POSTER CATALOGUE
WASP Winter Conference 2020

Poster Session 2
Tuesday 14 January 16.00-17.00

Poster no.	First name	Last name	Title of poster
2	Mattias	Tiger	Spatio-Temporal Learning and Reasoning for Situation Awareness in Robotics
6	Max	Åstrand	Short-term Mine Scheduling using Constraint Programming
10	Mia	Kokic	Learning Task-Oriented Grasping from Human Activity Datasets
14	-	-	-
18	Miquel	Martí i Rabadán	Semi-supervised multitask learning
22	Mohamadreza	Farid Ghasemnia	Robot Learning of Symbol Grounding in Multiple Contexts Through Dialog
26	Mårten	Lager	Smart Technologies for unmanned ships
30	Niklas	Gunnarsson	Real Time Image Guided Radiotherapy - Deep learning motion estimation
34	Niklas	Åkerblom	Robust Route Planning for Electric Vehicles
38	Olof	Zetterqvist	Learning with noisy labels
42	Olov	Andersson	Learning Safe Decision-Making for Autonomous Robots
46	Oskar	Ljungqvist	Motion planning and sensing-aware model predictive control of complex tractor-trailer combinations
50	Pavel	Anistratov	Autonomous-Vehicle Maneuver Planning Using Segmentation
54	Pavlo	Melnyk	Enabling Unsupervised Learning by Injecting Geometry
58	Per	Boström-Rost	Optimal range and beamwidth for radar tracking of maneuvering targets using nearly constant velocity filters
62	Per	Skarin	Automation of efficient, reliable, and elastic predictive control over the cloud
66	Petter	Restadh	Polytopes for Graphical Models
70	Piergiuseppe	Mallozzi	Engineering Trustworthy Self-Adaptive Autonomous Systems
74	Quantao	Yang	Reinforcement Learning in Continuous Spaces with Interactively Acquired Knowledge-based Models
78	Rasmus	Ros	Continuous Optimization of Software
82	Rebekka	Wohlrab	Living Boundary Objects to Support Agile Inter-Team Coordination at Scale *
86	Rita	Laezza	Reinforcement Learning for robotic manipulation of deformable objects
90	Robert	Gieselmann	Simulation and Manipulation of Deformable Objects using Deep Learning
94	Rui	Oliveira	Motion Planning for Self-Driving Long and Multi-Body Vehicles
98	Sebastian	Bujwid	Knowledge Transfer between Language and Vision *
102	Shahbaz	Abdul Khader	Learning Contact-rich Manipulation Skills
106	Shervin	Parvini Ahmadi	Distributed Localization of Sensor Networks Using Augmented Lagrangian *
114	Simon	Lindståhl	Configurable observability for 5G orchestration
118	Tomasz	Kosiński	Usable Privacy Controls for Interaction with Autonomous Agents *
122	Tommi	Nylander	Equivalent G/G/1 Modeling for Server Systems with Redundant Requests
126	Truls	Nyberg	Situational Learning and Decision Making for AVs
130	Vandana	Narri	Shared Situational Awareness Under Complex Traffic Scenarios
134	Victor	Fors	Yaw-Moment Control At-the-Limit of Friction Using Individual Front-Wheel Steering and Four-Wheel Braking
138	-	-	-
142	-	-	-
146	Vladislav	Polianskii	Voronoi Boundary Classification: A High-Dimensional Geometric Approach via Weighted Monte Carlo Integration
150	Yassir	Jedra	Reinforcement Learning and Statistical Learning for Control

