Figure 3**. Algorithmic steps for reasoning.*

For each lane (**Figure 3b**), the reachability is computed for the possible hidden obstacles (bright red in **Figure 3c**). New unseen regions are deemed free if they cannot have been reached since the last observation (the checkered region in **Figure 3d**). The result can be seen in **Figure 4**, where the time to traverse the intersection is **greatly reduced** by reasoning about possible hidden obstacles over time.

***Figure 4**. Simulation in CommonRoad of intersection in Fürstenfeldbruck, using the proposed algorithm.*

Oxenstierna, Johan
Lund University / Kairos Logic AB

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Towards data-driven Storage Location Assignment: New benchmarks and optimization model

The Storage Location Assignment Problem (SLAP) is concerned with the choice of locations for products in a warehouse. It is of primary significance for operational quality since the travel cost of order-picking vehicles is strongly related to where and how far they have to travel. Unfortunately, a generalized model of the SLAP poses a highly intractable problem. State-of-the-art optimization methods tend to be usecase-specific and there exists no standard benchmark dataset format. In this study we introduce new benchmark data on a modified TSPLIB format and demonstrate how instances can be approximately solved using a state-of-the-art Order Batching Problem (OBP) optimizer aided by a Quadratic Assignment Problem (QAP) surrogate. Our results show that the OBP optimizer can yield significantly better performance when it is aided by the surrogate.

Oxenstierna, Johan
Lund University / Kairos Logic AB

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Towards data-driven storage assignment: New benchmarks and optimization model

Johan Oxenstierna, Jacek Malec, Volker Krueger
Computer Science Dept., Lund University

The Storage Location Assignment Problem (SLAP) is concerned with the choice of locations for products in a warehouse. It is of primary significance for operational quality since the travel cost of order-picking vehicles is strongly related to where and how far they have to travel. Unfortunately, a generalized model of the SLAP poses a highly intractable problem. State-of-the-art optimization methods tend to be usecase-specific and there exists no standard benchmark dataset format. In this study we introduce new benchmark data on a modified TSPLIB format and demonstrate how instances can be approximately solved using a state-of-the-art Order Batching Problem (OBP) optimizer aided by a Quadratic Assignment Problem (QAP) surrogate. Our results show that the OBP optimizer can yield significantly better performance when it is aided by the surrogate.

Contributions

1. Introduction of SLAP benchmark data on the TSPLIB format.
2. Evaluation of the usage of QAP surrogates within SLAP optimization.

Problem formulation

The general SLAP objective is to assign locations for products such that the distance to pick the products, using order batching, is minimized [1]:

$$min \sum_{b \in \mathcal{B}} D(b) x_{mb}, m \in \mathcal{M}$$

where $\mathcal{B}$ denotes generated batches, where $D(b)$ denotes the distance to pick batch b (a solution to a Traveling Salesman Problem where the node-location pointers are mutable), where m denotes a vehicle and where x_{mb} denotes a binary variable that is 1 if vehicle m is assigned to pick b and 0 otherwise.

An optimizer would require significant computational time to find a value close to the minimum above since the problem is NP-hard. For a faster but less precise quality evaluation of a product to location assignment, a Quadratic Assignment Problem (QAP) surrogate is used [2]. The QAP uses distances and weights between products. The weights between two products is here defined as the number of times two products appear in the same order. The QAP function solution value is computed by:

$$\sum_{P_1 \in P} \sum_{\substack{p_2 \in P \\ p_1 \neq p_2}} \sum_{l_1 \in L} \sum_{\substack{l_2 \in L \\ l_1 \neq l_2}} w_{p_1 p_2} \, d_{l_1 l_2} x(p_1, l_1) x(p_2, l_2)$$

where w denotes weight, where d denotes distance and $x(p, l)$ a function which returns 1 if product p is located at location l and 0 otherwise.

Optimization model

An order batching optimizer (Single Batch Iterated) with a Quadratic Assignment Surrogate (SBI-QAS):

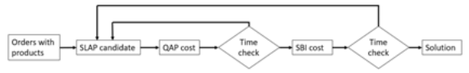

This model is delimited in two major ways: 1. No bootstrapping is used to generate SLAP candidates. 2. The QAP stopping criterion is time-based rather than convergence based. These are motivated since they are bias-imposing techniques [3] and they should only be implemented after the accuracy of the surrogate can be established through experimentation.

Surrogate evaluation metric

As shown in the above diagram, the slower OBP optimizer is run on the most promising set of candidate solution(s), whose selection is determined by ordering the QAP surrogate prediction values. Hence, to evaluate the surrogate quality, we are not interested in the QAP estimates themselves, but rather the surrogate's ability to correctly rank those estimates. We therefore use the Normalized Discounted Cumulative Gain [4]:

$$NDCG(k) = 1 - \frac{DCG(k)}{IDCG(k)}, \quad DCG(k) = \sum_{i=1}^{k} \frac{G_i}{log_2(i+1)}, \quad IDCG(k) = \sum_{i=1}^{|I(k)|} \frac{G_i}{log_2(i+1)}$$

where k denotes an index and G denotes gain. NDCG compares the ranking quality of a candidate ranker against a ground truth ranker.

Experiment

The aim of the experiment is to empirically test the QAP surrogate predictive strength against a ground truth ranker (SBI).

Benchmark data and format

The public TSPLIB format datasets in L6_203 and L09_251 are modified for the SLAP [5]. For analysis, the instances are divided into 7 classes according to number of orders in the test data (10 – 1000). The experiment is set up such that 60 randomly generated SLAP candidate solutions are generated for each instance. For each of these SLAP candidates the QAP and OBP estimates and CPU-times are recorded. In total there are 2274 candidates and corresponding predictions for the QAP and OBP modules.

Experiment result

On average, the QAP surrogate predictive error ranges from 35 - 44% (the data is normalized such that anything below 50% means the proposed algorithm is successful), with standard deviations in the range 12-16%. A slight pattern for larger predictive errors for larger instance sizes can be observed, but this is inconclusive.

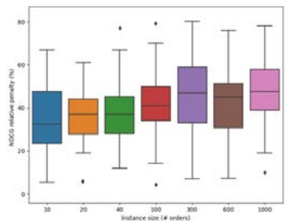

Figure 1: Box-plot showing instance sizes on the x-axis (in terms of number of orders) and NDCG on the y-axis (how wrong the QAP ranking is on average). The inner and outer boxes represent 95% and 99% confidence intervals respectively.

This result provides evidence that the QAP surrogate can be successfully utilized within the larger algorithm. For future work the distribution of work between the QAP surrogate and the SBI optimizer could be investigated. It was found that the QAP surrogate requires around 60X less CPU-time than SBI, so an experiment could be set up with various number of surrogate solution candidates.

References

1. Gademann and Velde: Order batching to minimize total travel time in a parallel-aisle warehouse, IIE Transactions, 2005.
2. Kubler, Glock and Bauernhansl: A new iterative method for solving the joint dynamic storage location assignment, order batching and picker routing problem in manual picker-to-parts warehouses, Computers & Industrial Engineering, 2021.
3. Sutton and Barto: Reinforcement Learning: An introduction, MIT Press, 2018.
4. Busa-Fekete, Róbert, et al. "An apple-to-apple comparison of learning-to-rank algorithms in terms of normalized discounted cumulative gain." *ECAI 2012-20th, 2012.*
5. Oxenstierna et al: https://github.com/johanoxenstierna/OBP_instances, https://github.com/johanoxenstierna/L09_251, 2021.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Optimal Input Design Through Infinity Norm Minimization Using Proximal Mapping

To avoid non-convexity of the criterion, various relaxations are typically used in input design. For example, the input may be assumed to be stationary and the design problem may be formulated in terms of the correlation coefficients. Now, we instead propose a method to directly design the input sequence. This allows to maximize the information obtained from short-time (transient) experiments using non-stationary inputs. We do this by fitting the achieved Fisher matrix to a desired target matrix in a matrix sense, using the infinity norm. The target matrix can either be the desired Fisher matrix, obtained from quality considerations of the intended use of the model, or a matrix directly representing the performance of the application. An often used quantity is the Hessian of the so called the application cost. Thus, the method is formulated as a time domain optimization problem that is non-convex. This optimization problem is solved by alternative minimization and proximal mapping.

Optimal Input Design through Infinity Norm Minimization

Javad Parsa, PhD student, KTH University
Division of Decision and Control Systems
Supervisor: Prof. Håkan Hjalmarsson

Abstract

To avoid non-convexity of the criterion, various relaxations are typically used in input design. For example, the input may be assumed to be stationary and the design problem may be formulated in terms of the correlation coefficients. In this contribution, we instead propose a method to directly design the input sequence. This allows to maximize the information obtained from short-time (transient) experiments using non-stationary inputs. We do this by fitting the achieved Fisher matrix to a desired target matrix in a matrix sense, using the infinity norm. The target matrix can either be the desired Fisher matrix, obtained from quality considerations of the intended use of the model, or a matrix directly representing the performance of the application. An often used quantity is the Hessian of the so called the application cost. Thus, the method is formulated as a time domain optimization problem that is non-convex. This optimization problem is solved by alternating minimization and proximal mapping.

Methods

Linear regression model:

$$\mathbf{y} = \mathbf{\Phi}\theta + \mathbf{e}, \quad \mathbf{e} \sim \mathcal{N}(0, \lambda \mathbf{I})$$

where $\theta \in R^{n_\theta \times 1}$, $\mathbf{\Phi} \in R^{N \times n_\theta}$ and $\mathbf{e} \in R^{N \times 1}$.
Fisher Information Matrix (FIM):

$$\mathbf{I}_F = \frac{1}{\lambda}\mathbf{\Phi}^T\mathbf{\Phi}.$$

Identification set:

$$\varepsilon_{si} = \{\theta : (\theta - \theta_0)^T \mathbf{I}_F (\theta - \theta_0) \leq \chi^2_\alpha(n_\theta)\}$$

where $\chi^2_\alpha(n_\theta)$ is the α-percentile of the chi-square distribution with n_θ degrees of freedom and θ_0 denotes true parameter vector.
Application performance cost V_{app}: measures the performance degradation when the model parameter θ is used in a model based design.

$$\varepsilon_{app} = \{\theta : (\theta - \theta_0)^T V''_{app}(\theta_0)(\theta - \theta_0) \leq \frac{1}{\gamma}\}.$$

To guarantee that the estimates are inside the application set with high probability, we need to ensure the identification set is inside the application set [1], i.e.:

$$\mathbf{I}_F = \frac{1}{\lambda}\mathbf{\Phi}^T\mathbf{\Phi} \geq \gamma\chi^2_\alpha(n_\theta)V''_{app}(\theta_0).$$

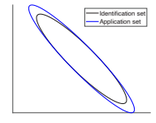

State-of-the-art: Per sample Fisher matrix (means that the input is stationary) designed using convex optimization. Results in a desired Fisher matrix $\mathbf{I}_F^d$. Will be denoted the Frequency Domain Method (T-FDM).

Proposed method:

To find the optimal regressor, the following optimization problem is proposed:

$$\min_{\mathbf{\Phi} \in \mathcal{D}} \|\mathbf{\Phi}^T\mathbf{\Phi} - \mathbf{T}\|_\infty.$$

in which the $\mathcal{D}$ denotes the set of $N \times n_\theta$ Toeplitz matrices and the infinity norm of a matrix means maximum of the absolute values of the elements of this matrix. To solve the above optimization problem, we use proximal mapping and a closed form solution to update the regressor $\mathbf{\Phi}$ can be found [2].

- Target matrix:
  1) $\mathbf{T} = \lambda\mathbf{I}_F^d$-(T-FDM)
  2) $\mathbf{T} = \lambda\chi^2_\alpha(n_\theta)\gamma V''_{app}(\theta_0)$- the scaled Hessian of APPlication cost (T-APP).

Selected Results

Feedforward control problem:

State Space Model:

$$\begin{bmatrix} x_1(t+1) \\ x_2(t+1) \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} u(t)$$

$$y(t) = \begin{bmatrix} \theta_1 & \theta_2 \end{bmatrix} \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix} + d(t), \quad d(t) = \begin{cases} 1, & t \geq 0 \\ 0, & t < 0 \end{cases}$$

Feedforward controller:

$$u(t) = \frac{-d(t)}{\theta_1 + \theta_2}$$

Application Cost:

$$V_{app}(\theta) = \frac{1}{B}\sum_{t=1}^{B} \big(y(t,\theta) - y(t,\theta_0)\big)^2.$$

Optimal input design:

$$\min_{\phi_u(\omega)} \quad \Big(\frac{1}{2\pi}\int_{-\pi}^{\pi} \phi_u(\omega)\mathrm{d}\omega\Big)$$
$$\text{st} \quad \mathbf{I}_F \geq \chi^2_\alpha(n_\theta)\gamma V''_{app}(\theta_0),$$

The outcome of above optimization problem (using the MOOSE2 toolbox) is the desired per sample matrix $\mathbf{I}_F^d$. Then, we take $\mathbf{T} = \mathbf{I}_F^d$ and T-FDM (Target matrix using per sample FIM from FDM) denotes this case.

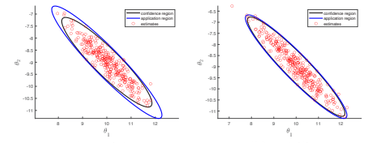

(a) T-FDM **(b)** T-APP

From these figures we see that approximately 99% and 98% of the estimates are inside the application cost for T-FDM and T-APP, respectively.

References

[1] System identification of complex and structured systems
H. Hjalmarsson, European Journal of Control, 2009.

[2] Optimal Input Design Through Infinity Norm Minimization Using Proximal Mapping, J. Parsa and H. Hjalmarsson, 60th IEEE Conference on Decision and Control conference, CDC 2021.

Mission-critical Applications at the Edge of 5G

In this work, we performed the evaluation of real-time HTTP applications that can be deployed at the edge of a complete mid-band stand alone 5G base station. We showed the advantages and impacts of a real 5G network system on mission-critical processes, which are envisioned to benefit from the ultra-low latency and extreme-high bandwidth of the 5G.

Value-Maximizing Combinatorial Assignment

We investigate value-maximizing combinatorial assignment—i.e., the problem of partitioning items into bundles among a set of alternatives to maximize some notion of social welfare. This problem is a major research challenge in computer science with many applications in for example operations research, economics, and artificial intelligence. Unfortunately, combinatorial assignment problems are in general both NP-hard and inapproximable.

Value-Maximizing Combinatorial Assignment

Fredrik Präntare (fredrik.prantare@liu.se)
Department of Computer Science (AIICS/ReaL), Linköping University

Background & Motivation

We investigate *value-maximizing combinatorial assignment*—i.e., the problem of partitioning items into bundles among a set of alternatives to maximize some notion of social welfare. This problem is a major research challenge in computer science with many applications in for example operations research, economics, and artificial intelligence. Unfortunately, combinatorial assignment problems are in general both NP-hard and inapproximable.

Contributions

Our contributions include developing the state-of-the-art algorithms for combinatorial assignment, as well as theoretical and empirical advances. Our algorithms have also been applied to (and are being used in) a commercial real-world application with more than 1 million users (Fig. 1).

Fig. 1: Our algorithms are used in the commercial strategy game *Europa Universalis 4* (a game with more than 1 million players) to coordinate and deploy armies to different regions.

Publications

[1] Präntare, F., Ragnemalm, I., & Heintz, F. (2017).
Lilla Polhemspriset
An Algorithm for Simultaneous Coalition Structure Generation and Task Assignment. In *International Conference on Principles and Practice of Multi-Agent Systems*.

[2] Präntare, F., & Heintz, F. (2018).
Best Student Paper Award
An Anytime Algorithm for Simultaneous Coalition Structure Generation and Assignment. In *International Conference on Principles and Practice of Multi-Agent Systems*.

[3] Präntare, F., & Heintz, F. (2019).
An Anytime Algorithm for Optimal Simultaneous Coalition Structure Generation and Assignment. In *Journal of Autonomous Agents and Multi-agent Systems*.

[4] Präntare, F., & Heintz, F. (2020).
Best Student Paper Award
Dynamic Programming for Optimal Simultaneous Coalition Structure Generation and Assignment. In *International Conference on Principles and Practice of Multi-Agent Systems*.

[5] Präntare, F., Appelgren, H., & Heintz, F. (2021).
Anytime Heuristic and Monte Carlo Methods for Large-Scale Simultaneous Coalition Structure Generation and Assignment
In *AAAI Conference on Artificial Intelligence*.

[6] Präntare, F., Tiger, M., Bergström, D., Appelgren, H., & Heintz, F. (2022).
Learning Heuristics for Combinatorial Assignment Problems by Optimally Solving Subproblems.
In *International Conference on Autonomous Agents and Multi-Agent Systems*.

[7] Präntare, F., Osipov, G., Eriksson, L. (2022).
Concise Representations and Complexity of Combinatorial Assignment Problems.
To be published.

Power Efficient Multi DNN Accelerator for Future IoT Devices

Future IoT devices and autonomous systems would have widespread use of deep learning. It is increasingly important to do those heavy computations on the device itself without offloading them to the cloud. Those devices often use specialized domain-specific DNN accelerators to achive high computation efficiency. However, those devices are battery-powered and power and area constrained. On the other hand, future applications would require simultaneous execution of multiple network models on the same device concurrently. Further due to being chip area constrained those devices could not be provisioned with large compute and memory resources. Limited on-chip memory leads to frequent off-chip memory accesses which dominates the power consumption in such devices and multi DNN acceleration aggravates this problem where the on-chip memory should be shared between multiple DNNs. Therefore this research aims at exploring architectural techniques for achieving better power efficiency for multi DNN acceleration such as stream caching and efficient memory management for such accelerators which could pothentially reduce the number of off-chip accesses.

AUTONOMOUS SYSTEMS (AS) | Page 134 B
Ranawaka, Piyumal
Chalmers
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Topological Method for fMRI Analysis

Functional magnetic resonance imaging (fMRI) is used to measure brain activity due to tasks or stimuli. Resting-state measurements are used to provide a subject's baseline. The signal is prone to noise from various sources. Random brain activity and noise, from the scanner, can reach a strength comparable to the signal itself. Thus, extracting the underlying signal is a challenging process typically approached by applying statistical methods. The goal of this study is to investigate possibilities to recover information from the signal using topological feature vectors directly based on the raw signal without any medical pre-knowledge. The goal is to recover, the temporal development of brain activations, connectivity between these activations, and their relation to these cognitive tasks

Topological Method for fMRI Analysis

Farhan Rasheed, Linkoping University
Department of Science and Technology
LINKÖPING UNIVERSITY
Main advisor: Ingrid Hotz

Motivation & Research goals

Functional magnetic resonance imaging (fMRI) is used to measure brain activity due to tasks or stimuli. Resting-state measurements are used to provide a subject's baseline. The signal is prone to noise from various sources. Random brain activity and noise, from the scanner, can reach a strength comparable to the signal itself. Thus, extracting the underlying signal is a challenging process typically approached by applying statistical methods. The goal of this study is to investigate possibilities to recover information from the signal using topological feature vectors directly based on the raw signal without any medical pre-knowledge. The goal is to recover, the temporal development of brain activations, connectivity between these activations, and their relation to these cognitive tasks

Data and Methods

Functional MRI

fMRI measures BOLD signals that are coupled with neural activation.
Participants were scanned while solving 4 different task in two 3 min block for each task[1].

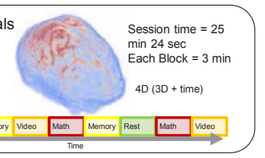

Topological feature vector – requirements:

1) represents the main characteristics of brain activity, 2) acquire directly from raw data, 3) reduce dimensionality significantly 4) has a fixed size so it is comparable across time steps.

Proposed feature vector per time step

- Activity vector derived from the merge tree[2] (Req 1,2)
- Use persistence simplification as controlled tree simplification (Req 3)

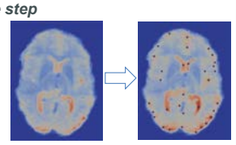

Reference brain (RF) – Feature correspondence (Req 4)

- RF combines all activity sites of one subject from all time-steps
- RF introduces a **global index set** that allows
- Morse-mapping[3] is used to establish temporal correspondence.