*) Poster not available in this catalogue

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Wasp

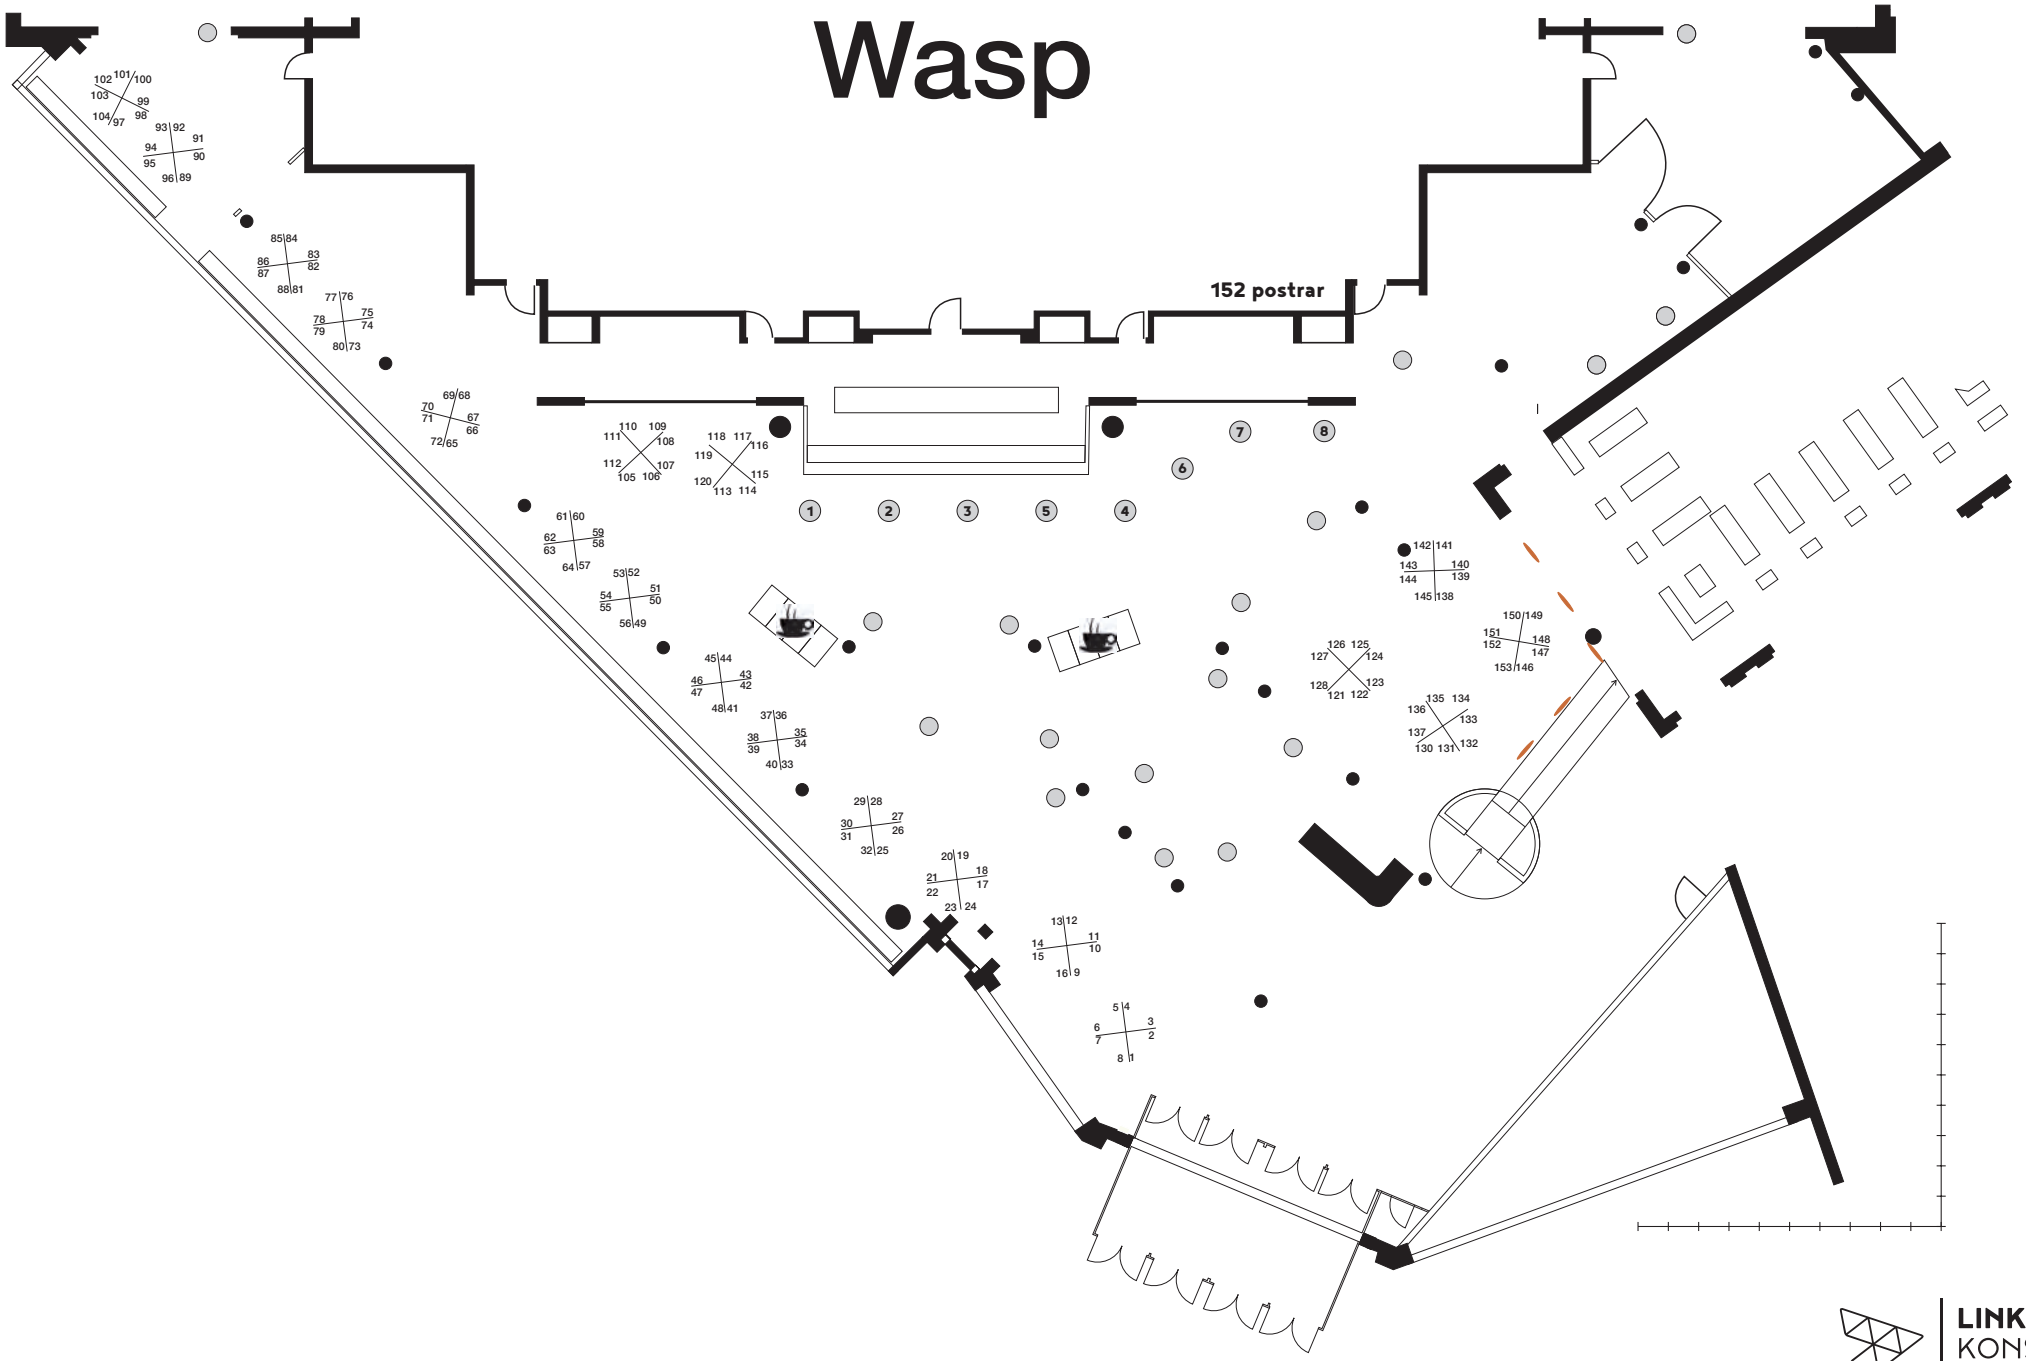

LINKÖPING
KONSERT &
KONGRESS

Spatio-Temporal Learning and Reasoning for Situation Awareness in Robotics

Mattias Tiger, Linköping University

Awareness of entity motion is central to robot safety. This is true regardless of whether the robot is driving, flying or working alongside humans. It is important to recognize common motion patterns and to discover new ones. Furthermore, the important ability to anticipate what will happen next requires good predictions of future movement. It is also useful to be aware of how predictable other entities are, and to detect abnormal behaviors. Accurate awareness of the current situation is necessary for robust, efficient and safe decision making and control. Our focus is on trajectory-based approaches for continuous learning using probabilistic machine learning techniques, and on integrating probabilistic and logical stream reasoning.

Probabilistic and predictive stream reasoning for runtime verification

Logical constraints can be used to specify safe operating behaviors and assumptions about the environment. We have proposed P-MTL **[3]**, a metric temporal logic for stream reasoning (incremental reasoning over rapidly-changing information) over probabilistic and predicted states. This addresses two important problems in AI; integrating logical and probabilistic reasoning and integrating reasoning over observations and predictions. This allows a robot to explicitly reason about the uncertainty of the world, the expected change of the world and the quality of its observations and predictions.

Reasoning over Uncertainty | Reasoning over Predictions

Example formula for monitoring prerequisites for safe motion

$$\Box\,(Pr(insideRelative(\bullet_{0:0}\text{Position}[\mathbf{uav1}], \mathbf{1mCircle})) > 0.99 \;\wedge\; Pr(insideRelative(\bullet_{10}\text{Position}[\mathbf{uav1}], \mathbf{1mCircle})) > 0.95 \;\wedge\; \bigcirc_3\,(similarity(\bullet_{0:0}\text{Position}[\mathbf{uav1}], \bullet_{0:-3}\text{Position}[\mathbf{uav1}]) > 0.5))$$

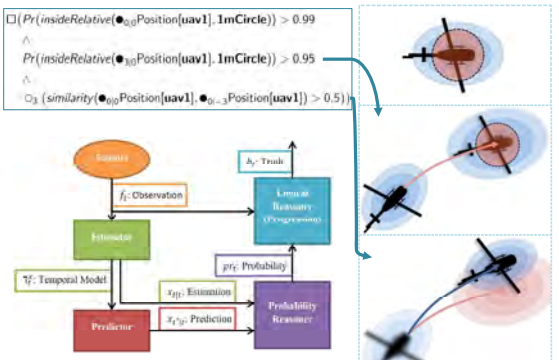

Computational Environment

Learning and recognizing trajectory-based spatio-temporal activities such as motion patterns

We have proposed an online unsupervised framework that learns a probabilistic representation of observed spatio-temporal activities and their transitions from observed trajectories **[1]**. It can recognize activities and predict activity chains. Bayesian networks are used to model the transition graph and Gaussian processes (GP) to model atomic activities.

We have developed Gaussian process trajectory modelling tools for merging and separating serially and in parallel connected models. This includes an efficient technique for approximating mixtures-of-GPs and efficient online trajectory model learning using GPs **[2]**.

We have further proposed extended GP-based models and methodology to motion patterns in large-scale complex road structures **[4]**. Such topology necessitates sequential local models.

Motion pattern recognition includes:
clustering, classification, prediction, abnormality detection

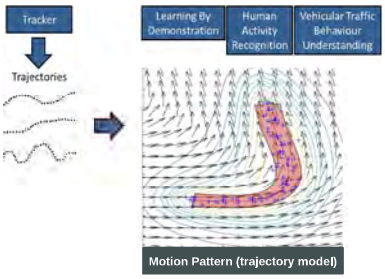

Motion Pattern (trajectory model)

Common motion patterns in a crossing

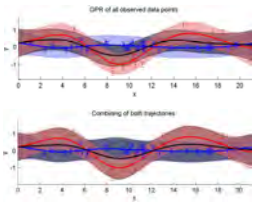

Mixture-of-GPs approximation

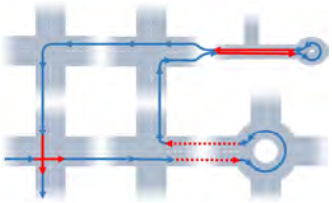

Large-scale modelling of normative motion patterns

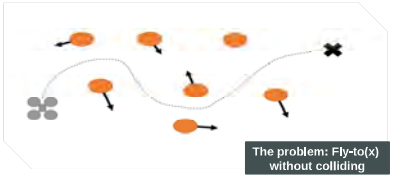

The problem: Fly-to(x) without colliding

DJI M100 UAV platform

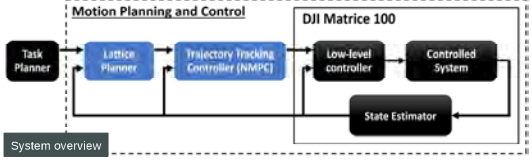

System overview

Example where we plan to wait for an obstacle to pass us by

Trajectory-based motion planning and control in large-scale, complex and dynamic environments

We have proposed **[5]** a principled solution to motion planning with dynamic obstacle avoidance by using a unified optimization-based motion planning and control architecture, where both layers use the system dynamics to generate and execute feasible trajectories in real-time.

The plans are made with respect to time in several ways: In terms of predicted future movements of agents, time duration of actions and by the ability to plan to wait.

Motion primitives (full trajectories) are generated offline using numerical optimal control and then used in online graph-search for finding feasible and cost-efficient plans.