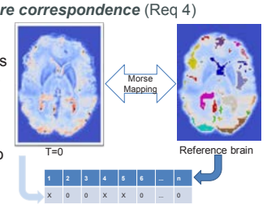

References

1. Javier Gonzalez-Castillo et al, Tracking ongoing cognition in individuals using brief, whole-brain functional connectivity patterns, *Proceedings of the National Academy of Sciences (2015)*
2. Carr et al, Computing contour trees in all dimensions, Computational Geometry (2003)
3. W. Engelke et al, Topology-Based Feature Design and Tracking for Multi-Center Cyclones, Topological Methods in Data Analysis and Visualization VI, Eds. Springer, (2021)

Selected Results

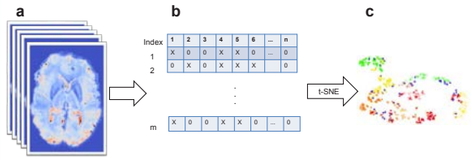

Pipeline: (a) raw data (m time steps) (b) Feature matrix representing the activity levels at the n dominant activity sites (c) 2d embedding, color coding according to the tasks, blue refers to instruction times.
Feature vectors mapping into two dimensions space reveals the different activity tasks and their transition

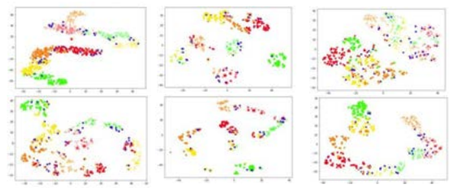

Fig: Six subjects in comparison, it can be observed that the feature vector coherency is differently strongly expressed for the subjects. For subject 1 one can observe an almost continuous change of the brain activity over time. In contrast subject 4 shows clear clusters per task.

Activation sites connectivity

For a brain connectivity analysis we focus on the correlation of individual activation sites and relate their behavior to each other. This includes analyzing the overall activation level.

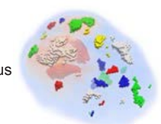
Spatial rendering of the activity sites

Chord diagram visualize the relation between the activation sites (indicated by index).The ribbon connections in chord diagram Connects regions with correlation value higher than 60%


Cord diagram

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Dual Control by Deep Reinforcement Learning using a Deep Hyperstate Transition Model

A method is proposed for performing dual control using a deep reinforcement learning algorithm in combination with a neural network model trained to represent hyperstate transitions. The method is evaluated on a simple nonlinear system suggested as a suitable benchmark for such problems, but can scale to high-dimensional systems.
The hyperstate is compactly represented as the parameters of a mixture model, which is fitted to Monte Carlo samples from the hyperstate using the Expectation Maximization algorithm. This compact representation is then used to train a hyperstate transition model, which is used by a s tandard reinforcement learning algorithm to find a dual control policy. It is demonstrated that the method is able to learn a probing technique that reduces hyperstate uncertainty, yielding improved control performance.

Dual Control by Deep Reinforcement Learning using a Deep Hyperstate Transition Model

Christian Rosdahl, Anton Cervin, Bo Bernhardsson
Department of Automatic Control, Lund University
LUND UNIVERSITY

Abstract

A method is proposed for performing dual control using a deep reinforcement learning algorithm in combination with a neural network model trained to represent hyperstate transitions. The method is evaluated on a simple nonlinear system suggested as a suitable benchmark for such problems, but can scale to high-dimensional systems.

The hyperstate is compactly represented as the parameters of a mixture model, which is fitted to Monte Carlo samples from the hyperstate using the Expectation Maximization algorithm. This compact representation is then used to train a hyperstate transition model, which is used by a standard reinforcement learning algorithm to find a dual control policy. It is demonstrated that the method is able to learn a probing technique that reduces hyperstate uncertainty, yielding improved control performance.

Example System

We try out the method on systems

$$x_{k+1} = x_k + u_k + w_k, \quad y_k = |x_k| + e_k,$$

with a linear state equation and a simple but nonlinear measurement equation. The hyperstate transition model is first tested for a case where the state x_k, input u_k and disturbances w_k and e_k are discrete-valued, since the hyperstate in this case can be computed exactly, and used for comparison with the model output. Then, we allow the signals to be continuous-valued, train a new model, and apply reinforcement learning with the estimated hyperstates, to find a control policy $\pi(\xi_k)$.

The Dual Control Problem

We consider a system with state and measurement equations

$$x_{k+1} = f(x_k, u_k, \theta_k^f, w_k),$$
$$y_k = g(x_k, \theta_k^g, e_k),$$

with parameters $\theta_k = (\theta^f_k, \theta^g_k)$ and disturbances w_k and e_k. The state x_k and parameter vector θ_k are not deterministically known, but their initial probability distribution $P(x_k, \theta_k)$ is. The goal is to determine a control policy π which minimizes some cost function J, given past inputs and outputs $\mathcal{D} = \{y_k, u_{k-1}, y_{k-1}, u_{k-2}, \ldots\}$. We define the hyperstate as

$$\xi_k = \mathbb{P}\{x_k, \theta_k \mid y_k, u_{k-1}, y_{k-1}, u_{k-2}, \ldots\},$$

i.e., the probability distribution of the state and parameter vector given past data. The goal can then be described as finding a policy $\pi(\xi_k) = \operatorname{argmin}_\pi J(\pi, \xi_k)$ that minimizes a cost

$$J(\pi, \xi_k) = \mathbb{E}\left\{ \sum_{i=k}^{k+T-1} c_i(x_{i+1}, u_i = \pi(\xi_k)) \,\middle|\, \xi_k \right\}.$$

Hyperstate Transition Model

The probability density function (PDF) for the hyperstate ξ_k is approximately represented by a mixture model

$$f(\xi) = \sum_{i=1}^{c} \lambda_k^i f_i(\xi; \beta_k^i),$$

where f_i are some simple PDFs, such as normal distributions, $\Sigma_i \lambda^i = 1$, and the distribution parameter vector

$$\alpha_k = \begin{bmatrix} \lambda_k^1 & \ldots & \lambda_k^{c-1} & (\beta_k^1)^T & \ldots & (\beta_k^c)^T \end{bmatrix}^T$$

defines the approximate PDF.

Using a system model for simulation, we collect data for how the hyperstate parametric representation α_k evolves for different inputs u_k, and the corresponding measurement signals y_{k+1}.

This data is used for training a hyperstate transition model, which yields a prediction of the new hyperstate representation α_{k+1}, given the one at the previous timestep α_k, the applied input u_k, and the measurement signal y_{k+1}. The hyperstate predictions from the model are then used as the state in a reinforcement learning procedure, which is used for determining an approximately optimal policy $\pi(\xi_k)$.

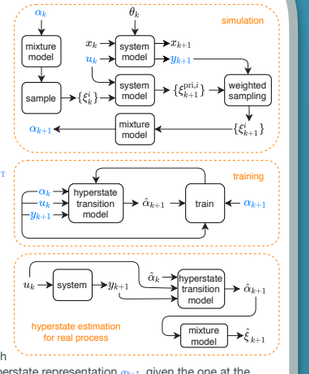

Hyperstate Transition Model Results

Some predictions (---) of the next hyperstate representation $\hat{\xi}_{k+1}$ (—) by the hyperstate transition model, given the current hyperstate ξ_k (—), the measurement y_{k+1} (---), and the input u_k for the discrete-valued case.

The real hyperstates are, for the illustrated example, discrete, and can therefore be represented exactly, as the bars in the diagrams. The vertical lines represent the real states x_k and x_{k+1}.

The evaluation cost is the Wasserstein L_1-distances between the predicted and real hyperstate representations.

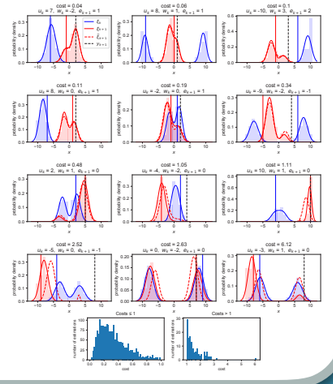

Control Results

We study a case where the system is run for episodes of length 10. Each episode is started at $x_0 = 1$. In this case, it is impossible for a controller to know whether a positive or a negative control signal should be applied to reduce state error, and small controls do not lead to large enough changes to resolve the issue, due to the measurement noise. Therefore, some exploration, by applying some larger control signal, could be useful.

Using $u_k = -x$ as input, where x is the maximum-likely state, i.e. the state corresponding to the maximum of the probability density for the approximation of the hyperstate ξ_k, does not give better performance than using $u_k = 0$. However, using the hyperstate approximation in combination with a simple reinforcement learning algorithm, to get an action-value function Q, leads to a control policy which gives a smaller cost.

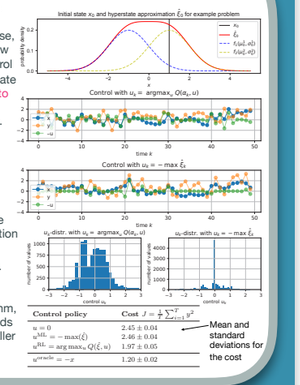

Control policy	Cost $J = \frac{1}{N}\sum_{k=1}^{N} y^2$
$u = 0$	2.45 ± 0.04
$u^{\text{ML}} = -\max(\hat{\xi})$	2.46 ± 0.04
$u^{\text{ML}} = \arg\max_u Q(\hat{\xi}, u)$	1.97 ± 0.03
$u^{\text{oracle}} = -x$	1.20 ± 0.02

Mean and standard deviations for the cost

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Fluid Models for Cloud Service Graphs

Resource management in cloud computing is a difficult problem, as one is often tasked with balancing between adequate service to clients and cost minimization in dynamic environments of many inter-connected components. To make correct decisions in these environments, good performance models are necessary. A common modeling methodology is to use networks of queues, but as these are prohibitively expensive to evaluate for many real-time applications, different approximation methods for important metrics are frequently employed. Here we build on one such method, the fluid model, to generate a time-dynamic model for mixed networks with general phase-type services times and show how these can be extracted from tracing data of a service graph.

AUTONOMOUS SYSTEMS (AS) | Page 137 B
Ruuskanen, Johan
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Fluid Models for Cloud Service Graphs

Johan Ruuskanen, Lund University
Department of Automatic Control

Modeling performance metrics in modern cloud applications

Resource management in cloud computing is a difficult problem, as one is often tasked with balancing between adequate service to clients and cost minimization in dynamic environments of many interconnected components. To make correct decisions in these environments, good performance models are necessary. A common modeling methodology is to use networks of queues, but as these are prohibitively expensive to evaluate for many real-time applications, different approximation methods for important metrics are frequently employed. Here we build on one such method, the fluid model, to generate a time-dynamic model for mixed networks with general phase-type services times and show how these can be extracted from tracing data of a service graph.

Mean-field fluid model

Below follows an illustration of the mixed multiclass network.

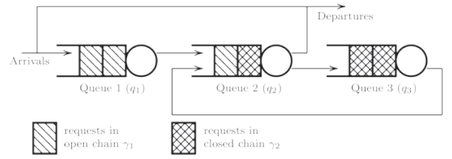

It is possible to show that such a network, assuming that the queues follow either the processor sharing and/or delay disciplines, and where the service times have a general phase-type distribution, fulfills the so-called Kurtz's Theorem. This implies that a fluid model of the mean queue lengths can be obtained via the mean-field approximation.

By stacking the parametrization matrices of the phase-type distributions into block diagonals and combining with the routing matrix ***P***, the mean-file fluid model becomes

$$\dot{\boldsymbol{x}} = \left(\boldsymbol{\Psi} + \boldsymbol{B}\boldsymbol{P}\boldsymbol{A}^T\right)^T \theta\left(\boldsymbol{x}\right) + \boldsymbol{A}\boldsymbol{\lambda}$$

Kurtz's theorem states that the queue lengths converges to the solution of this ODE as the system size scales to infinity. At lower system sizes, this model can however be inaccurate. Instead, it can be improved by using an inverse p-norm smoothing

$$\hat{\theta}_{i,r,a}(\boldsymbol{x}, \boldsymbol{p}) = \frac{x_{i,r,a}}{\left(1+\left(k_i^{-1}\sum_{j,b} x_{i,j,b}\right)^{p_i}\right)^{1/p_i}}$$

Given such a fluid model, it is further possible to retrieve a closed form approximation of the entire response time CDF over almost any subset of classes in the network

$$\Phi_{\mathcal{C}_R}\left(t \mid \beta\right) \approx 1 - \beta^T \boldsymbol{A}^T \exp\left[D^{\hat{g}(\boldsymbol{P}^*)}\boldsymbol{W}_R t\right] \mathbb{1}$$

References

- Johan Ruuskanen, Tommi Berner, Karl-Erik Årzén, Anton Cervin. *Improving the mean-field fluid model of processor sharing queueing networks for dynamic performance models in cloud computing*. Performance Evaluation. 2021
- Johan Ruuskanen, Haorui Peng, Alfred Åkesson, Lars Larsson, Maria Kihl. *FedApp: a Research Sandbox for Application Orchestration in Federated Clouds using OpenStack.* https://github.com/JohanRuuskanen/FedApp. 2021

Queuing network from trace

It is possible to retrieve a rudimentary multiclass network model directly from tracing data by assuming that each service follows the processor sharing discipline.

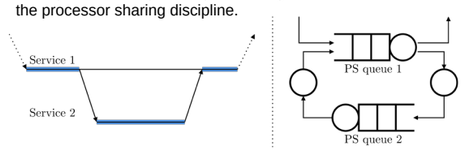

To evaluate such a model, we created the following example application performing face detection in the cloud.

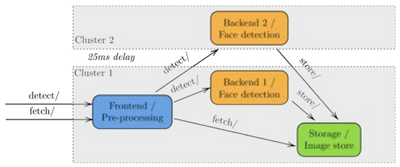

Using the sandbox FedApp, the example application was then deployed in two Kubernetes clusters with Istio to handle routing and trace logging, and intercluster delay emulated by TC Netem. The two call types where then loaded simultaneously with both open and closed connections, yielding the following results.

Arrival rate / client wait time	10 / 1	20 / 0.5	35 / 0.25	65 / 0.125
$\mathbb{E}$ [utilization]	0.19	0.36	0.59	0.84
max [utilization]	0.30	0.57	0.85	0.98

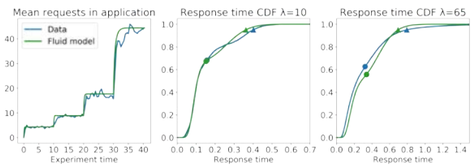

Due to modeling errors, the service time distributions will shift slightly depending on the load. In this example, we have dealt with this by refitting the distributions at each of the four operating points to demonstrate the ability to capture system metrics.

Saleh Sedghpour, Mohammad Reza
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Self-Driving Microservices

Recently, there has been a paradigm shift in software architectures from large monolithic applications into graphs of hundreds of loosely coupled microservices. The combination of this architectural transition and the DevOps movement with CI/CD has blurred the border between software development and IT operations. For IT operations, the new microservice paradigm results in a constantly evolving infrastructure landscape of software components. Ensuring of performance, reliability, and cost efficient operations in such dynamic environments is too complex for human operators, but autonomic computing mechanisms are required to make the systems increasingly manage themselves. The use of a service mesh enables outstanding observability without imposing any particular implementation costs during the development process, which suggests that it may be beneficial to develop methods for autonomous control of traffic management policies in the service mesh.

AUTONOMOUS SYSTEMS (AS) | Page 138 B

Saleh Sedghpour, Mohammad Reza
Umeå University

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Self-driving Microservices

Mohammad Reza Saleh Sedghpour, Umeå University
Department of Computing Science

Motivation

Recently, there has been a paradigm shift in software architectures from large monolithic applications into graphs of hundreds of loosely coupled microservices. The combination of this architectural transition and the DevOps movement with CI/CD has blurred the border between software development and IT operations. For IT operations, the new microservice paradigm results in a constantly evolving infrastructure landscape of software components. Ensuring of performance, reliability, and cost efficient operations in such dynamic environments is too complex for human operators, but autonomic computing mechanisms are required to make the systems increasingly manage themselves. The use of a service mesh enables outstanding observability without imposing any particular implementation costs during the development process, which suggests that it may be beneficial to develop methods for autonomous control of traffic management policies in the service mesh.

State of the Art

The herein proposed research project on autonomic self-management for service mesh clusters, extends on early efforts on service mesh. Since the incarnation of autonomic computing [1], there have been substantial efforts within autonomic management of cloud infrastructures. Selected references include surveys by Jennings and Stadler [2] and Mani and Shyam [3]. The rather recent microservice concept has been studied from an architectural perspective e.g., by Jamshidi et al. [4] and Mendonça at al. [5] who propose architecture based self-adaption. Gariga defines a taxonomy of microservice architectures and identifies several challenges for autonomic management of microservices [6]. Tofetti et al. suggest an architectural framework for microservice self-management [7].

Methodology

The management system gathers monitoring information from the whole stack. The gathered information, stored in the Data collection, consists mostly of time series data, but also discrete event data. Similarly, Actuation of management decisions may occur at different levels, e.g., circuit breaking and retry mechanism in the service mesh. In comparison with traditional resource management research in clouds, this microservice software stack gives both increased observability and additional actuators. The service meshes may be combined with distributed tracing that allows individual service invocations to be tracked across the microservice topology. This enables breakdown analysis of response times and thus greatly simplifies root cause analysis, at the expense of having to modify applications.

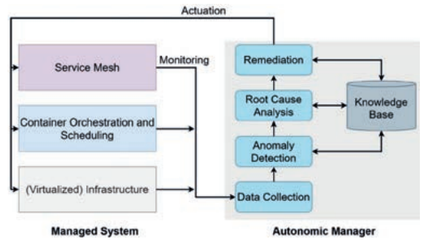

Selected Results

Preliminary results related to the proposed project include two recent papers on service mesh. In the first paper [8], we proposed a controller to manage the circuit breaker adaptively in order to maximize throughput while maintaining response time of single service in service mesh. In the second paper [9], We studied the impact of various tuning parameters for circuit breaking and retry mechanisms empirically.

References

1. Kephart, J., & Chess, D. (2003). The vision of autonomic computing. Computer, (1), 41-50
2. Jennings, B., & Stadler, R. (2015). Resource management in clouds: Survey and research challenges. Journal of Network and Systems Management, 23(3), 567-619
3. Manvi, S. S., & Shyam, G. K. (2014). Resource management for Infrastructure as a Service (IaaS) in cloud computing: A survey. Journal of network and computer applications, 41, 42444.
4. Jamshidi, P., Pahl, C., Mendonça, N. C., Lewis, J., & Tilkov, S. (2018). Microservices: The journey so far and challenges ahead. IEEE Software, 35(3), 24-35
5. Mendonça, N. C., Garlan, D., Schmerl, B., & Cámara, J. (2018). Generality vs. reusability in architecture-based self-adaptation: the case for self-adaptive microservices. European Conference on Software Architecture: (p. 18-24)
6. Garriga, M. Towards a taxonomy of microservices architectures. In International
Conference on Software Engineering and Formal Methods (pp. 203-218).
7. Toffetti, G., Brunner, S., Blöchlinger, M., Dudouet, F., & Edmonds, A. (2015, April). An
architecture for self-managing microservices. In Proceedings of the 1st International
Workshop on Automated Incident Management in Cloud (pp. 19-24). ACM
8. Saleh Sedghpour, M.R., Klein, C., Tordsson, J. (2021). Service mesh circuit breaker: From panic button to performance management tool. 1st Workshop on High Availability and Observability of Cloud Systems (HAOC '21): (p. 4-10)
9. Saleh Sedghpour, M.R., Klein, C., Tordsson, J. (2021). An empirical study of service mesh traffic management policies for microservices. Submitted to 13th ACM/SPEC International Conference on Performance Engineering.

UMEÅ UNIVERSITY

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Joint Stiction Avoidance with Null-Space Motion in Real-Time Model Predictive Control for Redundant Collaborative Robots.

Null-space motion has been used in a Franka Emika Panda robot, a redundant collaborative robot, to ensure a continuous movement of all joints during an entire trajectory execution as an approach to avoid joint stiction and allow accurate kinesthetic teaching. As is conventional for collaborative and industrial robots, the Panda robot is equipped with an internal controller, which allows to send position and velocity references directly to the robot. Therefore, null-space motion can be added directly to the velocity references, which we generate using Model Predictive Control. The observed trajectory deviation caused by discretization approximations of the Jacobian matrix when implementing null-space motion has been corrected experimentally using sensor feedback for the real-time velocity-reference recalculation and by performing a fast sampling of the null-space vector. Null-space motion has been experimentally seen to contribute to reducing the friction torque dispersion present in static joints.

AUTONOMOUS SYSTEMS (AS) | Page 139 B
Salt Ducaju, Julian M.
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Joint Stiction Avoidance with Null-Space Motion in Real-Time Model Predictive Control for Redundant Collaborative Robots [1]

LUNDS UNIVERSITET
J. M. Salt Ducaju, B. Olofsson, A. Robertsson, R. Johansson
Department of Automatic Control, LTH, Lund University, Sweden

Motivation

In kinesthetic teaching applications, the human operator should be comfortable with the physical interaction with the robot. Thus, it is important to be familiar with the force/torque required for leading the robot. Since the necessary force should not vary greatly between different human interventions, joint stiction should be avoided.

Problem Formulation

This research experimentally analyzed the use of null-space motion to avoid joint stiction in a redundant robot. By adding null-space motion to a generated trajectory reference we ensured that no joint remained still during the trajectory execution and facilitate kinesthetic teaching.

Moreover, we evaluated the use of sensor feedback from joint angular position when online recalculating a point-to-point time-constrained trajectory using Model Predictive Control (MPC) [2] to address possible undesired side-effects of the null-space motion addition.

Franka Emika robot (7-DoF collaborative robot) used in the experiments.

Null-Space Motion Addition

The null-space unitary vector, $\dot{q}_{nsu}$, must be scaled before being added to the MPC-generated reference, $\dot{q}_{MPC}$, and sent to the robot as a velocity reference, $\dot{q}_{ref}$.

$$\dot{q} = J^{\dagger}(q)\gamma + N(q)\dot{q}_a \qquad N(q) = I_7 - J^{\dagger}(q)J(q)$$

$$\dot{q}_{\text{nsu}} = \frac{N(q)\dot{q}_a}{\|N(q)\dot{q}_a\|} \qquad \dot{q}_{\text{NS}} = \dot{q}_{\text{nsu}}\alpha \sin\left(\frac{2\pi t}{T_F}\right) \qquad \dot{q}_{ref} = \dot{q}_{\text{MPC}} + \dot{q}_{\text{NS}}$$

(where the matrix $N(q)$ projects the additional arbitrary joint angular velocity, $\dot{q}_a$, into the null-space so that it is independent of the end-effector Cartesian motion)

References

[1] J. M. Salt Ducaju, B. Olofsson, A. Robertsson and R. Johansson, "Joint Stiction Avoidance with Null-Space Motion in Real-Time Model Predictive Control for Redundant Collaborative Robots," in IEEE International Conference on Robot & Human Interactive Communication (RO-MAN), 2021, pp. 307-314.
[2] M. Ghazaei Ardakani, B. Olofsson, A. Robertsson, and R. Johansson, "Model predictive control for real-time point-to-point trajectory generation", *IEEE Transactions on Automation Science and Engineering*, vol. 16, no. 2, pp. 972-983, Apr. 2019.
[3] M. Linderoth, A. Stolt, A. Robertsson, and R. Johansson, "Robotics force estimation using motor torques and modeling of low velocity friction disturbances", *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Tokyo, Japan, Nov. 3-7, 2013, pp. 3550-3556.

Results

A reduction of friction-torque dispersion has been experimentally observed thanks to the addition of null-space motion in a joint that otherwise would be static.

$$\hat{\tau}_{\text{fric}} = \hat{\tau}_{\text{ext}} - \left(M(\hat{q})\ddot{\hat{q}} + C(\hat{q},\dot{\hat{q}})\dot{\hat{q}} + g(\hat{q})\right)$$ (from the rigid-body dynamics of the robot)

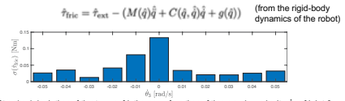

Standard deviation of the torque friction as a function of the angular velocity $\dot{\theta}_3$ of joint 3.

We investigated trajectory deviations caused by discretization approximations of the Jacobian matrix when implementing null-space motion.

END-EFFECTOR'S FINAL CARTESIAN POSITION ERRORS [MM]

	CL - NS	CL	OL - NS	OL
Mean	1.43	0.48	5.97	1.18
Std. Dev.	0.78	0.25	0.39	0.11

(*CL = closed-loop, OL = open-loop, NS = null-space*)

The effects of these approximations have been reduced by:
-Performing a fast sampling of the null-space vector.
-Using sensor feedback for the online trajectory-reference recalculation.

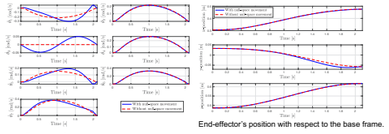

Evolution of commanded joint angular velocities.

End-effector's position with respect to the base frame.

Conclusion

Null-space motion addition can avoid joint stiction, while also avoiding the appearance of robot vibrations that may appear with dithering [3]. However, such a null-space approach is only applicable to redundant robots. Moreover, the use of sensor feedback for online trajectory recalculation has shown to address the side-effects caused by null-space motion addition.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Decentralized Multi-Agent Strategy Synthesis via Exchange of Least-Limiting Advisers

We propose a decentralized solution to a highlevel task-planning problem for a multi-agent system under a set of possibly dependent LTLf specifications. We propose an approach where the problem is turned into a number of individual two and a half player stochastic games with reachability objectives. If almost-surely winning strategies cannot be found for them, we deploy so-called least-limiting advisers to restrict agents' behaviours. A key step is treating safety and liveness separately, by synthesizing necessary safety and fairness assumptions and iteratively exchanging them in the form of advisers between the agents. We avoid the state-space explosion problem by computing advisers locally in each game, independently of the model and specification of other agents.
The solution is sound, but conservative. We demonstrate its scalability in a series of simulated scenarios involving cleaning of an office-like environment.

Decentralized Multi-Agent Strategy Synthesis via Exchange of Least-Limiting Advisers

Georg Friedrich Schuppe and Jana Tumova
KTH Royal Institute of Technology
{schuppe, tumova}@kth.se

Motivation

Heterogeneous robots in shared environment might occasionally be required to collaborate, even though they were originally not deployed to operate as a team.

Example:

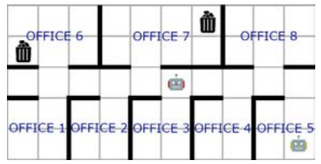

Figure 1: A partitioned office-like environment. A state of a robot is determined by its orientation and the cell it occupies. In each state, a robot can choose to stay, move forward, or turn 90°.

- Bin-emptying robots are tasked to empty k specific bins in the offices
- Cleaning robots need to clean detected spillages in certain offices and guarantee that none of the bin-emptying robots enter the affected office in order to prevent further damage

How do we express their interdependent tasks? How do we ensure that the tasks are accomplished?

Problem Formulation

- Each agent modelled an MDP $\mathcal{M}_i$

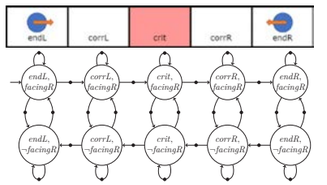

Figure 2: A small example of an MPD modeling the left robot in the corridor illustrated above. A state of a robot is determined by its orientation and the cell it occupies, the actions are to move, turn around, or stay.

- Each agent given an LTL_f specification ϕ_i

$$\longrightarrow \qquad \phi_i = \bigwedge_{\forall k \in \{1,\dots,\ell\}} Fbin_{i,k},$$

- Develop an efficient procedure to synthesize reactive strategies for all $\mathcal{M}_i$ such that all ϕ_i are satisfied, i.e. avoiding to construct a centralized Product MDP

Contributions and Approach

- A *reactive synthesis*-based approach for multi-agent high-level task planning
- A novel, decentralized approach via exchange of least-limiting advisers
- Demonstrating the scalability of the approach on selected use-cases

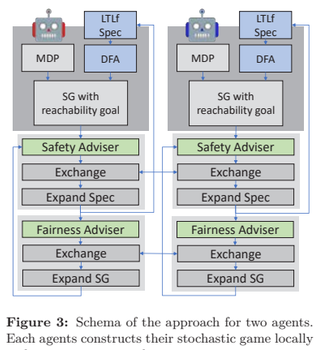

Figure 3: Schema of the approach for two agents. Each agents constructs their stochastic game locally and computes minimal, necessary assumptions on the behaviour of other agents. In an iterative process, agents incorporate advice from each other and compute additional advisers, if necessary.

Safety Advisers

Definition 1 (Minimality)**.** *A safety assumption E_s is minimal if $|E'_s| \leq |E_s|$ for all safety assumptions $E'_s \in E_2$. The unique, minimal safety assumption can be computed as*

$$E_s = \{(s,s') \in E_2 \mid s \in \langle\langle 1,2\rangle\rangle \psi \text{ and } s' \notin \langle\langle 1,2\rangle\rangle \psi\},$$

The assumption E_s cannot be directly communicated as an adviser to the other agents. Instead, we communicate the advice in the form of an adviser:

Definition 2 (Safety Adviser)**.** *A* safety adviser *is a set of tuples:*

$$SafeAdv = \{(pre, \sigma) \mid pre \in AP_i, \sigma \in \widehat{\Sigma}_i\}$$

Given that agent i satisfies pre, other agents should not satisfy σ in their next state.

Safety Advisers are implemented by expanding the specification formula of affected agents:

$$\phi_{(pre,\sigma),i} = G(pre \rightarrow \neg X proj_i(\sigma)),$$

We can incorporate all advisers from all agents into the specification of the agent i through conjunction:

$$\phi_{s,i} = \bigwedge_{\forall (pre,\sigma) \in SafeAdv_j, j \in N} \phi_{(pre,\sigma),i}$$

Fairness Advisers

Since computing a minimal fairness assumption is NP-hard, we compute *locally minimal* fairness assumptions instead. Similarly to safety assumptions, we transform fairness assumptions into Fairness Advisers:

Definition 3 (Fairness Adviser)**.** *A* fairness adviser *is a set of tuples:*

$$FairAdv = \{(pre, \sigma) \mid pre \in AP_i, \sigma \in \widehat{\Sigma}_i\}$$

Given that agent i satisfies pre, other agents should satisfy σ with non-zero probability in their next state.

Fairness Advisers are implemented through explicit modification of the stochastic games.

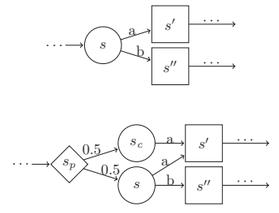

Figure 4: Enforcing fairness on (s, a, s') by preprending a probabilistic state.

Results

- The solution is sound, but conservative
- Conservativeness stems from the information gap between agents and the implementation of fairness advisers
- When dependencies between agent specifications are low, the computation time depending on the number of agents behave almost linearly
- Conveying advisers to humans poses an interesting line of future research

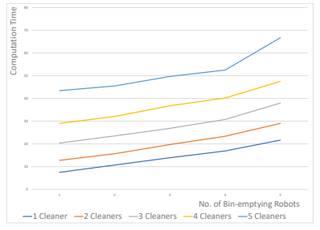

Acknowledgement

This work is partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation and the Swedish Research Council (VR) (project no. 2017-05102).

Complexity certification of Mixed-Integer Quadratic Programming

In hybrid model predictive control (MPC), a non-convex optimization problem has to be solved at each time step, which in real-time applications makes it important to solve these efficiently and to have good upper bounds on worst-case solution time. For linear hybrid MPC problems, the optimization problem is often a multi-parametric mixed-integer quadratic program (mp-MIQP) that depends on parameters such as system states and reference signals.
The aim of the research is to certify the complexity of MIQPs by computing which sequence of subproblems are required to solve in the branch and bound (B\&B) method for every parameter of interest. These sequences can be used to compute the worst-case bounds on how many iterations, floating-point operations and, ultimately, the maximum solution time, the B\&B algorithm would require to converge online.

Complexity certification of Mixed-Integer Quadratic Programming

Shamisa Shoja, Linköping University
Division of Automatic Control, Department of Electrical Engineering (ISY)
Supervisor: Daniel Axehill

LINKÖPINGS UNIVERSITET

Motivation & Research Goals

In hybrid model predictive control (MPC), a non-convex optimization problem has to be solved at each time step, which in real-time applications makes it important to solve these efficiently and to have good upper bounds on worst-case solution time. For linear hybrid MPC problems, the optimization problem is often a multi-parametric mixed-integer quadratic program (mp-MIQP) that depends on parameters such as system states and reference signals. The aim of the research is to certify the complexity of MIQPs by computing which sequence of subproblems are required to solve in the branch and bound (B&B) method for every parameter of interest. These sequences can be used to compute the worst-case bounds on how many iterations, floating-point operations and, ultimately, the maximum solution time, the B&B algorithm would require to converge online.

Methods

- **Problem Formulation**
  - mp-MIQP

$$\begin{aligned} \min_{x} \quad & \frac{1}{2}x^T H x + f^T x + \theta^T f_\theta^T x, \\ \text{s.t.} \quad & Ax \leq b + W\theta, \\ & x_i \in \{0,1\}, \quad \forall i \in \mathcal{B} \end{aligned}$$

  - $x = [x_c^T, x_b^T]^T \in R^{n_c} \times \{0,1\}^{n_b}$: state vector
  - $\theta \in \Theta_0 \subset R^{n_\theta}$: parameter vector

- **B&B method**

  Solving a sequence of relaxed convex mp-QP problems by fixing a binary variable to 0 and 1, forming nodes in the B&B search tree and cut a node if the solution of a relaxation is
  - *infeasible*
  - *does not provide better solution*
  - *integer feasible*

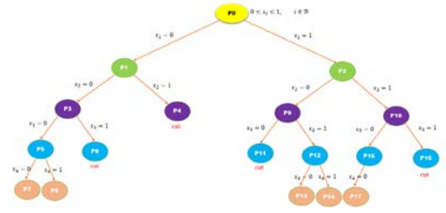

Contribution:
An algorithm for computing a useful upper bound of the worst-case computational complexity for solving any possible MIQP that can arise from a specific parameter in a polyhedral parameter ser

References

[1] A parametric branch and bound approach to suboptimal explicit hybrid MPC
Daniel Axehill, Thomas Besselmann, Davide Martino Raimondo, Manfred Morari
Automatica,2014

[2] A unifying complexity certification framework for active-set methods for convex quadratic programming
Daniel Arnström, Daniel Axehill
IEEE Transactions on Automatic Control,2022

[3] Integer programming
Laurence A. Wolsey
John Wiley & Sons, 2020

Results

Partitioning the parameter space based on the total number of QP iterations, i.e., the total number of linear system of equations solved, for a random example with with $n_c = 2$, $n_b = 4$, $n_\theta = 2$, from the proposed certification algorithm. Points with the same color share the same number of complexity measure.

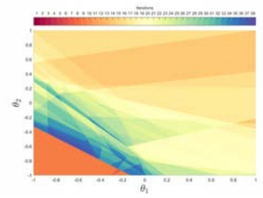

The total QP iteration number for 10000 samples specified by * in the parameter space derived by applying online B&B to the same example.

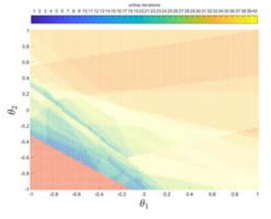

The complexity certification result coincides with the online algorithm in all sample points, despite that the conservative upper bound is used in the certification method, using depth-first search strategy

Ongoing & Future works:

- Exact complexity certification of mixed-integer linear programming (MILP) (*ongoing*)
- Complexity certify the B&B method for different node selection strategies such as best-first strategy
- Certification of the warm-started algorithm to decrease the computational complexity

WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Critical Scenario Identification for Realistic Testing of Autonomous Driving Systems

Testing of autonomous vehicles involves enormous challenges for the automotive industry. The number of real-world driving scenarios is extremely large, and choosing effective test scenarios is essential, as well as combining simulated and real-world testing. We present an industrial workbench of tools and workflows to generate critical test scenarios for active safety and autonomous driving functions in an efficient way. The workbench is based on existing engineering tools and helps smoothly integrate simulated testing, with real vehicle parameters and software. We validate the workbench with real autonomous driving systems and demonstrate its effectiveness for the realistic testing of such systems.

Software Testing of Autonomous Systems

Qunying Song, Lund University
Department of Computer Science
LUNDS UNIVERSITET

Concepts in Testing of Autonomous Systems

Testing of autonomous systems is extremely important as many of them are both safety-critical and mission-critical, yet it is still an open challenge on how to test such systems effectively and efficiently. To gain a better understanding of autonomous systems practice and facilitate testing of different autonomous systems, we conduct an exploratory study [1] by synthesizing existing academic literature with a focus group discussion and interviews with industry practitioners. As a result, we present a conceptualization of autonomous systems, classifications of challenges and current practices as well as of available techniques and approaches for testing of autonomous systems.

Critical Scenario Identification for Realistic Testing of Autonomous Driving Systems

The number of real-world operational scenarios for autonomous systems is extremely large, and choosing effective test scenarios is essential, as well as combining simulated testing and real-world testing. We focus on a common area within autonomous systems – autonomous vehicles and establish an industrial workbench to generate efficient and effective scenarios for testing such systems [2]. The workbench consists three existing engineering tools and a workflow, and helps smoothly integrate simulated testing, with real vehicle parameters and software. We also demonstrate the effectiveness of the workbench by using two real autonomous driving systems from industry by collaborating with Volvo Cars.

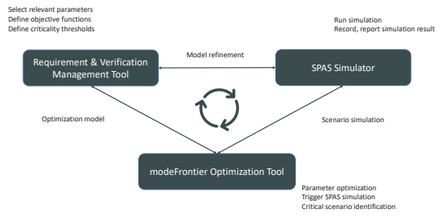

Figure-1. Tools and workflow involved in the workbench for critical scenario identification for autonomous driving systems

A Vehicle-Pedestrian Time-To-Collision Model for Testing of Autonomous Driving Systems

While autonomous driving systems are expected to reduce road accidents and improve traffic safety, understanding of the intensive and complex traffic situations is prerequisite to enable testing of such systems in a realistic traffic setup. We propose a model that predicts the worst-case distribution of TTC (Time-to-Collision) for vehicle-pedestrian interactions at unsignalized crossings, based on the traffic density. We validate the model using real traffic data collected in Sweden. We also demonstrate its use for testing of autonomous driving systems by connecting the model to critical test scenario identification for an autonomous emergency braking function from the industry.

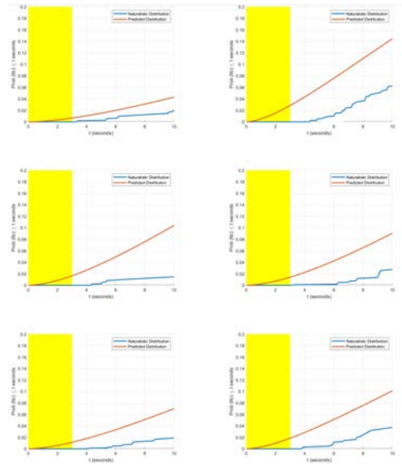

Figure-2. Model validation with naturalistic traffic data collected by Viscando in Linköping, Sweden

References

1. Song, Qunying, Emelie Engström, and Per Runeson. "Concepts in Testing of Autonomous Systems: Academic Literature and Industry Practice." In *WAIN'21 1st Workshop on AI Engineering*. IEEE Computer Society, 2021.
2. Song, Qunying, Kaige Tan, Per Runeson, and Stefan Persson. "An Industrial Workbench for Test Scenario Identification for Autonomous Driving Software." In *2021 IEEE International Conference on Artificial Intelligence Testing (AITest)*, pp. 81-82. IEEE, 2021.

Svahn, Caroline
Linköping University / Ericsson AB

WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Bayesian Prediction with Covariates Subject to Detection Limits

Missing values in covariates due to censoring by signal interference or lack of sensitivity in the measuring devices are common in industrial problems. We propose a full Bayesian solution to the prediction problem with an efficient Markov Chain Monte Carlo (MCMC) algorithm that updates all the censored covariate values jointly in a random scan Gibbs sampler. We show that the joint updating of missing covariate values can be at least two orders of magnitude more efficient than univariate updating. This increased efficiency is shown to be crucial for quickly learning the missing covariate values and their uncertainty in a real-time decision making context, in particular when there is substantial correlation in the posterior for the missing values. The approach is evaluated on simulated data and on data from the telecom sector. Our results show that the proposed Bayesian imputation gives substantially more accurate predictions than naïve imputation, and that the use of auxiliary variables in the imputation gives additional predictive power.

Svahn, Caroline
Linköping University / Ericsson AB

WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Bayesian Prediction with Covariates Subject to Detection Limits

Caroline Svahn[†,‡], Mattias Villani[†,§]

[†]IDA, Linköping University: {firstname.lastname}@liu.se
[‡]Ericsson Research, Linköping: {firstname.lastname}@ericsson.com
[§]Dept of Statistics, Stockholm University: {firstname.lastname}@stat.su.se

LINKÖPING UNIVERSITY

WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Description

Missing values in covariates due to censoring by signal interference or lack of sensitivity in the measuring devices are common in industrial problems. We propose a full Bayesian solution to the prediction problem with an efficient Markov Chain Monte Carlo (MCMC) algorithm that updates all the censored covariate values jointly in a random scan Gibbs sampler. We show that the joint updating of missing covariate values can be at least two orders of magnitude more efficient than univariate updating. This increased efficiency is shown to be crucial for quickly learning the missing covariate values and their uncertainty in a real-time decision making context, in particular when there is substantial correlation in the posterior for the missing values. The approach is evaluated on simulated data and on data from the telecom sector. Our results show that the proposed Bayesian imputation gives substantially more accurate predictions than naïve imputation, and that the use of auxiliary variables in the imputation gives additional predictive power.