Example of motion primitives

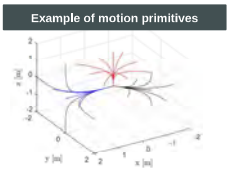

Full trajectory planning and tracking
State: 3D position, 3D velocity, roll, pitch, thrust

Does not get stuck locally

[1] M. Tiger and F. Heintz, *Towards unsupervised learning, classification and prediction of activities in a stream-based framework*, **SCAI** 2015.
[2] M. Tiger and F. Heintz, *Online sparse Gaussian process regression for trajectory modeling*, **FUSION** 2015.
[3] M. Tiger and F. Heintz, *Stream Reasoning using Temporal Logic and Predictive Probabilistic State Models*, **TIME** 2016.
[4] M. Tiger and F. Heintz, *Gaussian Process Based Motion Pattern Recognition with Sequential Local Models*, **IV** 2018.
[5] O. Andersson, O. Ljungqvist, M. Tiger, D. Axehill, F. Heintz, *Receding-Horizon Lattice-Based Motion Planning with Dynamic Obstacle Avoidance*, **CDC** 2018
[6] M. Selin, M. Tiger, D. Duberg, F. Heintz, P. Jensfelt, *Efficient Autonomous Exploration Planning of Large Scale 3D-Environments*, **RA-L** & **ICRA** 2019

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Short-term Mine Scheduling using Constraint Programming

Max Åstrand, Mikael Johansson, Alessandro Zanarini
max.astrand@se.abb.com, mikaelj@kth.se, alessandro.zanarini@ch.abb.com

ABSTRACT

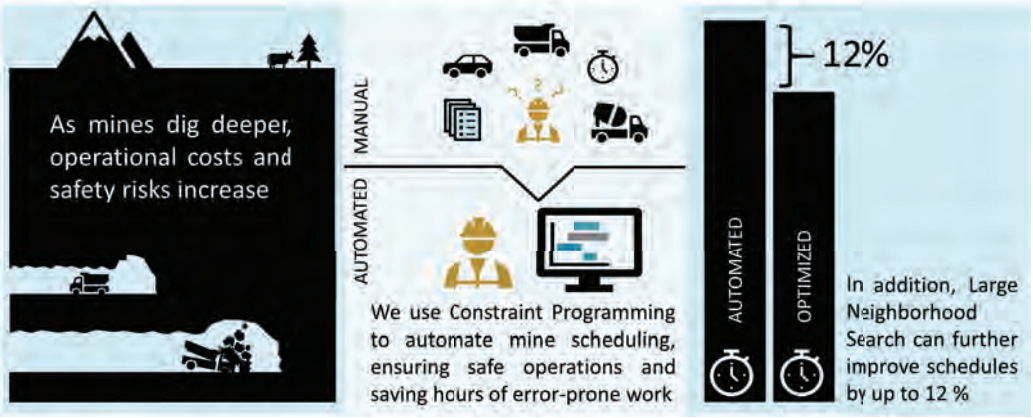

Manual short-term scheduling in underground mines is a time-consuming and error-prone activity. We use Constraint Programming to automate the scheduling process: deciding *what* to do *where* and *when*.

We extend previous work by accounting for fleet travel times, and by introducing a new model based on solving a related scheduling problem and transforming its solution back to the original domain. In addition, a neighborhood definition is introduced to optimize using Large Neighborhood Search.

Results show that the proposed method scales to realistic problem sizes, and that the solutions obtained are of high-quality.

KEY RESULTS

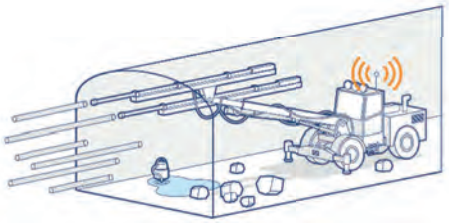

- The mine scheduling problem resembles a rich variant of a k-stage flow shop, with a mix of interruptible and uninterruptible jobs, periodically induced machine unavailabilities, after-lags in some stages, sharing of (certain) machines between stages, and sequence-dependent setup times due to the travel times of the mobile machines [1].
- Underground mines can have road networks spanning several hundreds of kilometers. Therefore, to ensure that schedules are feasible to operationalize, we extend previous work [2] by including travel times of the mobile machines in the constraint model.
- In addition, we propose a new approach based on first generating solutions to a modified *uninterruptible* scheduling problem without blast windows. A post-processing step inserts blast windows and transforms the solutions to solve the original problem. To further improve the obtained schedules, Large Neighborhood Search is used with a domain-specific neighborhood definition based on relaxing all variables corresponding to jobs scheduled at a random subset of production areas.
- We can find high-quality schedules to realistic instances, generated using data from an operational mine, including more than 200 jobs. Compared with a common constructive heuristic [3], solutions are found within minutes exhibiting $\sim 7\%$ lower objective value. Studying the optimal solution to a relaxed problem, we note that on a realistic instance we are at most $\sim 12\%$ away from optimality.

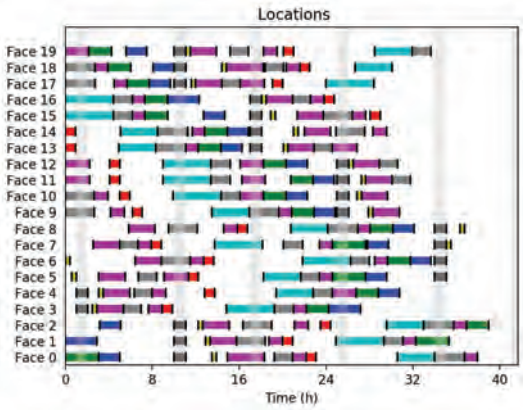

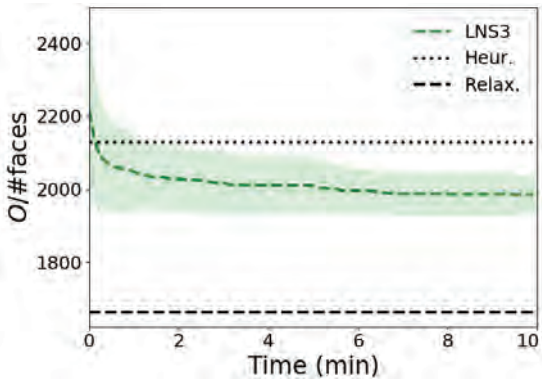

REFERENCES:

[1] Åstrand, M., Johansson, M., & Greberg, J. (2018). Underground mine scheduling modelled as a flow shop: a review of relevant work and future challenges. Journal of the Southern African Institute of Mining and Metallurgy, 118(12), 1265-1276.