Background & Motivation

While frequentist approaches generally have the advantage of being relatively fast, Bayesian methods can quantify the uncertainty for both parameters and predictions in a way that is directly usable for decision making under uncertainty. This is clearly crucial in safety critical scenarios where faulty decisions have severe consequences, but also in less dramatic but often occurring decisions, such as in wireless telecommunications where a faulty decision may disconnect users from the network [1]. The existing Bayesian literature use Gibbs sampling algorithms to simulate from the joint posterior of the model parameters and the missing values. The proposed Gibbs samplers update the missing covariate values in an observation conditional on all other missing values, see e.g. [2] and [3]. This can be highly inefficient when the missing values are highly correlated in the posterior.

Methods & Results

The complete model can be written

$$y_i = \beta_0 + \tilde{\boldsymbol{\beta}}^\top \boldsymbol{x_i} + \varepsilon_i$$
$$\boldsymbol{x_i} = \boldsymbol{\Gamma}^\top \boldsymbol{w_i} + \boldsymbol{v_i},$$

with $\varepsilon_i \overset{iid}{\sim} \mathcal{N}(0, \sigma^2)$, $\boldsymbol{v_i} \overset{iid}{\sim} \mathcal{N}(0, \boldsymbol{\Omega})$ and where x_{ij} is unobserved if $x_{ij} < c_{ij}$.

We take a Bayesian approach and assume the following prior with independence between the blocks of parameters

$$\beta_0 \sim \mathcal{N}(0, \tau^2_{\beta_0})$$
$$\tilde{\boldsymbol{\beta}} \sim \mathcal{N}(0, \tau^2_\beta \boldsymbol{I_p})$$
$$\sigma^2 \sim \mathrm{IG}(a, b)$$
$$\boldsymbol{\gamma}|\boldsymbol{\Omega} \sim \mathcal{N}(0, \boldsymbol{\Omega} \otimes \tau^2_\gamma \boldsymbol{I_r})$$
$$\boldsymbol{\Omega} \sim \mathrm{IW}(\boldsymbol{A}, \kappa),$$

where $\gamma = \mathrm{vec}\,\boldsymbol{\Gamma}$ stacks the columns of $\boldsymbol{\Gamma}$ in a vector, $\mathrm{IG}(a, b)$ is the inverse Gamma distribution and $\mathrm{IW}(\boldsymbol{A}, \kappa)$ is the inverse Wishart distribution with κ degrees of freedom.

The missing values in a given observation, $\mathbf{x}_i^{(\mathrm{m})}$, are conditionally independent of the missing values in all other observations. Each $\mathbf{x}_i^{(\mathrm{m})}$ vector can therefore be drawn in parallel from truncated multivariate normal distributions.

	Quantiles				
Dataset	0%	25%	50%	75%	100%
1	0.52	1.03	1.25	6.61	43.8
2	0.51	1.04	1.29	9.08	65.6
3	0.54	1.03	1.22	6.62	61.1
4	0.48	1.03	1.24	8.00	57.7
5	0.55	1.03	1.25	10.13	141.7

Table 1: Ratios of the effective sample size (ESS) comparing joint sampling of missing values to univariate updates. Each row represents the quantiles of the ratio $\mathrm{ESS}_{\mathrm{multi}}/\mathrm{ESS}_{\mathrm{uni}}$ for a simulated dataset with $n = 1000$ observations on $p = 40$ covariates and approximately 40 % censored values.

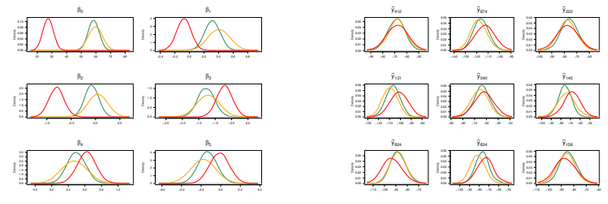

Figure 1: Posterior densities for $\boldsymbol{\beta}$ and a subset of predictive distributions. We compare our Bayesian imputation (yellow) to two baselines. The green line is an idealized model using the uncensored (complete) data, the red represent a naïve imputation strategy (red) where

$$\mathbf{x}_i^{(\mathrm{m})} = \max(\mathbf{x}_i^{(\mathrm{o})}) - \Delta,$$

meaning that all missing values are imputed to the lowest limit of detection.

Research Goal & Question

Efficient sampling of missing values is particularly important in the prediction phase where the missing covariates for a new observation must be learned quickly in a real-time context. We therefore develop a fast and efficient Markov Chain Monte Carlo (MCMC) algorithm that samples all missing covariates jointly. The joint sampling is performed using the recently proposed and highly efficient truncated multivariate normal simulation algorithm in [4] and we additionally propose a random scan implementation [5] to further increase the speed of the missing covariate updating step. The data are censored according to the following principle, which aims to mimick the censoring due to interference from the strongest signal among a set of signals:

$$x_{ij} = \begin{cases} x_{ij} & \text{if } x_{ij} \geq \max(\mathbf{x}_i^{(\mathrm{o})}) - \Delta \\ \max(\mathbf{x}_i^{(\mathrm{o})}) - \Delta & \text{otherwise,} \end{cases} \quad (1)$$

where Δ is a known distance from the strongest signal for which covariates of lower amplitude are still detectable.

Bibliography

[1] Rydén, H., Berglund, J., Isaksson, M., Cöster, R. and Gunnarsson, F., Predicting Strongest Cell on Secondary Carrier using Primary Carrier Data, IEEE Wireless Communications and Networking Conference Workshops (WCNCW): 7th International Workshop on Self-Organizing Networks (IWSON), pp. 137–142, 2018.

[2] Yue, Yu Ryan and Wang, Xiao-Feng, Bayesian inference for generalized linear mixed models with predictors subject to detection limits: An approach that leverages information from auxiliary variables, Statistics in medicine, vol. 35(10), 2016, pp. 1689–1705.

[3] Wu, Huiyun and Chen, Qingxia and Ware, Lorraine and Koyama, Tatsuki, A Bayesian Approach for Generalized Linear Models with Explanatory Biomarker Measurement Variables Subject to Detection Limit - an Application to Acute Lung Injury, Journal of applied statistics, vol. 39, 2012, pp. 1733–1747.

[4] Botev, Zdravko I, The normal law under linear restrictions: simulation and estimation via minimax tilting, Journal of the Royal Statistical Society, B, vol. 79, 2016, pp. 1–24.

[5] Amit, Yali and Grenander, Ulf, Comparing sweep strategies for stochastic relaxation, Journal of multivariate analysis, vol. 39(2), 1991, pp. 197–222.

Robust Control of Large-scale Networks

Networks in various domains such as district heating, power systems and transportation have grown increasingly complex over the years. It is therefore of interest to control the system in a manner which scales well with the network size. In this project, we are particularly concerned with scalable means of suppressing the influence of disturbances on the desired output. A recent result features a distributed controller which optimally suppresses worst-case disturbances for a class of nonlinear systems.

AUTONOMOUS SYSTEMS (AS) | Page 144 B
Vladu, Emil
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Robust Control of Large-scale Networks

Emil Vladu, Lund University
Dept. of Automatic Control
Main supervisor: Anders Rantzer

LUND UNIVERSITY

Motivation & Research Goals

Networks in various domains such as district heating, power systems and transportation have grown increasingly complex over the years. It is therefore of interest to control the system in a manner which scales well with the network size. In this project, we are particularly concerned with developing scalable means of suppressing the influence of disturbances on the desired output.

Methods

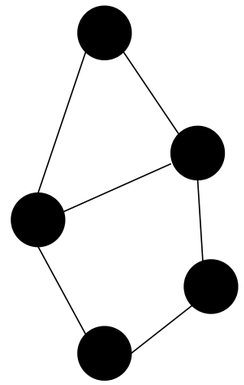

The control of networks becomes more *computationally demanding* as the system increases in size, with many nodes and interconnections as a result. For example, it may not be feasible for each controller in a network to have access to data from all nodes in the network. Such limitations can be circumvented by maintaining good control using only *local* information. In this project, we are particularly concerned with the control of *nonlinear* systems in the face of *disturbances*.

Goal: suppress the impact of disturbances on a desired output in a way suitable for large-scale systems.

We use techniques from various areas both within and outside of control, including:

1. Robust Control
2. Nonlinear Control
3. Optimization
4. Positive Systems
5. Network Dynamics

References

1. E. Vladu, C. Bergeling and A. Rantzer, *Global Solution to an H-infinity Control Problem with Input Nonlinearity*, **CDC**, 2021

Recent Result

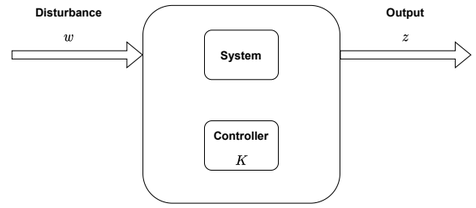

Problem: find controller K to minimize the worst-case disturbance impact on actuator effort and desired output

$$\min_K \max_w \frac{||z||}{||w||}$$

Systems considered: nonlinear systems with state-dependent input matrix and constant, stable and symmetric state matrix, where the actuator influence is the weakest at the desired output.

Result[1]: closed-form expression for an optimal controller with potential for scalability.

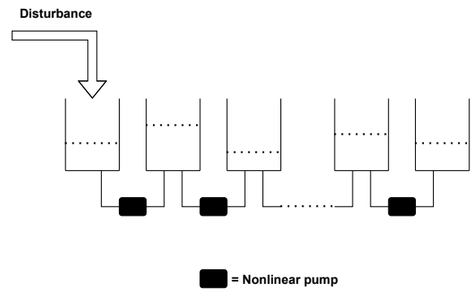

Application: consider a network of buffers with some contents, e.g. water, connected by flow links.

Goal: find a controller to minimize worst-case disturbance impact on pump effort and deviations from a desired water level.

Complication: nonlinear pumps.

Application of the result[1] gives: an optimal controller such that each pump only requires info from its neighboring buffers.
$\longrightarrow$ Useful for large-scale networks!

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Exploring autonomous USVs - Planning and Manoeuvring

The usage of USVs in a search-and-rescue scenario may be to assist in searching an area, either as a stand-alone vessel or in collaboration with other vessels. This collaboration calls for knowledge sharing and situational awareness together with automated planning to cover the search area in an efficient way between possibly heterogeneous agents/vessels. Furthermore, the concept of obstacle avoidance at sea is important in order to perform relocations in a safe manner. This includes performing preventive actions in compliance with maritime rules and regulations (COLREGs), which means that these rules must be incorporated both into the motion planning of the search task and as reactive patterns for replanning with respect to other sea traffic entering the area.

Exploring autonomous USVs – Planning and Manoeuvring

Birgitta Wingqvist, Lund University
Department of Automatic Control

Motivation & Research goals

The usage of unmanned surface vessels, USVs, in a search-and-rescue scenario may be to assist in searching an area, either as a stand-alone vessel or in collaboration with other vessels. This collaboration calls for knowledge sharing and situational awareness together with automated planning to cover the search area in an efficient way between possibly heterogeneous agents/vessels. The goal of having unmanned vessels assisting in this scenario is to increase efficiency and to relieve the humans involved. Having the vessels collaborating on low-level tasks, leaves the operator out of continuous control and the human-machine system becomes an autonomous hybrid discrete event system. As communication bandwidth is limited, local low-level processing of information is desirable [Lager, 2021] before transmitting as well as automating the continuous control.

Figure 1. Heterogeneous vessels at WARA-PS 2020.

Furthermore, the concept of obstacle avoidance at sea is important in order to perform relocations in a safe manner. This includes performing preventive actions in compliance with the traffic rules at sea, COLREGs, which means that these rules must be incorporated both into the motion planning of the search task and as reactive patterns for replanning with respect to other sea traffic entering the area. Previous work in COLREGs-compliant trajectory planning can be found in [Bergman et. al., 2020].

Traffic rules at sea

The Convention on the International Regulations for preventing Collisions at Sea, by International Maritime Organization (IMO), COLREGs, defines the traffic rules at sea. Given that the vessel is moving along a straight path, modified manoeuvres can be suggested by using predictive control [Hagen et. al., 2018]. When another vessel is detected the situation and applicable rule is identified. The movement of the own vessel is predicted with a discrete set of modifications in speed and direction, creating a set of trajectories for evaluation. For each trajectory, costs are associated with collision risk, COLREGs compliance as well as on the control signal. The cost is to be minimized. The scenarios were tested experimentally at WARA-PS 2020, see Figure 2.

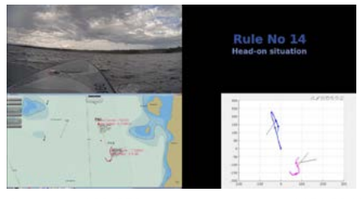

Figure 2. Initial test of COLREGs compliant manoeuvres at WARA-PS 2020 with two vessels involved. Here, a head-on situation where the vessels are both suggested to turn starboard (right). The upper left image shows the situation as seen from the magenta track in the lower right plot.

Path following

For path following, a Model Predictive Controller, MPC, with a simplified model was developed and compared with a non-model-based PID controller. The model-based controller is using the Serret-Frenet frame. Two versions of the MPC were developed, one using a non-linear model and one using a linearised version for speeding up the optimization. The PID controller is based on a Line-of-Sight, LOS, controller [Fossen, 2021] with a fixed look-ahead distance, *l*. This means that the reference heading value is the bearing to a point on the path located a distance *l* ahead on the intended path as presented in Figure 3a. Comparisons in simulation show that the path-following performance is similar under the simulated conditions.

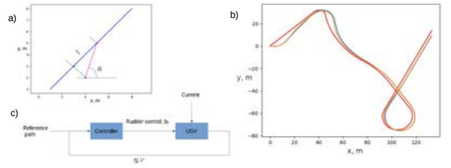

Figure 3. a) Look-ahead point. A point at a look-ahead distance, *l*, ahead is in the LOS controller the point to aim at each moment.
b) and c) Path-following controller and resulting path in simulation. Reference track in red, non-linear MPC in blue and MPC using a simplified linearised version in yellow. The prediction horizon is 10 s and a sea current is added as a disturbance.

References

1. Lager M., Digital Cognitive Companions for Marine Vessels: On the Path Towards Autonomous Ships, PhD Thesis, Print ISBN 978-91-7895-608-1, Lund University, 2021
2. Bergman, K., Ljungqvist, O., Linder, J., & Axehill, D., A COLREGs-Compliant Motion Planner for Autonomous Maneuvering of Marine Vessels in Complex Environments. arXiv preprint arXiv:2012.12145, 2020
3. Hagen I. B., Kufoalor D. K. M., Brekke E. F., Johansen T. A., MPC-based Collision Avoidance for Existing Marine Vessel Guidance Systems; IEEE International Conference on Robotics and Automation (ICRA), May 21-25, Brisbane, Australia, 2018
4.Fossen, T. I., Handbook of Marine Craft Hydrodynamics and Motion Control, Wiley, Hoboken, NJ, 2nd edition, 2021

Xie, Yiping
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Bathymetry Reconstruction from Sidescan Sonar

In recent years there has been increasing interest in the use of sidescan to reconstruct bathymetric maps. Bathymetric maps are usually constructed with high-end multibeam echo sounders (MBES), which are normally mounted on survey vessels or large autonomous underwater vehicles (AUVs). However, such MBES are relatively large and expensive compared to sidescan sonars, thus not suitable for smaller AUVs. Furthermore, sidescans generally have a wider swath range than multibeam and can produce images with a much higher resolution. If information about the seafloor's slope changes can be inferred from sidescan images, a low-cost and efficient method to construct high-resolution bathymetric maps would result and be of great benefit to many applications using smaller AUVs.

Xie, Yiping
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Bathymetry Reconstruction from Sidescan Sonar

Yiping Xie, KTH Royal Institute of Technology
Robotics Perception and Learning Lab
Main supervisor: John Folkesson

Motivation & Research Goals

In recent years there has been increasing interest in the use of sidescan to reconstruct bathymetric maps. Bathymetric maps are usually constructed with high-end multibeam echo sounders (MBES), which are normally mounted on survey vessels or large autonomous underwater vehicles (AUVs). However, such MBES are relatively large and expensive compared to sidescan sonars, thus not suitable for smaller AUVs. Furthermore, sidescans generally have a wider swath range than multibeam and can produce images with a much higher resolution. If information about the seafloor's slope changes can be inferred from sidescan images, a low-cost and efficient method to construct high-resolution bathymetric maps would result and be of great benefit to many applications using smaller AUVs.

Methods

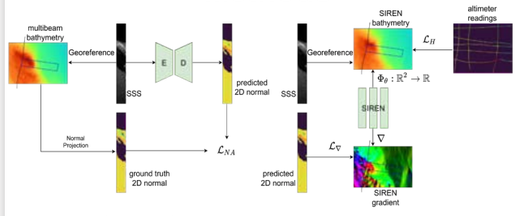

Overview: The method has two stages. Stage 1: Given the sidescan sonar (SSS) data and the mesh formed from MBES data, we can georeference the sidescan images to create training data for learning the inverse sensor model that estimates the surface normal from sidescan with a CNN. Such CNN (with an encoder "E" and a decoder "D") is trained with normal-aware loss, which can focus both the high normal area and the low normal area. Stage 2: Once the CNN is trained, we use the surface normal predicted from the CNN to constrain the gradient of the SIREN MLP, i.e., the to be estimated bathymetry. At the same time, we also use the altimeter readings to constrain the SIREN directly.

Instead of representing the bathymetry with explicit methods (meshes or grids), we use a function Φ_θ, parameterized by a fully connected neural network with parameters θ to represent the bathymetry. The fully connected neural network is a variant of MLP with sinusoidal activation functions, known as SIREN [1], mapping 2D spatial coordinates to the corresponding seafloor height, $\Phi_\theta : \mathbb{R}^2 \to \mathbb{R}$, the same as in our previous work [2]. Note that the representation is continuous, differentiable and capable of producing high-quality derivatives with respect to the 2D spatial coordinates, allowing us to supervise the derivatives during training. In addition to constraining the derivatives of the bathymetry, we also need some boundary conditions which come from the altimeter readings. Here we assume to have access to high-quality navigation data.

References

[1] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," Proc. NeurIPS, vol. 33, 2020.

[2] N. Bore and J. Folkesson, "Neural shape-from-shading for survey-scale self-consistent bathymetry from sidescan-submitted to ieee joe," IEEE J. Ocean. Eng.

Selected Results

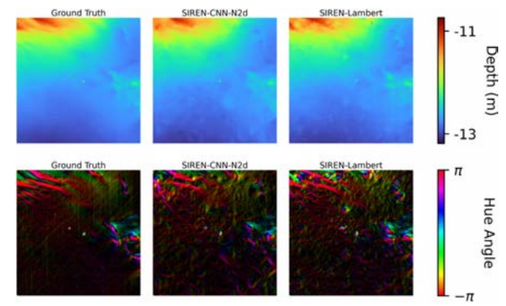

The results for the final bathymetry maps from Dataset II, another place in freshwater. Top: bathymetry maps (approximately 300m×300m); bottom: gradient. Column 1, 2 and 3: MBES Ground Truth, the SIREN trained with CNN normals from all the sidescan lines, the SIREN trained with Lambertian models with all the sidescan lines. Note that the CNN only has been trained on Dataset I, a different natural environment in seawater. The results show one application of the proposed method, that is, reconstructing survey-scale bathymetry from another place with the same ship and equipment once the CNN is trained in one place. This then shows that once trained the approach will work with only a sidescan sonar.

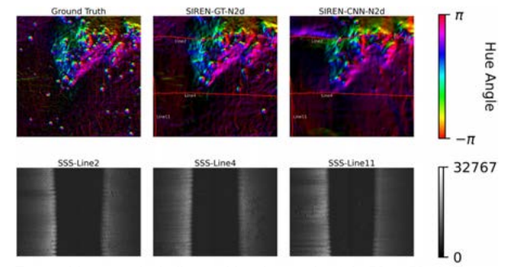

Zoomed-in areas including a ridge (approximately 50m×50m). Top: gradient of the bathymetric map in HSV; left: ground truth gradient of our bathymetry from MBES data; middle: using normal computed from the MBES data to constrain the SIREN; right: using normal predicted from the CNN to constrain the SIREN. Bottom: corresponding SSS images from different lines.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Zhu, Xiaomeng
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Automatic quality inspection based on Computer Vision and Unsupervised Domain Adaptation

Deep learning-based computer vision technologies could offer a possible solution for automatic quality inspection with their outperformance. However, most deep learning methods currently implemented in production are based on supervised learning, which requires a large amount of labeled training data that is time-consuming and expensive to collect in the industry.
This research aims to solve this problem by utilizing unsupervised domain adaptation (UDA) models. The models can be trained on annotated synthetic images generated from CAD models and unannotated images captured from cameras. They achieve promising results on an industry case study of pedal car front-wheel assembly. Furthermore, since the models do not require manually annotated images, they are less time-consuming to implement in production.

AUTONOMOUS SYSTEMS (AS) | Page 147 B

Zhu, Xiaomeng
KTH

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Automatic quality inspection based on Computer Vision and Unsupervised Domain Adaptation

Xiaomeng Zhu, Scania & KTH
Robotics, Perception, and Learning

Motivation & Research goals

Deep learning-based computer vision technologies with their outperformance could offer a possible solution for automatic quality inspection. However, most deep learning methods currently implemented in production are based on supervised learning, which requires a large amount of labeled training data that is **time-consuming** and **expensive** to collect in the industry.
This research aims to solve this problem by utilizing **unsupervised domain adaptation (UDA) models**. The models can be **trained on annotated synthetic images generated from CAD models and unannotated images captured from cameras**. Since the models do not require any manually annotated images, they are less time-consuming to implement in production.

Selected Results

Research in 3D point cloud UDA

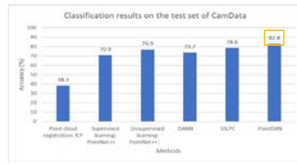

Research in 2D image UDA

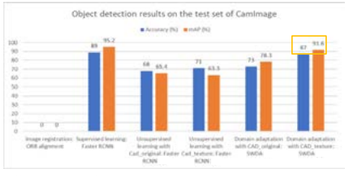

2D & 3D results comparison

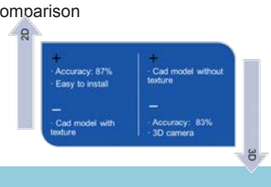

Contribution:

- End-to-end methods for automatic assembly quality inspection based on unsupervised domain adaptation with promising results.
- 2D and 3D.
- Easy to adapt to different production projects.
- Does not require does not require manual labeling work.

Method

Automatic data generation from CAD models

- Data can represent difference assembly end positions.
- Data have label.

Model implementation

- Choose an unsupervised domain adaptation model to implement.
- Comparing the unsupervised domain adaptation method with:
  - Non-deep learning method;
  - Supervised learning method;
  - Unsupervised learning method.

Research on 3D point cloud UDA [1]

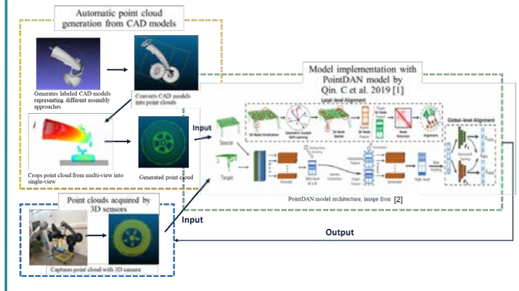

Research on 2D image UDA [3]

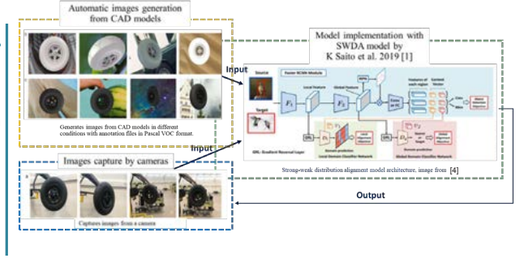

References

[1] X. Zhu, H. Manamasa, J.L. Jiménez Sánchez, A. Maki, L. Hanson, Automatic assembly quality inspection based on an unsupervised point cloud domain adaptation model, Procedia CIRP. 104 (2021) 1801–1806.
[2] C. Qin, H. You, L. Wang, C.-C.J. Kuo, Y. Fu, PointDAN: A Multi-Scale 3D Domain Adaption Network for Point Cloud Representation,ArXiv:1911.02744 [Cs]. (2019).
[3] X. Zhu, A. Maki, L. Hanson, Unsupervised domain adaptive object detection for assembly quality inspection, CIRP ICME '21 Virtual Conference, 14-16 July, Procedia CIRP, Elsevier, ISSN: 2212-8271
[4] K. Saito, Y. Ushiku, T. Harada, K. Saenko, Strong-Weak Distribution Alignment for Adaptive Object Detection, ArXiv:1812.04798 [Cs]. (2019).

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Pose Estimation from RGB Images of Highly Symmetric Objects using a Novel Multi-Pose Loss and Differential Rendering

We propose a novel multi-pose loss function to train a neural network for 6D pose estimation, using synthetic data and evaluating it on real images. Our loss is inspired by the VSD (Visible Surface Discrepancy) metric and relies on a differentiable renderer and CAD models. This novel multi-pose approach produces multiple weighted pose estimates to avoid getting stuck in local minima. Our method resolves pose ambiguities without using predefined symmetries. It is trained only on synthetic data. We test on real-world RGB images from the T-LESS dataset, containing highly symmetric objects common in industrial settings.
We show that our solution can be used to replace the codebook in a state-of-the-art approach. So far, the codebook approach has had the shortest inference time in the field. Our approach reduces inference time further while a) avoiding discretization, b) requiring a much smaller memory footprint and c) improving pose recall.

AUTONOMOUS SYSTEMS (AS) | Page 148 B
Åström, Hampus
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Pose Estimation from RGB Images of Highly Symmetric Objects using a Novel Multi-Pose Loss and Differential Rendering

Hampus Åström, Lund University, Computer Science

Adapted from paper by Stefan Hein Bengtson and Hampus Åström, Thomas B. Moeslund, Elin A. Topp, Volker Krueger - IROS 2021

We propose a novel multi-pose loss function to train a neural network for 6D pose estimation, using synthetic data and evaluating it on real images. Our loss is inspired by the VSD (Visible Surface Discrepancy) metric and relies on a differentiable renderer and CAD models. This novel multi-pose approach produces multiple weighted pose estimates to avoid getting stuck in local minima. Our method resolves pose ambiguities without using predefined symmetries. It is trained only on synthetic data. We test on real-world RGB images from the T-LESS dataset, containing highly symmetric objects common in industrial settings.

The code for our project is available at https://github.com/shbe-aau/multi-pose-estimation

Overview

We propose an adaptation of the 6D pose estimation approach in [1,2], that relies on an autoencoder for feature extraction in a codebook-based approach. By replacing their codebook with a neural network and utilizing differential rendering [3], we provide a solution that:

- has improved pose recall when tested on the T-LESS dataset.
- is faster at inference
- has a significantly smaller memory footprint

Our solution does not require discretizing poses and it is therefore easily extendable. It is trained on synthetic RGB images (no depth information required) rendered from CAD models or reconstructions and requires no labelled data or predefined symmetries.

Method

Our pose estimation method utilizes synthetic data (RGB images from CAD models) and a pre-trained encoder to train a regression network. With differential rendering we manage ambiguities with symmetries. Multiple weighted pose estimates overcome problems with local minima that stem from low output dimensionality.

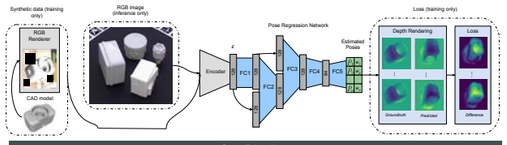

Our Pipeline

Error is measured by a difference in differentially rendered depth maps for each pose versus the ground truth. This inherently handles symmetries as it relies on physical appearance.

Each pose estimate is describes the pose as a rotation matrix. Due to the low dimensionality of the output, steepest descent get problems with local minima. Our network outputs multiple weighted pose estimates to avoid those problems.

Local Minima Visualization

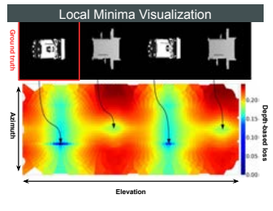

Multi-Pose Loss

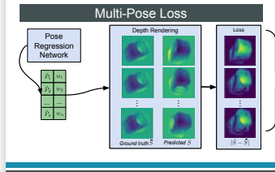

The final loss is a weighted average of the depth map differences plus a term that forces the pose estimates to be spread out.

$$= L_{\text{pose}}(\hat{P}) + \sum_{i=1}^{n} L_{\text{single}}(\hat{S}_i, \bar{S}) \cdot (\gamma + w_i)$$

Performance

On the highly symmetric dataset T-LESS [4], we achieve better recall by replacing the codebook in [1,2] with our network.

Our method operates at similar inference speed while drastically reducing the amount of GPU memory needed. We can handle cluttered scenes with occlusions as seen on the right.

The novel multi-pose loss significantly improves recall and we perform better on objects with continuous symmetries..

Inference Speed

On a GTX 1060 GPU:
6.2 ms (ours)
7.0 ms (codebook)

Memory Usage

	Encoder	Codebook	Pose Regression Network	Total
Codebook [4]	15 MB	30×45 MB	-	1365 MB
Ours		-	30×0.6 MB	33 MB

Multi-Pose Improvement

	1 pose	10 poses	improvement
Continuous symmetries	57.37 ± 1.6	**67.23** ± 2.7	9.86
Discrete symmetries	50.62 ± 0.9	**59.51** ± 0.3	8.89
All objects	53.10 ± 0.6	**62.34** ± 0.9	9.24

Recall

Obj.	Codebook	Ours	Obj.	Codebook	Ours
01	37.82	**51.84** ± 2.8	19	51.19	**54.15** ± 1.7
02	51.88	**63.74** ± 1.8	20	**40.71**	35.96 ± 1.6
03	62.87	**71.53** ± 3.3	21	43.25	**43.31** ± 1.4
04	56.00	**62.66** ± 3.5	22	**38.15**	32.03 ± 0.5
05	77.18	**80.82** ± 0.3	23	39.18	**56.68** ± 1.1
06	**68.04**	66.71 ± 4.6	24	58.97	**61.93** ± 3.3
07	65.18	**65.68** ± 4.9	25	**69.86**	63.08 ± 1.6
08	**63.11**	61.21 ± 0.8	26	57.94	**58.87** ± 2.3
09	**68.96**	55.66 ± 0.5	27	68.09	**77.62** ± 1.2
10	**58.55**	54.14 ± 2.0	28	68.06	**73.33** ± 1.3
11	**52.15**	51.48 ± 2.4	29	76.43	**80.67** ± 0.7
12	**62.19**	56.58 ± 1.6	30	77.81	**83.41** ± 2.1
13	63.56	**64.21** ± 5.0	mean	57.47	**60.09** ± 0.4
14	57.29	**63.01** ± 1.2			
15	64.91	**66.37** ± 3.8			
16	**75.82**	73.16 ± 2.7			
17	76.62	**77.72** ± 0.9			
18	**71.26**	62.71 ± 2.0	All	Codebook	Ours
mean	62.97	**63.85** ± 1.2	mean	60.77	**62.34** ± 0.9

Current and Future Work

In our paper the regression network determines the rotation from a cropped image. Translation can then be determined from the bounding box. In our current work we are extending the regression network to also determine the translation directly by providing it with bounding box information. We do this by adding additional outputs. In a similar way it could be possible to do pose predictions for flexible objects by adding additional degrees of freedom to the output.

The current version of the code has a shared autoencoder for multiple objects, but individual regression networks for each object type. Those could possibly be merged to improve training time and scalability further.

References

1. M. Sundermeyer, Z.-C. Marton, M. Durner, M. Brucker, and R. Triebel, "Implicit 3d orientation learning for 6d object detection from rgb images," in ECCV, September 2018.
2. M. Sundermeyer, M. Durner, E. Y. Puang, Z.-C. Marton, N. Vaskevicius, K. O. Arras, and R. Triebel, "Multi-path learning for object pose estimation across domains," in CVPR, June 2020.
3. N. Ravi, J. Reizenstein, D. Novotny, T. Gordon, W.-Y. Lo, J. Johnson, and G. Gkioxari, "Pytorch3d," https://github.com/facebookresearch/pytorch3d, 2020.
4. T. Hodaň, P. Haluza, Š. Obdržálek, J. Matas, M. Lourakis, and X. Zabulis, "T-LESS: An RGB-D dataset for 6D pose estimation of texture-less objects," WACV, 2017.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Batkovic, Ivo
Chalmers / Zenseact AB

Using MPC to Enable Safe Autonomous Driving

The rapid development of autonomous driving technologies in the past decades has been driven by the objectives of enabling safer and more efficient transportation. However, in order to enable such automated systems to be deployed on a global scale, problems regarding safety must be addressed. In particular, a self-driving vehicle must be able to safely interact with a surrounding environment consisting of other road users, whose intentions cannot be perfectly known. In this poster, we briefly mention how Model Predictive Control (MPC) can be used to ensure safe autonomous driving in uncertain environments.

AUTONOMOUS SYSTEMS (AS) | Page 202 B

Batkovic, Ivo
Chalmers / Zenseact AB

Using MPC to Enable Safe Autonomous Driving

Ivo Batkovic, Ind. PhD, Zenseact AB and Chalmers University of Technology
Dept. of Electrical Engineering, Mechatronics group
Supervisors: Prof. Paolo Falcone (CTH) and Dr. Mohammad Ali (Zenseact AB)

Motivation & Research Goals

In the past decade both the research community and industry have spent a vast amount of time and resources to further develop autonomous driving technologies with the objective of increasing safety and efficiency of passengers and goods transportation. However, in order to fully deploy highly automated driving functionalities, vehicles need not only to reliably sense their surrounding environment, but also *safely* interact with it.
In order to overcome such problems one has to address the question of how to design a vehicle controller that is safe by design, but also what requirements need to be set on the sensor-suite and prediction algorithms in order to enable safe autonomous driving. This poster presents an MPC-based approach to ensure safe autonomous driving in uncertain environments by slightly modifying the standard MPC controller design.

Methods

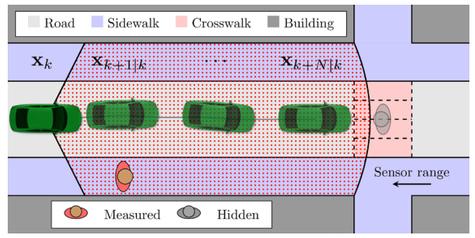

The objective is to control the autonomous vehicle, given by a nonlinear model $\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k)$, such that the a-priori *known* constraints $h(\mathbf{x}, \mathbf{u}) \leq 0$ and a-priori *unknown* constraints $g(\mathbf{x}, \mathbf{u}) \leq 0$ are satisfied. In this setting, function h can model actuator limitations or the allowed distance to the lane boundaries, and is known beforehand. The unknown constraint g on the other hand models the uncertainty and collision-avoidance w.r.t moving obstacles.
In order to ensure that the vehicle always plans a *safe* trajectory that satisfies the a-priori known and unknown constraints, we formulate the vehicle controller as the following *Model Predictive Control Problem*

$$V(\mathbf{x}) := \min_{\mathbf{x},\mathbf{u}} \sum_{n=k}^{k+N-1} q(\mathbf{x}_{n|k} - \mathbf{r}^{\mathbf{x}}_{n|k}, \mathbf{u}_{n|k} - \mathbf{r}^{\mathbf{u}}_{n|k}) + p(\mathbf{x}_{k+N|k} - \mathbf{r}^{\mathbf{x}}_{k+N|k})$$

$$\begin{aligned}
\text{subject to } & \mathbf{x}_{k|k} = \mathbf{x}_k, \\
& \mathbf{x}_{n+1|k} = f(\mathbf{x}_{n|k}, \mathbf{u}_{n|k}), && \forall n \in [k, k+M-1] \\
& h_n(\mathbf{x}_{n|k}, \mathbf{u}_{n|k}) \leq 0, && \forall n \in [k, k+M-1] \\
& g_{n|k}(\mathbf{x}_{n|k}, \mathbf{u}_{n|k}) \leq 0, && \forall n \in [k, k+M-1] \\
& \mathbf{x}_{n|k} \in \mathcal{X}^{\mathrm{s}}_{\mathbf{r}}, && \forall n \in [k, k+M-1] \\
& \mathbf{x}_{k+M|k} \in \mathcal{X}_{\mathrm{safe}},
\end{aligned}$$

where $(\mathbf{r}^{\mathbf{x}}, \mathbf{r}^{\mathbf{u}})$ is a predefined reference, $\mathcal{X}^{\mathrm{s}}_{\mathbf{r}}$ is a standard stabilizing set, and $\mathcal{X}_{\mathrm{safe}}$ is a safe set. By placing mild assumptions on the structure of the unknown constraint $g_{n|k}$ and on the existence of a *safe* terminal set $\mathcal{X}_{\mathrm{safe}}$, we show that recursive feasibility (safety) can be proven using a controller based on MPC [1,2].

Selected Results

In order to apply the theory from [1], one must first design a safe set, which essentially is a robust invariant set where the a-priori unknown constraints must be inactive. In other words, we assume the following.

> There exists a robust invariant safe set $\mathcal{X}_{\mathrm{safe}}$ such that for all $\mathbf{x}_{n|k} \in \mathcal{X}_{\mathrm{safe}}$, $h(\mathbf{x}, \mathbf{u}) \leq 0$ and $g(\mathbf{x}, \mathbf{u}) \leq 0$, and there exists a *safe* control action that entails that $f(\mathbf{x}_{n|k}, \mathbf{u}_{\mathrm{safe}}) \in \mathcal{X}_{\mathrm{safe}}$.

We deploy the safe MPC framework in a real Volvo XC90 vehicle, and assume that a suitable safe set for urban autonomous driving situations is given by a vehicle that has come to a complete stop. The motivation behind this choice is that practical settings, where safety is emphasized, typically consider a system to be safe at steady-state, i.e., a vehicle that is parked in a safe configuration (e.g., a parking lot) is not responsible for collisions with other road users.

The framework was verified in a four-way intersection at a test track, where the vehicle had to track a predefined reference (black line below) while avoiding collisions with moving pedestrians. To do so successfully, the future pedestrian motion had to be predicted (seen by the red and gray boxes) so that the a-priori unknown constraints g could be formulated. The framework demonstrated real-time capabilities, while providing a comfortable driving behavior and avoiding collisions with the moving pedestrians.

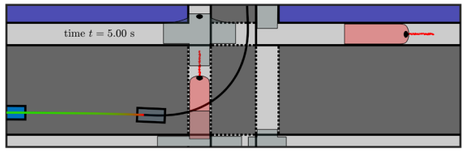

References

[1] Safe Trajectory Tracking in Uncertain Environments
Ivo Batkovic, Mohammad Ali, Paolo Falcone, and Mario Zanon
Provisionally accepted to IEEE Transactions on Automatic Control

[2] Experimental Validation of Safe MPC for Autonomous Driving in Uncertain Environments
Ivo Batkovic, Paolo Falcone, and Mario Zanon
To be submitted to IEEE Transactions on Control Systems Technology

Machine Learning Algorithms for Automatic Labelling

Obtaining labels for semi-supervised learning can be an extravagant and tedious task because of manual labeling. Because of this, industries are looking for automated solutions for data labeling. Semi-supervised machine learning algorithms (SSL) are used to automatically label datasets where few labels are available. However, it is time-consuming for practitioners in industry and academia to choose the optimal labeling algorithm for a particular problem. Therefore it is relevant to provide research that provides practitioners knowledge to choose the optimal algorithms for their specific use cases.

Machine Learning Algorithms for Automatic Labelling

CHALMERS
Teodor Fredriksson, Chalmers
Computer Science and Engineering

Motivation

- Supervised classification tasks requires labeled data.
- In industrial settings, datasets are rarely labeled.
- Data Labeling might be costly in terms of time and money.
- Automatic Labeling approaches exists but are not widely used in industry.
- Lack of research to help new practitioner's choose optimal AL approach based on their situation.

Semi-Supervised Learning

- **Semi-supervised classification algorithms learns from both labeled and unlabeled data.**
- **Assumes small amount of labeled data**

Future Research

In future research we wish to evaluate more state-of-the-art machine learning and deep learning algorithms for data labeling and evaluate them utilizing different datasets and settings based on industry.