[2] Åstrand, M., Johansson, M., & Zanarini, A. (2018, June). Fleet Scheduling in Underground Mines Using Constraint Programming. In International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research (pp. 605-613). Springer, Cham.

[3] Pinedo, M. (2012). Scheduling (Vol. 29). New York: Springer.

Learning Task-Oriented Grasping from Human Activity Datasets

Mia Kokic, Danica Kragic, Jeannette Bohg

Motivation

Task-Oriented Grasping (TOG): Object grasped such that it can be used for a task.

To infer: *Where* and *how* to the place the hand.

But: Data Collection is tedious.

Core Question: Can we leverage RGB human activity datasets to teach robots TOG?

Yes! But we need to lift them to 3D.

Goal

1. Process RGB datasets to obtain 6D hand and object poses and shapes.
2. Use this data to learn task-compliant regions on object for grasping.

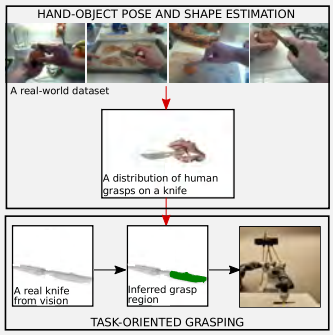

Contributions

A **framework** for estimating hand pose and configuration as well as object pose and shape from a single RGB image.
Competitive results to state-of-the-art[1] and generalization to novel object instances from three categories.
Learned TOG models from a challenging real-world dataset of humans manipulating objects.
Real robot demonstrations of task-oriented grasps of previously unseen object instances from a given category.

[1] Y. Hasson, G. Varol, D. Tzionas, I. Kalevatykh, M. J. Black, I. Laptev, and C. Schmid, "Learning joint reconstruction of hands and manipulated objects," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition , 2019, pp. 11 807–11 816.

System Overview

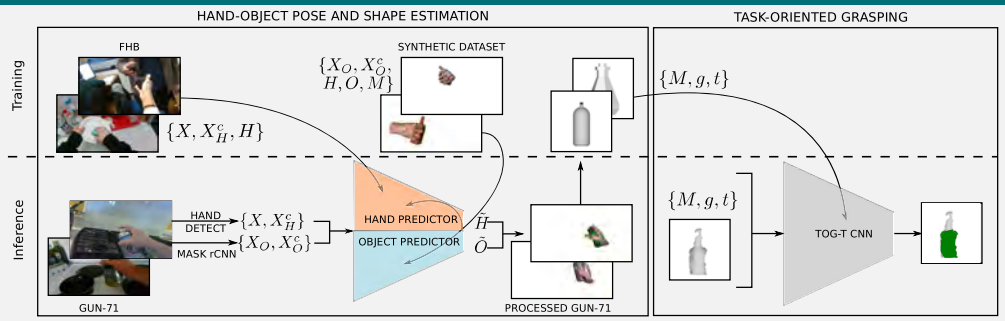

Datasets

Real data: *FHB [2]* contains video sequences from ego-centric perspective; hands are annotated with 3D joint positions but object annotations are available for only 4/26 instances.

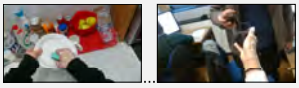

Synthetic data: easy to collect object annotations, but the data differs from real RGB images (sim-to-real gap)

A combination of **synth and real data** ensures generalization to novel objects and novel hand poses.

[2] G. Garcia-Hernando, S. Yuan, S. Baek, and T-K. Kim, "First-person hand action benchmark with rgb-d videos and 3d hand pose annotations," in CVPR , 2018, pp. 409–419.

Method

1. Train a CNN that predicts hand poses and configurations + object poses and shapes from RGB images.

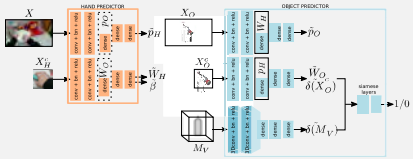

2. Process a real-world RGB human activity dataset [3].

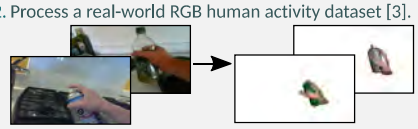

3. Train a category-specific CNN that predicts if a grasp on a novel object is suitable/unsuitable for a task.

[3] G. Rogez, J. S. Supancic, and D. Ramanan, "Understanding everyday hands in action from rgb-d images," in Proceedings of the IEEE international conference on computer vision , 2015, pp. 3889–3897.

Results

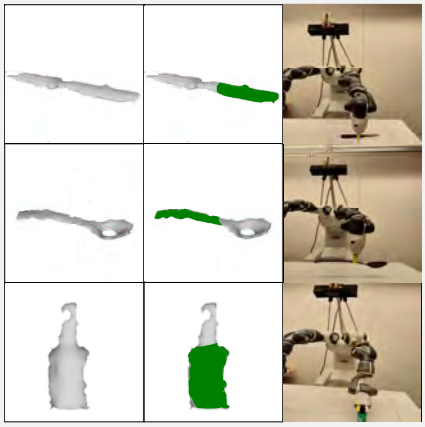

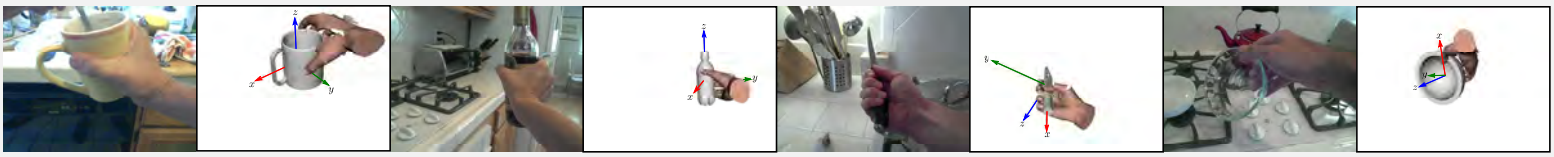

Semi-supervised multitask learning

Miquel Martí [1,2] Alessandro Pieropan[2] Hossein Azizpour[1] Atsuto Maki[1]

[1]KTH Royal Institute of Technology, Stockholm, Sweden [2]Univrses AB, Stockholm, Sweden

Motivation

Multitask learning offers efficient deployment of models for complex scene understanding but introduces new practical issues, which are key to enable its use in the industry.