Previous Work

- **Systematic Literature Review and Mapping Study [1], [2].**
  - In what research fields can we apply active and semi-supervised learning
  - What kind of machine learning algorithms are used?
  - What is the popularity of datatypes among the different methods?
  - What are the datasets used to evaluate these algorithms?
  - What algorithm(s) should be used for each application?
- **Case Study with Industry [3].**
  - What are the key-challenges that practitioners face in the process of labeling data?
  - What are the mitigation strategies that practitioners use to overcome these challenges?
- **Empirical Evaluation of Graph-based Semi-Supervised Learning Algorithms .**
  - Evaluates 13 different SSL algorithms on 24 different datasets divided into three datatypes, (numerical, text, images).
  - What is the ranking of the algorithms in terms of highest accuracy w.r.t aggregated data, manual effort and datatype?
- **Assessing the Sustainability of Semi-Supervised Learning Datasets using Item Response Theory [4].**
  - What datasets are suitable to compare different graph-based SSL algorithms compared.
  - How can different graph-based SSL algorithms be compared?

References

- [1] Fredriksson, T., Bosch, J. and Olsson, H.H., 2020. Machine Learning Models for Automatic Labeling: A Systematic Literature Review. In *ICSOFT* (pp. 552-561).
- [2] Fredriksson, T.A., Mattos, D.I., Bosch, J. and Olsson, H.H., 2020. Machine Learning Algorithms for Data Labeling: An Empirical Evaluation
- [3] Fredriksson, T., Mattos, D.I., Bosch, J. and Olsson, H.H., 2020, November. Data labeling: an empirical investigation into industrial challenges and mitigation strategies. In *International Conference on Product-Focused Software Process Improvement* (pp. 202-216). Springer, Cham.
- [4] Fredriksson, T., Mattos, D.I., Bosch, J. and Olsson, H.H., 2021, September. Assessing the Suitability of Semi-Supervised Learning Datasets using Item Response Theory. In *2021 47th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)* (pp. 326-333). IEEE

Decision Making for Design of Chemical Libraries

The need for data in a standardized format grows stronger within the machine learning modeling for chemistry in the pharmaceutical area. One of the popular formats for data generation are chemical libraries, which can now rapidly be designed by generative models such as RNNs. In my project I propose a method for filtering the output of the focused generative models (~10^5) down to the typical size used for library design in a lab (~10^2). This process combines the fields of generative modeling, retrosynthesis prediction, chemical property prediction and decision-making to filter the compound selection.

AUTONOMOUS SYSTEMS (AS) | Page 204 B
Johansson, Simon
Chalmers / AstraZeneca | WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Decision Making for Design of Chemical Libraries

Simon Johansson
Supervisors: Alexander Schliep, Morteza Chehreghani, Ola Engkvist
University of Gothenburg|Chalmers University of Technology|AstraZeneca
CHALMERS UNIVERSITY OF TECHNOLOGY Department of Computer science and Engineering|MolecularAI
UNIVERSITY OF GOTHENBURG

Abstract

The need for data in a standardized format grows stronger within the machine learning modeling for chemistry in the pharmaceutical area. One of the popular formats for data generation are chemical libraries, which can now rapidly be designed by generative models such as RNNs. In my project I propose a method for filtering the output of the focused generative models (~10^5) down to the typical size used for library design in a lab (~10^2). This process combines the fields of generative modeling, retrosynthesis prediction, chemical property prediction and decision-making to filter the compound selection.

Introduction

The development of strong data-driven models for chemistry in the pharmaceutical area has led to needs for more standardized data [1]. A **chemical library** is a collection of molecules synthesized under the same conditions with variations on the functional groups to represent a dense area in the chemical space.

Through **generative models** such as **LibINVENT** [2], thousands of molecules for libraries can be designed in an instant from a given core *scaffold*. By attaching **building blocks (BBs)** to this scaffold, we ensure a core similarity in the library.

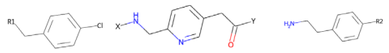

Building Block **Scaffold** **Building Block**

This generative throughput is larger than the number of molecules that can be synthesized physically and a data-driven **system for compound selection** is needed to filter the list of suggestions to a manageable library. This can be done using numerous selection criteria.

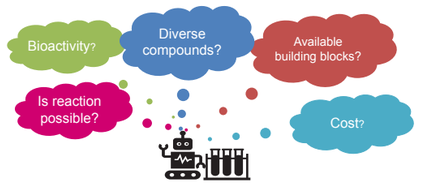

Methods

LibINVENT is trained using reinforcement learning for 1k epochs. Fragments using the amide coupling and Buchwald-Hartwig reactions to connect to the scaffold were targeted as the focus.

The saved BBs are then evaluated through Monte Carlo Tree Search (MCTS) [3] to explore possible synthesis routes. We compare the routes against a list of available stock.

A random forest model was trained to predict the activity of the generated compounds towards DRD2, and calibrated using a dataset split of 60:20:20 for training:calibration:test with the **VennABERs** predictor [4].

The selection method intened to be used is **Determninantal point processes** [5].

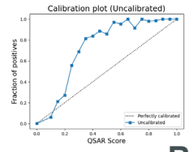

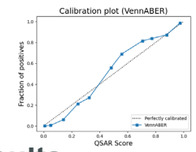

Results

After collecting all sampled suggestions, a total of 42,448 molecules had been generated. This went through a filtering process:

- BBs used in less than 5 molecules.
- BBs which could not perform the targeted 2 reactions.
- BBs that could not not be acquired within one synthesis reaction.

This yielded 100 BBs for amide coupling and 435 for Buchwald-Hartwig. However, several fragments yield the same products together with the scaffold. The number of unique products yielded were 13600, with a skew towards being active.

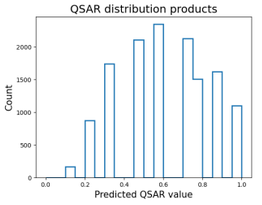

The decision-making process for this selection is still in progress.

References

1. *AI-Assisted Synthesis Prediction*, Johansson, S et al., *Drug Discovery Today: Technologies 32,* 65-72. (2019)
2. *LibINVENT: Reaction-based Generative Scaffold Decoration for* In Silico *Library Design,* Fialkova, V et al., *JCIM,* (2021) https://doi.org/10.1021/acs.jcim.1c00469
3. *AiZynthFinder: a fast, robust and flexible open-source software for retrosynthetic planning,* Genheden, S. et al, *JChemInf,* (2020), https://doi.org/10.1186/s13321-020-00472-1
4. *VennABERs predictors,* Vovk, V & Petej, I. (2014) arXiv:1211.0025
5. *The coincidence approach to stochastic point processes.* Macchi, O. (1975) *Advances in Applied Probability, 7*(1), 83-122. doi:10.2307/1425855

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Alternating Stutter Bisimulation

We want to use a fragment of Linear Temporal Logic (LTL) without the next operator to specify safety-critical requirements and synthesize a robust controller that fulfills those specifications on a discrete-time continuous state system that is subject to disturbances. Essentially, the robust controller must ensure that specific subsets of the state space are visited in an order which is allowed by the formal LTLnn specification.
However, the synthesis method cannot be applied directly on a continuous state space because it takes finite-state transition systems as input. One way to bridge this gap is to divide the continuous state space into a finite partition and let each block of the partition be one state in a transition system, which is called an abstract system. We introduce the alternating stutter bisimulation relation to be a basis for constructing the partion.

Alternating Stutter Bisimulation

Jonas Krook, Robi Malik, Sahar Mohajerani, Martin Fabian
krookj@chalmers.se, robi@waikato.ac.nz, mohajera@chalmers.se, fabian@chalmers.se

Defining a partition

For an equivalence relation $\mathcal{R} \subseteq S \times S$, the equivalence class of $s \in S$, denoted $[s]_\mathcal{R}$, is the set $\{s' \in S \mid (s, s') \in \mathcal{R}\}$. The equivalence classes of $\mathcal{R}$ form a partition of S, wherein they are referred to as blocks. We call the union of any number of equivalence classes for a superblock, and the set of all superblocks of $\mathcal{R}$ for $\mathrm{SB}(\mathcal{R})$.

A transition system is a tuple $G = \langle S, \Sigma, \delta, S^\circ, \mathrm{AP}, L\rangle$ where S is a set of states; Σ is a set of transition labels; $\delta \subseteq S \times \Sigma \times S$ is a transition relation; $S^\circ \subseteq S$ is a set of initial states; AP is a set of atomic propositions; $L\colon S \to 2^{\mathrm{AP}}$ is a state labelling function.

A path fragment of G is a sequence of states $\pi = s_1 s_2 s_3 \ldots \in S^*$ such that $(s_i, \sigma, s_{i+1}) \in \delta$ for some $\sigma \in \Sigma$ for all i.

We say that $\langle G, s_1\rangle \vDash [s_1]_\mathcal{R}\ \mathcal{U}\ T$, for $T \in \mathrm{SB}(\mathcal{R})$, if there exists an $i > 0$ for each infinite path fragment $\pi = s_1 \ldots s_i s_{i+1}$ such that $s_j \in [s_1]_\mathcal{R}$ for all $j \leq i$ and $s_{i+1} \in T$. Furthermore, $\langle G, s_1\rangle \vDash [s_1]_\mathcal{R}\ \mathcal{W}\ T$ allows also paths with $s_i \in [s_1]$ for all $i > 0$.

A controller for G is a function $C\colon S^+ \to 2^\Sigma$. A positional controller is a function $\bar{C}\colon S \to 2^\Sigma$. The transition system resulting from controlling G by C is denoted C/G.

Let $\mathcal{R}$ be an equivalence relation over S and let $(s, t) \in \mathcal{R}$. $\mathcal{R}$ is an alternating stutter bisimulation iff

(i) $L(s) = L(t)$

(ii) if, for some positional controller, $\langle \bar{C}_s/G, s\rangle \vDash [s]_\mathcal{R}\ \mathcal{U}\ T$ for some $T \in \mathrm{SB}(\mathcal{R})$, then there exists a positional controller $\langle \bar{C}_t/G, t\rangle \vDash [s]_\mathcal{R}\ \mathcal{U}\ T$

(iii) same as (ii) but with $\mathcal{W}$.

Abstraction that preserves $\mathrm{LTL}_{\setminus\circ}$ specifications under robust control

We want to use a fragment of Linear Temporal Logic (LTL) without the next operator to specify safety-critical requirements and synthesize a robust controller that fulfills those specifications on a discrete-time continuous state system that is subject to disturbances. Essentially, the robust controller must ensure that specific subsets of the state space are visited in an order which is allowed by the formal $\mathrm{LTL}_{\setminus\circ}$ specification. These subsets are represented by different color labels. For instance, we might want to implement a controller that forces the system to visit the red and green subsets infinitely often.

However, the synthesis method cannot be applied directly on a continuous state space because it takes finite-state transition systems as input. One way to bridge this gap is to divide the continuous state space into a finite partition and let each block of the partition be one state in a transition system, which is called an abstract system. We use the equivalence classes of an alternating stutter bisimulation as the partition. The transitions in the abstract transition system are then based on how the original, or concrete, dynamical system can be robustly controlled within and between the blocks.

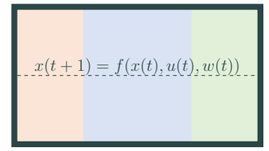

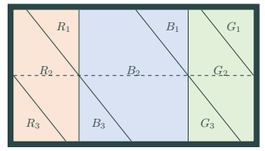

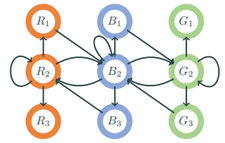

An abstract controller for the abstract system decides on control actions in the form of sets of allowed transitions, and the disturbance or process noise determines which of these allowed transitions are taken. The choices available to the abstract controller in an abstract state are based on the existence of concrete robust positional controllers that can robustly control the concrete system from the states in the corresponding source block to states in a target superblock. The blocks (equivalence classes) of alternating stutter bisimulations are defined in such a way that the system can be controlled to the same target superblocks from any of the source block's states.

For instance, if a concrete positional controller (e.g. $\bar{C}$) can control the system from one state in block B_2 to a set of states in the superblock $B_3 \cup R_2$, then an abstract controller can choose the abstract states B_3 and R_2 as the next possible states from abstract state B_2. This works since it must exist concrete controllers for any state in B_2 that can control to $B_3 \cup R_2$.

A self-loop is added to an abstract state if there exists a concrete positional controller that lets the concrete system remain in the corresponding block forever.

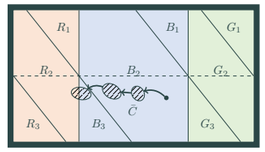

We can now synthesize a controller for the abstract system such that the red and green abstract states are visited infinitely often. Every abstract control action has a corresponding concrete controller forcing the transition, so a concrete controller fulfilling the requirement can be implemented as a sequence of concrete positional controllers.

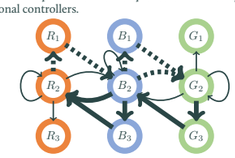

WASP WINTER CONFERENCE 2022

SOFTWARE

Evaluating Differences Between Rendered Images and Videos

In rendering research and development, it is important to have a formalized way of visualizing and communicating how and where errors occur when rendering with a given algorithm. Such evaluation is often done by comparing the test image or video to a ground-truth reference. We have presented a tool for comparing both low and high dynamic range images. Our tool is based on a perception-motivated image metric. Now, we are exploring how to extend that metric to also convey the differences in rendered videos, an extension that poses several challenges, as the presence of perceptual effects increase significantly when we consider spatiotemporal stimuli.

SOFTWARE | Page 149 B
Andersson, Pontus
Lund University / NVIDIA
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Evaluating Differences Between Rendered Images and Videos

Pontus Andersson, Lund University
Centre for Mathematical Sciences

Abstract

In rendering research and development, it is important to have a formalized way of visualizing and communicating how and where errors occur when rendering with a given algorithm. Such evaluation is often done by comparing the test image or video to a ground-truth reference. We have presented a tool for comparing both low and high dynamic range images. Our tool is based on a perception-motivated image metric. Now, we are exploring how to extend that metric to also convey the differences in rendered videos, an extension that poses several challenges, as the presence of perceptual effects increase significantly when we consider spatiotemporal stimuli.

Image Differences: ꟻLIP [1, 2]

Evaluates and visualizes the perceived differences observed when alternating between images

- Removes unperceivable details (spatial filtering)
- Perceptually uniform color space
- Enhancing edge and point errors
- Preferred viewing protocol in rendering

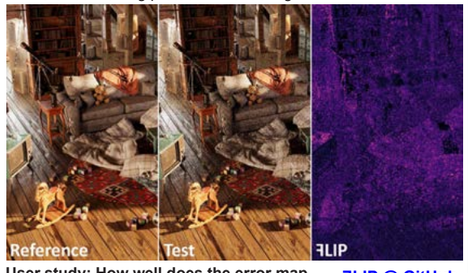

User study: How well does the error map correspond to the errors you perceive?
Results showed that ꟻLIP corresponds significantly better to the perceived error than any of the other metrics in the study.

Through an extension, ꟻLIP handles both high and low dynamic range imagery [2]

ꟻLIP @ GitHub

Video Differences: ?

Sequences of images are the main focus in rendering. Importantly, errors that are imperceptible in static images can be noticeable in videos, and vice versa.

Use cases for a video metric:

- Replace user studies and other expensive video comparisons
- Objective assessment of video generation algorithms
- Cost functions in, e.g., deep-learning-based techniques

Challenges in developing a video metric:

- *Motion*: How do we perceive errors on moving objects?
- *Flicker*: How can we determine if flickering is present?
- *Complexity*: A second of video could contain 240+ images
- *Evaluation*: Flipping not possible. What can we do instead?

Our research targets a new perception-based video metric:

- Current subproject: Temporal Edge Detection (TED):
  - Estimate the temporal edge detection filters in human vision
  - This has been done for spatial edge detection filters [3]
  - Result used to compare perceptible flicker between videos

References

1. Andersson, P., Nilsson, J., Akenine-Möller, T., Oskarsson, M., Åström, K., & Fairchild, M. D. (2020). ꟻLIP: A Difference Evaluator for Alternating Images. *Proceedings of the ACM on Computer Graphichs and Interactive Techniques*, *3(2)*, 15:1-15:23.
2. Andersson, P., Nilsson, J., Shirley, P., & Akenine-Möller, T. (2021). Visualizing Errors in Rendered High Dynamic Range Images. In *Eurographics Short Papers*.
3. Mcllhagga, W. (2018). Estimates of Edge Detection Filters in Human Vision. *Vision Research*, *153*, 30-36.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Building an data-structure selection tool for Java programs

Writing programs require using both algorithms and data-structures. Most programming languages provide implementations for some ”classical” data-structures (lists, maps, and sets). The developer chooses which implementations to use. Unfortunately, programmers are not that good at picking the data-structure that minimizes runtime.

There is existing work which uses machine learning to provide suggestions to C++ developers. Adapting this work to Java poses new challenges. In this poster, we evaluate how effective this tool is on Java programs.

SOFTWARE | Page 150 B
Couderc, Noric
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Noric Couderc, Lund University

LUNDS UNIVERSITET

I used machine learning to make Java programs faster, It didn't really work

My tool doesn't improve much over the original, exhaustive search does slightly better

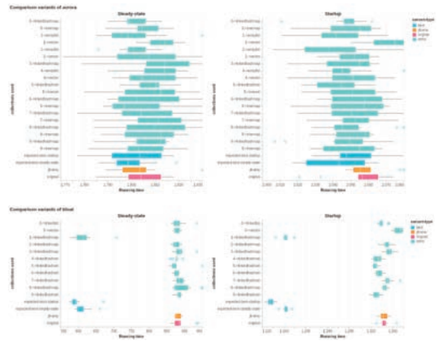

Exhaustive search finds a massive optimization, but my tool misses it

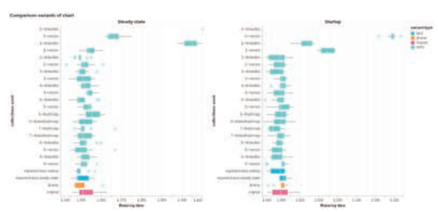

My tool doesn't improve over the original, neither does exhaustive search

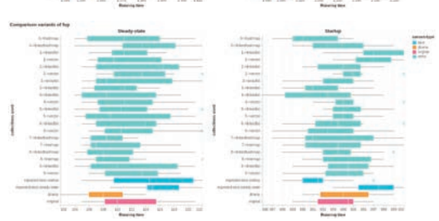

My tool does slightly better than exhaustive search on steady state performance

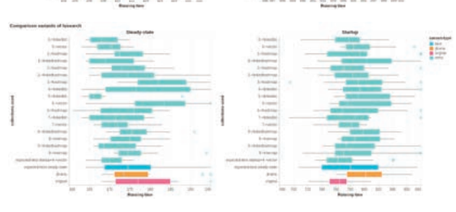

My tool does worse on startup than the original program

This is ok
It's a reproduction study!
Paper with details coming soon!

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Estimating the Potential of Program Repair Search Spaces with Commit Analysis

The most natural method for evaluating program repair systems is to run them on bug datasets, such as Defects4J. Yet, usingthis evaluation technique on arbitrary real world software projects requires heavy configuration. In this paper, we propose a newmethod, which is purely static, to evaluate the breadth of the search space of repair approaches. Our key insight is to encode thesearch spaces of repair approaches by specifying the repair strategies they employ. Next, we use the specifications to check whetheror not past commits lie in repair search spaces. For a repair approach, including many human-written past commits in its searchspace indicates its potential to generate useful patches. We implement our evaluation method in a tool, called LighteR. LighteRgets a Git repository as input and outputs a list of commits whose corresponding source code changes lie in the search spaces ofrepair approaches. Using LighteR, we conduct a study on 55,309 commits from the history of 72 Github repositories and showthat the precision and recall of LighteR are 77% and 92%, respectively. Overall, our experiments show that our novel method isboth lightweight and effective to study the search space of program repair approaches.

Estimating Program Repair Potential with Commit Analysis

Khashayar Etemadi, khaes@kth.se
Theoretical Computer Science Department @ EECS

Context: Running program repair tools is usually very costly [1]. Therefore, it is not an efficient way to assess the strength of program repair tools.
Contribution: We propose a new purely static method, which is purely static, to evaluate the breadth of the search space of repair approaches. Our key insight is to encode the search spaces of repair approaches by specifying the repair strategies they employ. Next, we use the specifications to check whether or not past human-made commits lie in repair search spaces. Our method is implemented in LighteR, with precision and recall of 77% and 92%, respectively. We find that 1.35% of 55,309 commits from 72 projects lie in search spaces of eight considered tools.