Multitask learning modes:

- (a) - Single set of inputs with different, overlapping sets of labels
- (b) - Multiple sets of inputs with different sets of labels, non-overlapping
- (c) - Same as (a) with labels for one task covering whole input set, for another only partially. Occurs naturally due to different labelling costs for different tasks.

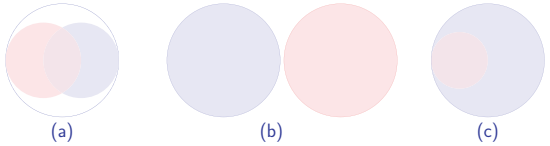

Figure: Different multitask cases for datasets (area within circles) and label sets (shaded discs).

Approach

Multitask models with shared encoder. We want to leverage unlabeled data for partially labeled task too. Use both distant supervision from multi-task learning and consistency regularization semi-supervised learning methods such as Virtual Adversarial Training [1] on such partially labeled samples for the task with missing label.

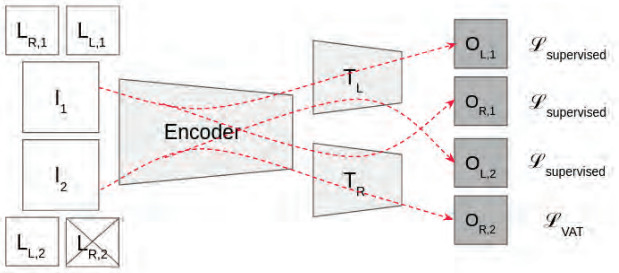

Figure: Batch of two images passed to the multi-task model with two tasks. For the outputs without corresponding labels the VAT loss is used.

Related work

Some works have addressed the problem of unbalanced label sets in multi-task learning before with a Pseudo-Labels [2] in Natural Language Processing [3] and in computer vision in [4], *pseudo-labeling* one new sample at each epoch and adding it to the training set, and in [5], using a knowledge distillation loss with previous snapshots of the model.

Figure: Samples from MultiMNIST dataset.

Experiments

We compare models on MultiMNIST[6,7], which consists of two classification tasks of digits partially overlaid on top of each other – *L* top-left digit, *R* bottom-right. We simulate the partial labelling for one task by keeping only a percentage of the labels for this task. For the sample and task label pairs where the label is missing we simply have zero loss. We consider three cases: fully labeled, 50% of labels and 10% of labels.

We use Virtual Adversarial Training[1] on the partially-labeled task only when there is no label for a given sample. We compare uniform task weights to searching over static weights. We report mean accuracy and standard deviation over 9 runs on the test set with the hyper-parameter configuration giving the best validation accuracy.

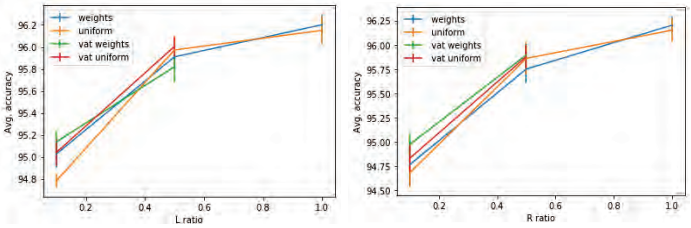

Figure: Mean average accuracy of both L and R tasks on 9 runs ± standard deviation for different ratios of labels for both L (left) and R (right) tasks for uniform multitask loss, weighted loss and both with VAT loss on unlabeled samples.

Conclusions and future work

- Differences between methods indicate that combining VAT and multi-task might lead to small improvements in most cases. However, results are within standard deviation over different seeds and increase the variance of the experimental results.
- Larger improvements happen in the experiments with tasks with fewer labels, suggesting that the method might be more effective in the less supervised regimes.
- More experiments with different datasets are needed in order to validate this preliminary results. Also, experiments for cases with even less labels for one of the labels should be carried out.

References

1. T. Miyato, S. Maeda, M. Koyama, and S. Ishii. *Virtual adversarial training: a regularization method for supervised and semi-supervised learning.* In ICLR 2016.
2. Dong-Hyun Lee. *Pseudo-Label : The Simple and Efficient Semi-Supervised LearningMethod for Deep Neural Networks.* In ICML 2013 Workshop.
3. Marek Rei. *Semi-supervised Multitask Learning for Sequence Labeling.* In ACL 2017.
4. N. Khosravan and U. Bagci. *Semi-supervised multi-task learning for lung cancer diagnosis.* In EMBC 2018.
5. D. Kim et al. *Disjoint Multi-task Learning between Heterogeneous Human-centric Tasks.* In WACV 2018.
6. Sabour, N. Frosst, and G. E. Hinton. *Dynamic routing between capsules.* In NIPS 2017.
7. O. Sener and V. Koltun. *Multi-task learning as multi-objective optimization.* In NeurIPS 2018.

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

Robot Learning of Symbol Grounding in Multiple Contexts Through Dialog

Mohamadreza Farid Ghasemnia
mohamadreza.farid@oru.se

Advisors: Alessandro Saffiotti, Lars Karlsson
Project partner: Bram Willemsen (KTH)
Center for Applied Autonomous Sensor Systems
Örebro University, SE-70182 Sweden

Problem Statement

How should a robot ground symbols of a situated dialog, and in multiple contexts?

User: Give me the red?
Robot: Okay, but which one do you mean now?

- **Plans for discourse:** Participants in a discourse use actions (utterances) to accomplish a goal (intent of dialog).
- **Dialog:** Discourse between two or more participants, eg. spoken dialog.
- **Common sense in dialog:** A robot should be able to understand the user's intent, and talk in such a way that the user understands. Participants in a dialog must have common sense, or they have to align their belief so they can arrive at a common interpretation.
- **Situated dialog:** A robot should be able to process its perception to understand its surrounding.
- **Context:** various information of a situated dialog can be found in "context", such as the position of the robot and user, information corresponding each object, user's profile, time and etc.
- **Learning symbol grounding in multiple contexts:** Grounding language in perception may be affected by different types of uncertainty.

  Which object the user is talking about?

  What is the semantic of a word?

  A robot should be able to learn how to remove such uncertainties.
- **Source of uncertainty:** In a situated dialog, different types of noise make a robot uncertain, eg. ASR . Moreover, different sources of information may be in conflict and may confuse the robot about which information to rely on. Possible sources include:
  - An *a-priori* knowledge base
  - User assertions in another context
  - User assertions in current context

Approach

Develop a precise formulation of context, ontology and their relation, which shows the most promising assertion in the ontology with respect to the context.

To use context to help the robot decide on which information to rely, or else to ask a question to obtain new information.

Current Setup with Pepper

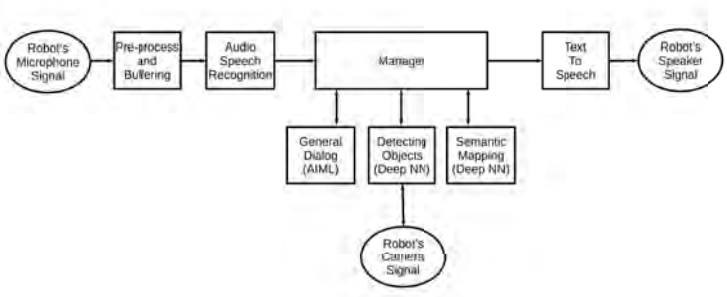

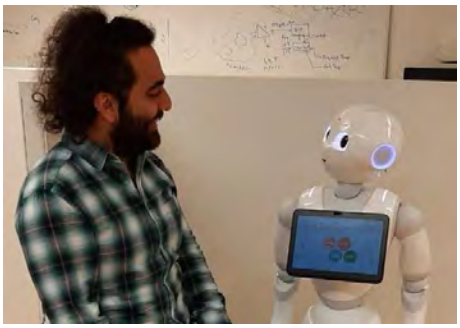

Current Platform

- Subscribing to Pepper camera and microphone
- Pre-processing of microphone signal
- Google ASR
- Managing multiple answer generators:
  - General dialog handler
  - Object descriptor handler
  - Perceived object handler
- Pepper Text to Speech

Current Showcase

Hello Pepper!

Hi there!

What can you see?

I see book, pen and desk.

The desk is out of order!

Oh! Ok, I save the desk as improper.

Conclusion

- Different types of uncertainty exist.
- Dialog is a shared plan among participants, for a common goal.
- Ontology and context are our two key concepts.
- Interpretation can be found using context.
- Interpretation changes upon context.
- A robot should know what does each word mean in each context.

Smart Technologies for unmanned ships

Mårten Lager, Lund University
Computer Science Department

Research Area

In this project, methods are presented to enhance the capability of two building blocks that are important for autonomous ships; a positioning system, and a system for remote supervision.

GPS-free navigation by fusion of bottom depth and magnetic field with CNN support

Overall description

Uses Bayesian calculations to compare the **bottom depth** and **magnetic field** measurements with known sea charts and magnetic field maps, in order to estimate the position. To optimize how the sensor data shall be fusioned, **CNN** adjusts the weights, after analyzing the map around the estimated position.

Our main contribution

State-of-the-art techniques for this method normally use low accuracy navigation sensors and high-resolution maps, which can hardly be used in real life. We rely on available normal **sea charts** and **low-resolution magnetic field maps** instead.

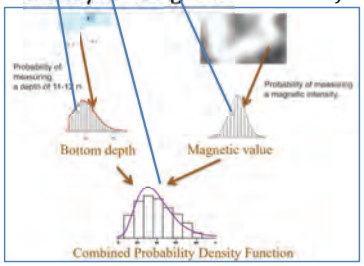

Results

- Good performance when comparing bottom depth with sea chart.
- Accuracy and robustness are increased when fusing bottom depth and magnetic field measurements.

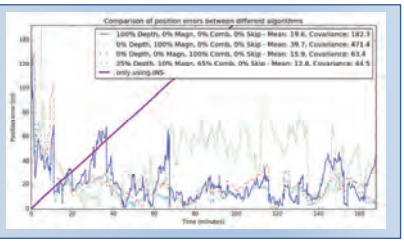

Remote supervision in VR

What has been done

Three GUI types have been created for **remote supervision** of an Unmanned Surface Vessel via a **low bandwidth** connection. Two in **3D** (presented in **VR** and on a **laptop**) and one **traditional 2D GUI**. The GUIs have been compared in a **user study**.

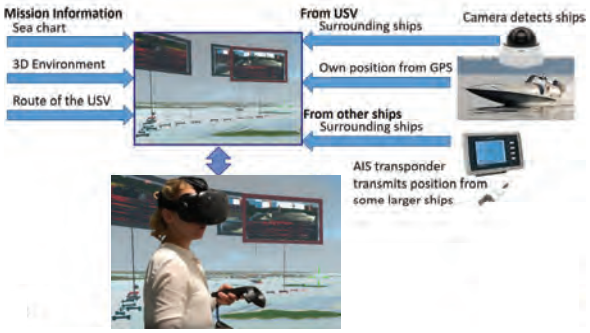

Why remote-supervision?

- Humans can handle complex dynamic environments

Why VR?

- Can provide a realistic environment comparable to what the operator is used to on a real ship
- Can augment information and guide the operator
- Do not need to transfer all videos, as most objects are already known in the Virtual World.

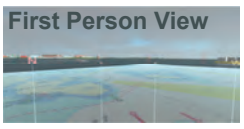

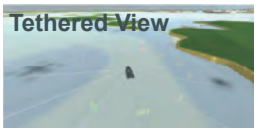

User Study Results

The users of the two 3D GUIs were **better at reacting to dangerous situations** than Traditional GUI users, and they could **keep track of the surroundings more accurately**. The users experienced the two 3D GUIs to be more **Easy to Use**, and believed the 3D GUIs, and especially the VR version to be the **Best Expert Tool** after several hours of training.