Overview of LighteR's Approach

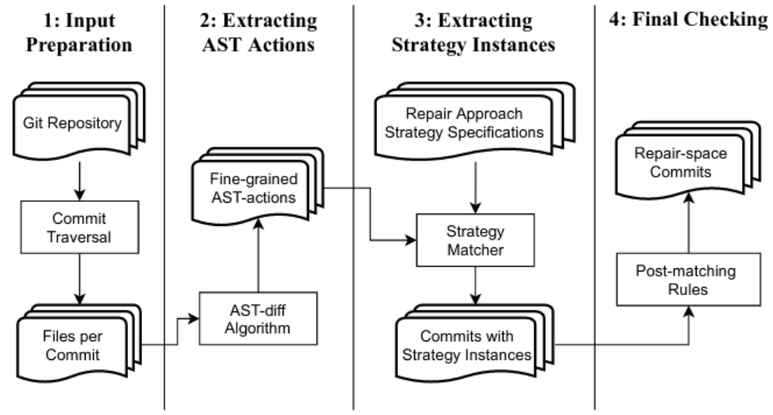

Identifying Repair-space Commits:

- Uses GumTree [2] to extract AST actions.
- Uses Coming [3] to match actions with specification.
- Uses post-matching rules to discard non-synthesizable patches.

Considered Tools:

Arja, Cardumen, Elixir, GenProg, jMutRepair, Kali, Nopol, NPEfix

Actionable Implications:

- Prototyping of New Repair Approaches by Researchers
- Evaluation of the Potential Value of Using Program Repair by Practitioners

Sample Strategy Specification for GenProg

```
<pattern name="genprog_simple">
    <entity id="1" type="*" role="Statement"/>
    <action entityId="1" type="*"/>
</pattern>
```

References

1- Durieux, Thomas, et al. "Empirical review of Java program repair tools: A large-scale experiment on 2,141 bugs and 23,551 repair attempts." Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering. 2019.
2- Falleri, Jean-Rémy, et al. "Fine-grained and accurate source code differencing." Proceedings of the 29th ACM/IEEE international conference on Automated software engineering. 2014.
3- Martinez, Matias, and Martin Monperrus. "Coming: A tool for mining change pattern instances from git commits." 2019 IEEE/ACM 41st International Conference on Software Engineering: Companion Proceedings (ICSE-Companion). IEEE, 2019.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Using Typed-Holes in Haskell for Property-Based Automatic Program Repair

This poster presents the use of typed-hole synthesis in PropR, a property-based automatic repair tool for Haskell that uses genetic programming to automatically repair Haskell programs. Haskell programs are often annotated with very specific types and come with a large suite of property-based tests in addition to unit tests. Using those properties and unit tests, we can isolate the parts of the code involved in a failing test case, and then we can leverage the available type information by integrating with the valid hole-fit synthesis in the GHC compiler to do accurate synthesis of well-typed programs as possible repairs for the fault-involved expressions.

SOFTWARE | Page 152 B
Gissurarson, Matthías Páll
Chalmers
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Using Typed-Holes in Haskell for Property-Based Automatic Program Repair

Matthías Páll Gissurarson,
Chalmers University of Technology
Department of Computer Science and Engineering

Abstract

We present the use of typed-hole synthesis in PropR, a property-based automatic repair tool for Haskell that uses genetic programming to automatically repair Haskell programs.

PropR: Property-Based Automatic Program Repair

Haskell programs are often annotated with very specific types and come with a large suite of property-based tests in addition to unit tests. Using those properties and unit tests, we can isolate the parts of the code involved in a failing test case, and then we can leverage the available type information by integrating with the valid hole-fit synthesis in the GHC compiler to do accurate synthesis of well-typed programs as possible repairs for the fault-involved expressions [1].

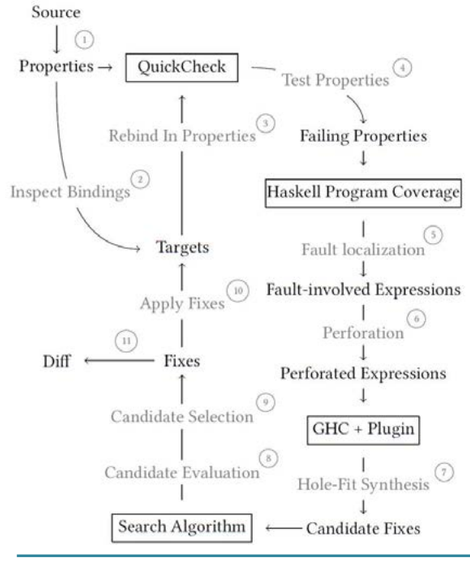

The Test-Localize-Synthesize-Rebind Loop

PropR is based on a Test-Localize-Synthesize-Rebind loop (TLSR), shown in the figure to the right [1]. We begin by parsing the source code and discovering ***properties***[1], as defined by the name and/or type of the functions in the provided source ①. By inspecting these properties ②, we determine top-level *targets* for repair. We rebind these to mutable expressions (making no changes initially, keeping the original target intact) ③. We test the properties using QuickCheck [2] to determine which ones fail ④, and use the generated counter-examples to determine the *fault-involved sub-expressions* in the targets ⑤. Then we *perforate* the targets by replacing each fault-involved sub-expression with a ***typed-hole***[2] ⑥. By using GHCs built-in ***valid hole-fit***[3] synthesis in conjunction with a hole-fit plugin, we synthesize candidate fixes in place of these holes based on their type ⑦ and mined expressions from the source code. We then evaluate the candidates by replacing the holes in the targets with candidate fixes ⑧ and apply genetic search to *select* ⑨ potential fixes based on how well the program would perform on the test-suite if they were to be applied. After selection, we apply the selected fixes to the program ⑩ by replacing sub-expressions in the targets using the selections. If all properties are now satisfied, we've found a repair, and output it as a diff ⑪, otherwise we apply the current fixes ③ to the program and repeat until we succeed or run out of search budget [1].

[1] ***Properties*** *are testable assertions in the code. They can either be unit-tests (i.e. an input-output pair) or more generic, e.g.* prop_flsPositive n = (f n) > 0, *is tested by generating random n and checking that the property holds. When QuickCheck finds a value for which the property does not hold, it* shrinks *it (i.e. minimizes it) and returns it as a counter-example [2].*

[2] *A* ***typed-hole*** *is a placeholder value for unknowns in code, and include a type inferred from the context and other constraints based on where it is placed in the source code. A typed-hole is represented by an underscore (_) in Haskell source-code.*

[3] *A* ***valid hole-fit*** *is an expression that matches the type of the hole. In GHCs, these can be synthesized by the compiler, either a simple fit (e.g.* sum *for* (_ :: [Int] -> Int)) *or a more complex* refinement hole-fit *(e.g.* (foldl (+) :: [Int] -> Int) [3].

References

1. Gissurarson, M. P., Leonhard, A., Panichella, A., Deursen, A., Sands, D. 2022. PropR: Property-Based Automatic Program Repair. To be published at the 44th International Conference on Software Engineering (ICSE 2022), Pittsburgh, PA, USA.
2. Claessen, K., & Hughes, J. 2000. QuickCheck: a lightweight tool for random testing of Haskell programs. In *Proceedings of the fifth ACM SIGPLAN international conference on Functional programming* (ICFP '00), Montreal, Canada.
3. Gissurarson, M. P. 2018. Suggesting Valid Hole Fits for Typed-Holes (Experience Report). In Proceedings of the 11th ACM SIGPLAN International Haskell Symposium (Haskell '18), St. Louis, MO, USA.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Closing the Feedback Loop in DevOps Through Autonomous Monitors in Operations

DevOps represent the tight connection between development and operations. To address challenges that arise on the borderline between development and operations, we conducted a study in collaboration with a Swedish company responsible for ticket management and sales in public transportation. The aim of our study was to explore and describe the existing DevOps environment, as well as to identify how the feedback from operations can be improved, specifically with respect to the alerts sent from system operations. Therefore, we design a solution to improve the alert management by optimizing when to raise alerts and accordingly introducing a new element in the feedback loop, a smart filter. Moreover, we implemented a prototype of the proposed solution design using a hybrid method that combines rule-based and unsupervised machine learning for operations data analysis.

Closing the Feedback Loop in DevOps Through Autonomous Monitors in Operations

Adha Hrusto, Lund University
Department of Computer Science

Abstract

DevOps represent the tight connection between development and operations. To address challenges that arise on the borderline between development and operations, we conducted a study in collaboration with a Swedish company responsible for ticket management and sales in public transportation. The aim of our study was to explore and describe the existing DevOps environment, as well as to identify how the feedback from operations can be improved, specifically with respect to the alerts sent from system operations. Therefore, we design a solution to improve the alert management by optimizing when to raise alerts and accordingly introducing a new element in the feedback loop, a smart filter. Moreover, we implemented a prototype of the proposed solution design using a hybrid method that combines rule-based and unsupervised machine learning for operations data analysis.

Research approach

Our study is a problem-driven design science approach as shown in *Figure 1*. We explored how the general problem, of incorporating feedback from **operations** in the **development**, manifests as a problem instance in the industrial context under study. For that purpose, we conducted *interviews* and performed *observations* in the case company to identify and articulate the main problems on which to focus further improvements. In the **problem conceptualization** step, we identified *three* problem instances related to *alert flooding***,** which is a phenomenon that appears in a case of a high number of alerts that are not properly managed. We provided a **conceptual design** for only one of the problem instances, alert flooding as an optimization problem since it causes the highest information overflow in the feedback loop. Moreover, alongside the proposed solution design, we implemented a **prototype instance** to get a better understanding of the opportunities of the available operations data, its type, and characteristics as well as the constraints of the context. We partially **evaluated** the implemented solution using the limited data set for implementation of the baseline anomaly detection method in a prototype environment.

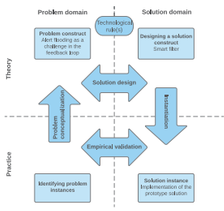

Figure 1 Overview of the Design Science approach

CASE DESCRIPTION

The system under study is a backend system of an **application** for **ticketing** and **payments** used in public transportation. It is a cloud-based system with a microservice architecture that consists of 20 services, developed using **Microsoft** tools and services. The health status of each service is monitored using the Azure Monitor through which various performance metrics and logs are available to use for alerting and visualization.

Selected Results

Problem conceptualization. We identified alert *targeting*, signal to noise *optimization*, and system *interoperability* as being three important problem instances of the general alert flooding problem in the feedback from operations to development.

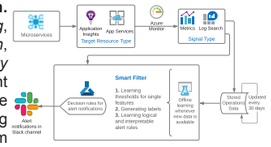

Figure 2 Overview of the proposed solution design

Solution design. We designed a solution for more effective processing of data available through the monitoring system in operations by introducing a smart filter in the feedback loop as shown in *Figure 2.* The smart filter is a unique technical solution that combines various systems' and applications' metrics for learning advanced alert rules (*see Table 1*).

Prototype implementation. We performed a pilot implementation of the proposed solution in the case environment as a proof of concept for further work. In the implementation of the prototype solution, we used unsupervised anomaly detection throughout the labeling process of unlabeled operations data while also considering the service vulnerability and observed metrics frequency (*see Table 1*). Further, for generating new advanced alert rules, a supervised tree-based machine learning technique was used.

Table 1 Overview of the selected data, service vulnerabilities and desired decision rules

Selected metrics	CPU Time	Num. of failed requests	Num. of exceptions	Num. of dep. failures	Http 4xx errors	Http 5xx errors	Num. of requests	Response time
Services with known vulnerabilities	Service B – buying tickets on vending machines		Service G – service for validating selected locations		Service M –> main service for ticketing		Service P –> bridge to an external payment service	
Example of a decision rule	**IF** num_of_failed_requests_SG > threshold_1 **AND** response_time_SB > threshold_2 **AND** num_of_Http500_SB > threshold_3 **THEN** send_notification							

Evaluation. We also implemented multivariate anomaly detection (MAD) to validate our prototype by comparing it with the pure unsupervised ML technique for detecting outliers, representing alerts, in multivariate unlabeled data set. The results revealed that the **MAD** trained model **does not scale** very well the number of predicted alerts, thus, producing the same level of noise and several alert floods. On the other hand, the **smart filter** produces **less noise** around actual failures and more accurately predicts isolated alerts in case of short system's glitches.

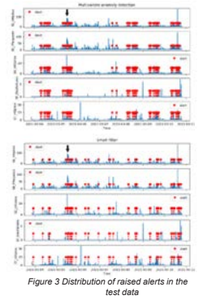

Figure 3 Distribution of raised alerts in the test data

References

1. Adha Hrusto, Per Runeson, and Emelie Engström (2021). **Closing the Feedback Loop in DevOps Through Autonomous Monitors in Operations.** SN Computer Science 2, 6 (Aug. 2021). https://doi.org/10.1007/s42979-021-00826-y

Timing-Attacks on Post-Quantum Cryptographic Primitives

Vulnerabilities found in crypto algorithms “HQC” and “BIKE”
Next-generation public-key encryption algorithms “HQC” and “BIKE”, which are code-based key encapsulation mechanisms, share a vulnerability due to the use of Rejection Sampling in their decapsulation mechanisms.
An attacker can use this vulnerability to craft special messages by which a complete secret-key recovery can be achieved.
The time-complexity of this attack is low enough to be practically managed within a couple of days, in an ideal lab scenario.

SOFTWARE | Page 154 B
Nilsson, Alexander
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Timing-Attacks on Post-Quantum Cryptographic Primitives

Alexander Nilsson, Lund University
Dept. Electrical and Information Technology, Faculty of Engineering LTH

LUND UNIVERSITY

advenica

Vulnerabilities found in crypto algorithms “HQC” and “BIKE”

Next-generation public-key encryption algorithms “HQC” and “BIKE”, which are code-based *key encapsulation mechanisms,* share a vulnerability due to the use of *Rejection Sampling* in their decapsulation mechanisms.

An attacker can use this vulnerability to craft special messages by which a complete secret-key recovery can be achieved.
The time-complexity of this attack is low enough to be practically managed within a couple of days, in an ideal lab scenario.

Post-Quantum Cryptography

Quantum computers threaten to break most of the encryption in use over the internet today. Therefore, new algorithms are needed.

NIST is currently in an open process to standardize a small number of new public-key primitives for encryption and digital signatures. Among the few remaining candidates are the code-based key encapsulation mechanisms “HQC” and “BIKE”.

KEM/PKE Scheme	Type	Vulnerable to Rejection Sampling Timing Attack
Finalists		
Classic McEliece	*Code-based*	NO
Kyber	Lattice	NO
NTRU	Lattice	NO
SABER	Lattice	NO
Alternates		
BIKE	*Code-based*	YES
FrodoKEM	Lattice	NO
HQC	*Code-based*	YES
NTRU Prime	Lattice	NO
SIKE	Supersingular elliptic curve isogeny	NO

Rejection Sampling

The Hamming Weight $hw(v)$ denotes the number of set bits in the bit vector v.

Both BIKE and HQC need to generate uniformly random error vectors e, where $hw(e) = T$, where T is a constant parameter.

Rejection sampling does this by iteratively generating random bit positions and rejecting invalid positions. This is an efficient way of ensuring uniform randomness of specific weights.

It is, however, very hard to implement it such that the run-time, or number of *samplings,* of the algorithm does not depend on its input (such as the seed for the random number generator).

Because of this the input to the rejection sampling algorithm must be comprised entirely of public values, or be derived from entirely public values.

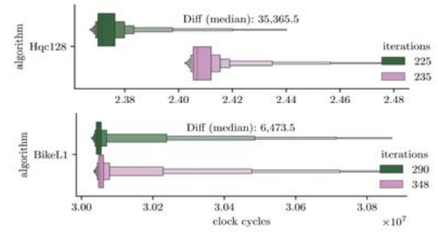

The Vulnerability [ref]

An important property in code-based schemes is the malleability of the ciphertexts. This means that it is inherently possible to slightly modify the ciphertext by a small amount and still decrypt to the very same plaintext.

This is an undesired property of highly secure KEM/PKE schemes, and the most common way to resolve the issue has been shown to open up PKE/KEM schemes to hitherto unknown timing attacks. [1]

In the current work [2] we show how to adapt the known attack to apply also to the Rejection Sampling algorithm employed by BIKE and HQC. In this work we rely on techniques from [3] and [4].

References

1. Qian Guo, Thomas Johansson, and **Alexander Nilsson**. "A key-recovery timing attack on post-quantum primitives using the Fujisaki-Okamoto transformation and its application on FrodoKEM." *Annual International Cryptology Conference*. Springer, Cham, 2020.
2. Clemens Hlauschek, Norman Lahr, Robin Leander Schröder, Qian Guo, Thomas Johansson, and **Alexander Nilsson,** "Key recovery attacks on BIKE and HQC, due to Rejection Sampling"
3. **Alexander Nilsson**, Thomas Johansson, and Paul Stankovski. "Error Amplification in Code-based Cryptography." IACR Transactions on Cryptographic Hardware and Embedded Systems (TCHES) 2019.1 (2018): 238-258.
4. Qian Guo, Thomas Johansson, and Paul Stankovski. "A key recovery attack on MDPC with CCA security using decoding errors". ASIACRYPT 2016, Springer, Heidelberg, December 2016

A Precise Framework for Source-Level Control-Flow Analysis

Static program analysis plays a fundamental role in software development and may help developers detect subtle bugs such as null pointer exceptions or security vulnerabilities. We present IntraCFG, a language-independent framework for constructing precise intraprocedural control-flow graphs (CFGs) superimposed on the Abstract Syntax Tree (AST). Source-level dataflow analysis permits easier integration with the IDEs and Cloud tools since the reports can be directly linked to the source code and do not require producing the Intermediate Representation (IR).

SOFTWARE | Page 155 B
Riouak, Idriss
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A Precise Framework for Source-Level Control-Flow Analysis

Idriss Riouak*, Christoph Reichenbach*, Görel Hedin*, and Niklas Fors
*idriss.riouak, christoph.reichenbach, gorel.hedin, and niklas.fors (@cs.lth.se)
Department of Computer Science, Lund University, Sweden
LUND UNIVERSITY
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

DESCRIPTION

Static program analysis plays a fundamental role in software development and may help developers detect subtle bugs such as null pointer exceptions or security vulnerabilities. We present IntraCFG, a language-independent framework for constructing precise intraprocedural control-flow graphs (CFGs) superimposed on the Abstract Syntax Tree (AST). Source-level dataflow analysis permits easier integration with the IDEs and Cloud tools since the reports can be directly linked to the source code and do not require producing the Intermediate Representation (IR).

OUR APPROACH

We build the CFGs on top of the AST using Reference Attribute Grammars (RAGs). Highlights of our approach:

- Handles implicit control flow
- Fully declarative specification using JastAdd2
- Overcomes the limitations of an earlier RAG framework, eliminating *misplaced* and *redundant* nodes in the constructed CFGs.

FRAMEWORK

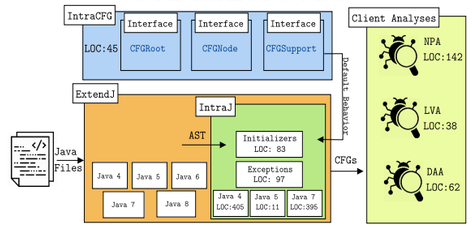

INTERFACE	ASTNODE
CFGRoot	MethodDecl, ConstructorDecl, …
CFGSupport	WhileStmt, IfStmt, …
CFGNode	All the ASTNodes that might appear in the CFGs.

The **IntraCFG** interfaces provide client APIs for the successor and predecessor relations, and default behaviour that simplifies constructing CFGs for a specific language. We used **IntraCFG** to construct high-precision CFGs for Java 7, extending the ExtendJ Java compiler.

EXPERIMENTS

We compared the results of IntraJ with:

- JastAddJ–Intraflow (JJI): a RAG based framework
- SonarQube: a highly tuned static analyser

We used as benchmarks:

- CFG size reduced by 30-40%

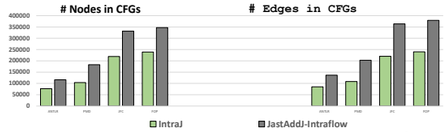

We compared the precision and the performance of IntraJ against SonarQube by implementing two dataflow analyses:

- Dead Assignment Analysis *[DAA]*
- Null Pointer Analysis *[NPA]*

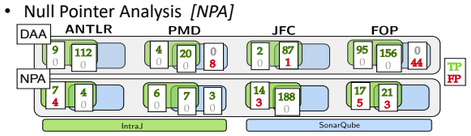

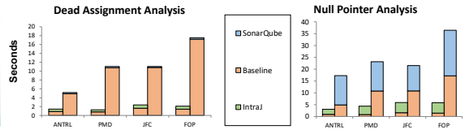

- Higher precision and better overall performance

CONCLUSIONS & FUTURE WORK

IntraCFG is a language-independent RAGs framework that overcomes the limitation of the earlier approaches:

- High-Precision
- ≥30% fewer nodes
- Concise CFG specification
- Competitive to SonarQube

We plan to:

- extend the support of IntraJ to Java 8
- extend **IntraCFG** to construct inter-procedural CFGs

LUND LTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A Domain-Specific language to express dynamic functional safety rules

Ensuring functional safety has become more challenging as robots work in a dynamic and unpredictable environment. A functionally safe autonomous system performs correctly given a set of inputs and if it receives an unknown input, then it fails predictably. While making a robot safe, sometimes we over-constrain the system that makes the robot incapable of doing anything useful. For example, turning off the robot whenever something unexpected happens is not a good recovery strategy. Specifying static safety constraints is a conservative approach. We develop a domain-specific language (DSL) that facilitates the user to specify dynamic safety specifications. A DSL allows us to reason at an abstract level making it easier to handle abstract domain-specific concerns like functional safety. Domain-specific modeling also allows us to validate if a system behaves as expected.

SOFTWARE | Page 156 B
Rizwan, Momina
Lund University
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

A ***static*** safety constraint hinders a robot to do anything ***useful***!

A Domain-Specific Language to express dynamic robot safety rules

Momina Rizwan, Christoph Reichenbach and Volker Krueger

Motivation & Research Goal

Ensuring functional safety has become more challenging as robots work in a dynamic and unpredictable environment.

"A functionally safe autonomous system performs correctly given a set of inputs and if it can't, it fails predictably."

Static safety rules over-constrain the system that makes the robot incapable of doing anything useful. Turning off the robot whenever something unexpected happens is not a good strategy. In our research, we try to replicate the work by [S.Adam et. al.] [1] and add dynamic safety rules.

Safety scenarios

1. Avoid damaging jerks while crossing uneven terrain (IMU)
2. Handle ramps in an industrial setting (IMU)
3. Handle partially closed doors (LIDAR)
4. Robot arm should never hit anything (Force Sensor)

Dynamic recovery strategy follow the rule:

"As soon as the input sensor reading is in safe-range, continue the task."

- When the ramp slope is gentle, slow down, go back and find a new path
- When the ramp slope is steep, speed up a bit so it can cross the ramp

Safety Node As a Filter

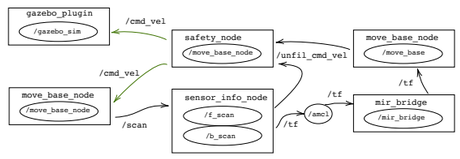

Why use a Domain Specific Language?

- To reason about domain specific concepts that are easier to understand by non-experts [2].
- One can do domain specific static analysis and report errors early.
- Better error reporting.

A code snippet (Syntax is in the style of Ulrik's work [1])

```
action moveBack ;
action lowSpeed ;
action increaseSpeed ;
const maxSpeed = 0.5 m/s
const reasonabletilt = 10 deg
const maxtilt = 20 deg
input orientation = topic
imu_information
entity imuSensorSystem
{
  gentleSlope :
    orientation.pitch() not in reasonabletilt
    for 4.0 sec;
  steepRamp :
    orientation.pitch() not in maxtilt
}
entity driveSystem {
  maxspeedExceeded :
    linearSpeed > maxSpeed for 2.0 sec ;
}
if imuSensorsystem.gentleSlope and driveSystem.moving
then { increaseSpeed ;} ;
if imuSensorsystem.steepRamp and driveSystem.moving
then { lowSpeed ; moveBack; } ;
```

*

References

[1] S. Adam, M. Larsen, K. Jensen, and U. P. Schultz, "Rule-based dynamic safety monitoring for mobile robots," *Journal of Software Engineering for Robotics*, vol. 7, no. 1, pp. 121–141, 2016.

[2] A. Nordmann, N. Hochgeschwender, and S. Wrede, "A survey on domain-specific languages in robotics," in *International conference on simulation, modeling, and programming for autonomous robots*, pp. 195–206, Springer, 2014.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Attacking and protecting GNSS receivers

Precise time and position obtained by GNSS receivers is an integral part of a wide gamut of strategic infrastructure. Demonstrations of attacks highlight the vulnerability of current civilian GNSS signals. Advanced countermeasures using external information and receiver properties can be used to detect advanced spoofer and recover from attacks. The poster describes both aspects, with a specific outlook on time focused receivers.

SOFTWARE | Page 157 B
Spanghero, Marco
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Attacking and protecting GNSS receivers

Marco Spanghero, KTH Royal Institute of Technology
Networked Systems Security (NSS) group, www.eecs.kth.se/nss
Main advisor: Panos Papadimitratos

Motivation & Research Goals

Precise time and position obtained by Global Navigation Satellite System (GNSS) receivers are an integral part of a wide gamut of strategic infrastructure. Demonstrations of attacks highlight the vulnerability of current civilian GNSS signals. Advanced countermeasures using external information and receiver properties can be used to detect advanced spoofer and recover from attacks. Specifically, integration of different time sources can tackle various spoofing scenarios, from simple cases of simulation to advanced signal lift-off.

Time based attack detection

We seek a general-purpose framework that enables validation of the GNSS time information [1]. **Fusion of multiple (secure) network time sources** [2] provides online verification, using different off-the-shelf technologies.
Local high-precision clocks can be leveraged as an ensemble to guarantee enhanced holdover under attack even when connectivity is not available, reducing the remote timing service polling frequency [3].

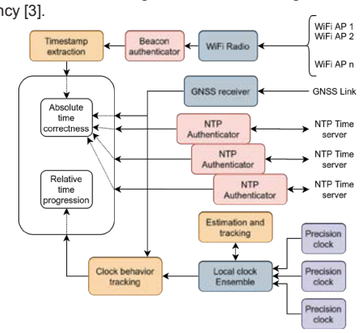

Combination of multiple time sources is not trivial and requires **knowledge of the source** characteristics: fusion can be possible with **stochastic filtering,** providing statistical indicators of GNSS time misbehavior (i.e. Kalman filtering).

Selected results:
- Detection based on opportunistic WiFi beacons with a 25µs threshold
- Multi-technology fusion framework for GNSS attack detection
- Local oscillator ensemble-based detection with 0.3µs threshold and dual frequency/phase indicator for advanced attack detection

References

[1] K. Zhang, M. Spanghero, and P. Papadimitratos, “Protecting GNSS-based Services using Time Offset Validation,” in 2020 IEEE/ION Position, Location and Navigation Symposium (PLANS), Portland, Oregon, April 2020.
[2] M. Spanghero, and P. Papadimitratos, “Detecting GNSS misbehaviour with high-precision clocks. WiSec 2021”, in Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks, Virtual.
[3] M. Spanghero, K. Zhang, and P. Papadimitratos, “Authenticated time for detecting GNSS attacks,” in Proceedings of the 33rd International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS+ 2020, Virtual
[4] M. Lenhart, M. Spanghero, and P. Papadimitratos, “Relay/replay attacks on gnss signals,” in Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks, Virtual.

GNSS Replay/Relay attacks

Replay of GNSS signals will be of increasing importance with the upcoming shift to authenticated signals [4]. Extending from simple meaconing to **over-the-network-meaconing** allows to **break free of the limitations** imposed by a physical connection.

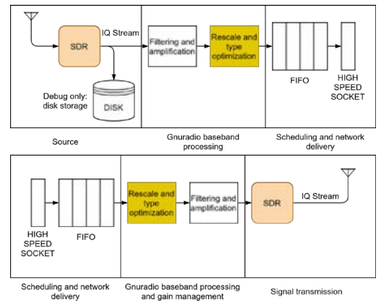

Selected results:
Meaconing of mobile and static targets was successful over **consumer 4G networks**, by either replaying the entire spectrum (bandwidth intensive) or by **surgically replaying navigation messages** using meacon and recreate strategies (future proof against authenticated navigation messages).

We developed a **future-proof, flexible and versatile GNSS testing platform**. Our modular prototype will enable research on security enhanced signals, allowing mobile scenario testing without requiring any regulatory permission.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Tests from Production Traces

Context: End users may interact with software in ways that are not well-tested.
Contribution: We propose to monitor software in production, in order to automatically improve the effectiveness of test suites. The generated tests can complement developer-written tests, and represent real usages of the application in production.

SOFTWARE | Page 158 B
Tiwari, Deepika
KTH
WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Tests From Production Traces

Deepika Tiwari deepikat@kth.se
Division of Software and Computer Systems,
School of Electrical Engineering and Computer Science

Context: End users may interact with software in ways that are not well-tested [1].
Contribution: We propose to monitor software in production, in order to automatically improve the effectiveness of test suites. The generated tests can complement developer-written tests [2], and represent real usages of the application in production [3].

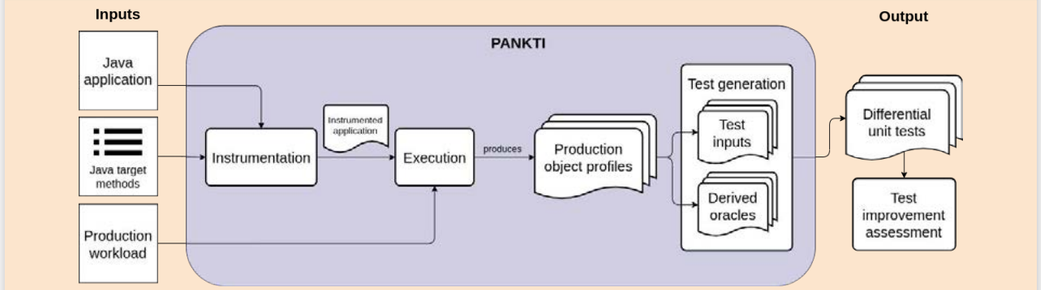

PANKTI: Automated unit test generation [2]
- Generate tests that improve the test quality of 53 weakly-tested methods in 3 open-source Java applications

AutoGraphQL: Tests from GraphQL requests [3]
- Find 8 schema faults in 1 industrial application
- Improve coverage in 1 open-source application

Future work
- Monitor the invocation of methods within target methods in production [4]
- Use the collected data for the automated generation of mocks

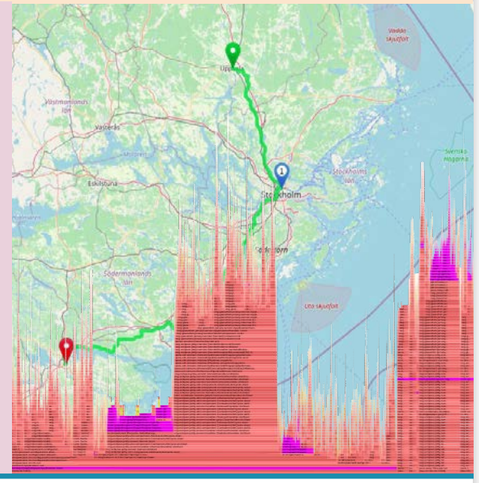

References
1. Wang, Q., Brun, Y., & Orso, A. (2017, March). Behavioral execution comparison: Are tests representative of field behavior?. In 2017 IEEE International Conference on Software Testing, Verification and Validation (ICST). IEEE.
2. Tiwari, D., Zhang, L., Monperrus, M., & Baudry, B. (2021). Production Monitoring to Improve Test Suites. IEEE Transactions on Reliability.
3. Zetterlund, L., Tiwari, D., Monperrus, M., & Baudry, B. (2022, April). Harvesting Production GraphQL Queries to Detect Schema Faults. In 2022 IEEE International Conference on Software Testing, Verification and Validation (ICST). IEEE.
4. Zhang, L., Tiwari, D., Morin, B., Baudry, B., & Monperrus, M. (2021). Automatic Observability for Dockerized Java Applications. Submitted to IEEE Transactions on Dependable and Secure Computing.

Waldemarson, Gustaf
Lund University / Arm

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Efficient GPU Programming for Visual and Autonomous Software Systems

This project aims to improve the efficiency of parallel programming of GPUs in heterogeneous architectures and environments, potentially leading to new frameworks and algorithms for, e.g., more realistic lighting in graphics systems or higher performance computing in real-time software systems. The aim is to enhance or produce programming tools that can be used to process large datasets from e.g. computer vision, deep learning, data visualization, or other graphics rendering or systems by using the potential available in modern platforms that contain GPUs and other accelerators.

SOFTWARE | Page 159 B

Waldemarson, Gustaf
Lund University / Arm

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Efficient GPU Programming for Visual and Autonomous Software Systems

Gustaf Waldemarson Ind. PhD, Arm Ltd and Lund University
Dept. of Computer Science, Lund University Graphics Group
Supervisors: Michael Doggett (LU) and Simone Pellegrini (Arm Ltd)

LUND UNIVERSITY

arm

Motivation & Research Goals

This project aims to improve the efficiency of parallel programming of GPUs in heterogeneous architectures and environments, potentially leading to new frameworks and algorithms for, e.g., more realistic lighting in graphics systems or higher performance computing in real-time software systems. The aim is to enhance or produce programming tools that can be used to process large datasets from e.g. computer vision, deep learning, data visualization, or other graphics rendering or systems by using the potential available in modern platforms that contain GPUs and other accelerators.

Methods

When devising high-performance algorithms, it is often beneficial to decouple the algorithm itself from scheduling and performance aspects. As an example, the Halide [3] framework (see below) has successfully done just this for the domain of image filters and as a part of this research we are investigating similar DSLs for generating 3D content, such as rasterization or ray tracing.

```
Func blur_3x3(Func input) {
  Func blur_x, blur_y;
  Var x, y, xi, yi;
  // The algorithm - no storage or order
  blur_x(x,y) = (input(x-1,y) + input(x,y) + input(x+1,y))/3;
  blur_y(x,y) = (blur_x(x,y-1) + blur_x(x,y) + blur_x(x,y+1))/3;
  // The schedule - defines order, locality; implies storage
  blur_y.tile(x, y, xi, yi, 256, 32)
        .vectorize(xi, 8).parallel(y);
  blur_x.compute_at(blur_y, x).vectorize(x, 8);
  return blur_y;
}
```

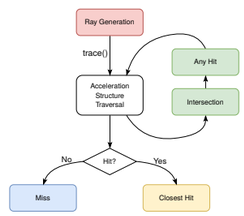

As seen above, one of the core components in ray tracing is the acceleration structure. It is the single greatest thing that determines how well ray tracing algorithms perform [2]. Thus, a lot of research has been devoted to understanding how to make this better. As a part of our research, we are investigating how to specialize this structure for particular use cases.

Physically Based Rendering

Nature is a complicated beauty with many areas that we still do not fully understand. Even when we limit ourselves to only light transport, there are many things that we are not modeling yet. Thus, another part of this research is to extend the toolkit of physical phenomena we can simulate using ray-tracing or rasterization based methods.

$$L_o(\mathbf{x}, \omega_o) = L_e(\mathbf{x}, \omega_o) + \int_\Omega f_r(\mathbf{x}, \omega_i, \omega_o) L_i(\mathbf{x}, \omega_i)(\omega_i \cdot \mathbf{n}) \, \mathrm{d}\omega_i$$

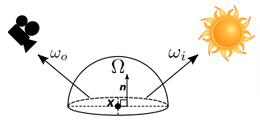

The Rendering Equation, the basis for almost all graphics algorithms.

Selected Results

One of the physical phenomena that is modeled as a part of this project is the transient light known as Cherenkov radiation. A few examples of this phenomenon can be seen below and for more details please see our work in **1**.

Recent research has been focused on creating efficient rendering algorithms for the so called *glow discharge* phenomenon, seen here.

References

[1] G. Waldemarson and M. Doggett, "Photon Mapping Superluminal Particles," in *Eurographics 2020 - Short Papers*, A. Wilkie and F. Banterle, Eds., The Eurographics Association, 2020, ISBN: 978-3-03868-101-4. DOI: 10.2312/egs.20201004

[2] D. Meister, S. Ogaki, C. Benthin, *et al.*, "A Survey on Bounding Volume Hierarchies for Ray Tracing," *Computer Graphics Forum*, 2021, ISSN: 1467-8659. DOI: 10.1111/cgf.142662

[3] J. Ragan-Kelley, A. Adams, D. Sharlet, *et al.*, "Halide: Decoupling algorithms from schedules for high-performance image processing," *Commun. ACM*, vol. 61, no. 1, pp. 106–115, Dec. 2017, ISSN: 0001-0782. DOI: 10.1145/3150211. [Online]. Available: https://doi.org/10.1145/3150211

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Application-level Chaos Engineering

Chaos engineering is a new scientific method within software engineering that consists in specifying and evaluating resilience hypotheses by 1) injecting faults in a production system, 2) observing the impact of such faults, and 3) building new knowledge about the strengths and weaknesses of the resilience of the system. Chaos engineering can be applied at different levels such as network level and infrastructure level. In order to provide more concrete and application-specific insights for developers, our research work focuses on using application-level chaos engineering to address the following challenges: C1-How to evaluate different aspects of resilience, C2-How to automate the chaos experiments, and C3-How to improve the efficiency of the chaos engineering experiments.

Application-level Chaos Engineering

Long Zhang <longz@kth.se>, KTH
Main advisor: Martin Monperrus

Abstract

Chaos engineering is a new scientific method within software engineering that consists in specifying and evaluating resilience hypotheses by 1) injecting faults in a production system, 2) observing the impact of such faults, and 3) building new knowledge about the strengths and weaknesses of the resilience of the system. Chaos engineering can be applied at different levels such as network level and infrastructure level. In order to provide more concrete and application-specific insights for developers, our research work focuses on using application-level chaos engineering to address the following challenges: **C1-How to evaluate different aspects of resilience, C2-How to automate the chaos experiments, and C3-How to improve the efficiency of the chaos engineering experiments.**

ChaosMachine

Building confidence in system behavior through EXPERIMENTS in RUNTIME

Java Virtual Machine | Java Virtual Machine | Javaagent, Java byte-code, ASM

In order to address C1, we propose to use different perturbation models at the application level [1,2]. This is because: application-level perturbation models are closer to the application's source code, which helps developers to locate the improvement target, and such models simulate more concrete failure scenarios for this specific application.

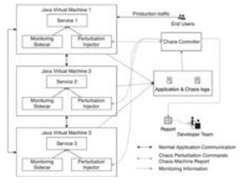

- Input
  - Arbitrary software in Java
  - Hypotheses
- Architecture
  - Monitoring sidecars
  - Perturbation injectors
  - Chaos controller
- Output
  - Resilience report

Perturbation model: try-catch block short-circuit testing
- A corresponding exception at the beginning
- The whole try block is made invalid

Hypotheses
- RH (Resilience hypothesis)
- OH (Observability hypothesis)
- DH (Debug hypothesis)
- SH (Silence hypothesis)

Evaluation
- 3 large-scale and well-known Java applications totaling 630k lines of code

POBS

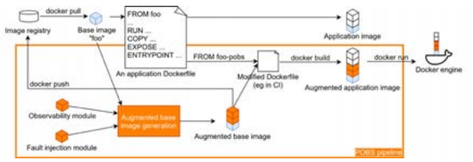

In order to address C2, we propose to design orchestration frameworks to connect monitoring, injection, and analysis. For most of the existing chaos engineering tools, there are several limitations: 1) steady state and hypotheses have to be manually defined, 2) the installation of a tool may be complicated, and 3) the setup of an experiment may be tedious.

Thus we propose a technique called POBS (imProved OBServability) to statically analyze and transform Docker configuration files of Java applications in order to inject observability capabilities [3].

For example, POBS allows developers to observe the JVM memory or CPU usage of their application **with minimal effort: a single line change in the Docker configuration**.

Phoebe

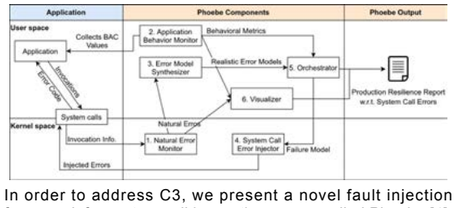

In order to address C3, we present a novel fault injection framework for system call invocation errors, called Phoebe [4]. Phoebe is unique as follows. First, Phoebe enables developers to have full observability of system call invocations. Second, **Phoebe generates error models that are realistic in the sense that they mimic errors that naturally happen in production.** Third, Phoebe is able to automatically conduct experiments to systematically assess the reliability of applications with respect to system call invocation errors in production.

References

1. Zhang et al, *A Chaos Engineering System for Live Analysis and Falsification of Exception-Handling in the JVM,* IEEE TSE, 2019.
2. Zhang et al, *TripleAgent: Monitoring, Perturbation and Failure-Obliviousness for Automated Resilience Improvement in Java Applications,* IEEE 30th International Symposium on Software Reliability Engineering (ISSRE), 2019.
3. Zhang et al, *Automatic Observability for Dockerized Java Applications,* arXiv preprint:1912.06914, 2019.
4. Zhang et al, *Maximizing Error Injection Realism for Chaos Engineering with System Calls,* IEEE TDSC, 2021.