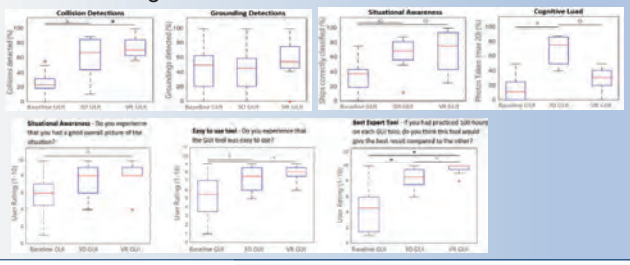

marten.lager@cs.lth.se
Department of Computer Science
LTH, Lund University
Sweden

Webpage:

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Real Time Image Guided Radiotherapy
- Deep learning motion estimation

UPPSALA UNIVERSITET

Niklas Gunnarsson, Thomas Schön, Jens Sjölund

Department of Information Technology, Uppsala University

Introduction and background

With the new radiotherapy device Elekta Unity, a linac with an integrated MRI unit, it is possible to do imaging of the patient during beam delivery. This makes it possible to adjust the treatment in real time in response to patient and organ motion.

Figure 1: Elekta Unity.

A big advantage of this type of real-time adaptive radiotherapy is that the radiation field can be focused on the actual tumour, avoiding margins which are currently being applied to compensate for e.g. organ motion.

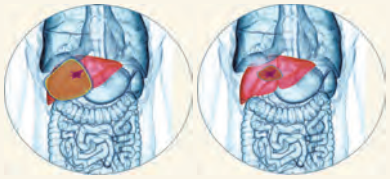

Figure 2: Left image: Currently a larger radiation field is used to compensate for motions. Right image: If the motion is known the margin can be decreased.

This requires fast and reliable algorithms for motion modelling. As a first step towards a full 3D motion model, we have investigated a state-of-the-art model in computer vision to estimate the motion between sequential 2D images on a pixel level. Once the estimation of the motion in 2D is good enough, the next step will be to reconstruct the entire 3D anatomy using the estimated 2D motion as input. The 3D images are needed in order to perform treatment plan adaptation and, as a QA tool, dose calculation in real time.

Method

The model we have investigate is the PWC-Net [5]. PWC-Net is a end-to-end deep learning algorithm based on techniques from traditional optical flow methods. The model takes two frames from e.g. a video sequence as inputs and the output is the estimated motion/flow for every pixel between the frames. An illustration of the net is shown in Figure 3

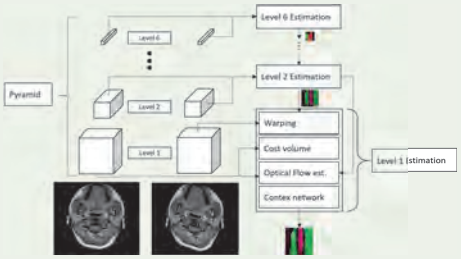

Figure 3: Illustration of PWC-Net.

The different components in the PWC-Net are:

- **Pyramid:** The pyramid contains of convolutional layers that downsamples the frames to six feature levels. The number of channels in each layer is 16, 32, 64, 128 and 196.
- **Warping:** Warping by use the features of the second frame and the upsampled estimated flow from the previous level.
- **Cost volume:** Correlation between features in the first frame and warped features from the second frame.
- **Optical Flow estimator:** The cost volume, features of the first frame and the upsampled estimated flow from the previous level are inputs to a multi-layer CNN. The output is the estimated flow at specific feature level with dimension height × width × 2. The optical flow estimator can also use DenseNet connections [3].
  - **DenseNet:** The inputs to every convolutional layer are the output and the input from its previous layer.

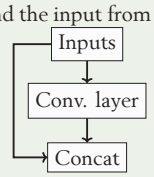

Figure 4: DenseNet connection.

- **Context network:** Contains of seven convolutional layers with different dilation constants.

Implementation

Our implementation of PWC-Net follows the standard procedure described in the article [5]. The implementation was made using Tensorflow 2 and trained on a Lambda Labs Quad workstation with a GeForce GTX 1080 Ti graphic card. One of the challenging part for optical flow and supervised machine learning methods in general is the lack of ground truth training data. For that purpose Dosovitski et al. created different synthetic datasets [1]. Some of the datasets are FlyingChairs, ChairsSDHom, FlyingThings3D, Driving and Monkaa. For our model we used a mixed of FlyingChairs and FlyingThings3D for training. The total size of the dataset is 43 010 where 97% was used for training and 3% for validation. To fit the investigated domain of medical images the dataset was converted to grayscale images and scaled to 256 × 256 pixels. Figure 5 shows one example from each set.

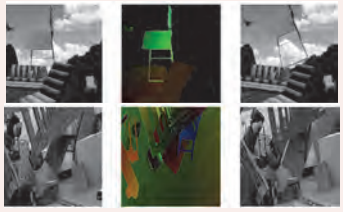

Figure 5: First frame, ground truth optical flow and second frame is shown for one sample for each dataset.

The loss function for used for training is a weighted cost between the estimated flow, $\boldsymbol{w}_\theta^l$ and ground truth flow, $\boldsymbol{w}_{\mathrm{GT}}^l$ in each level ($\boldsymbol{l}$) of the pyramid

$$\boldsymbol{L}(\Theta) = \sum_{l=l_0}^{L} \alpha_l \sum_{x} |\boldsymbol{w}_\theta^l(x) - \boldsymbol{w}_{\mathrm{GT}}(x)^l|_2 + \gamma|\Theta|_2. \quad (1)$$

where Θ is all parameters in the model. $\boldsymbol{l}_0$ is set to 2 and the weights $\alpha_6 = 0.32$, $\alpha_5 = 0.08$, $\alpha_4 = 0.02$, $\alpha_3 = 0.01$, $\alpha_2 = 0.005$. For the regularization term $\gamma = 0.0004$. Another metric to measure the quality for estimated optical flow is to use the endpoint error (EPE) between the estimated flow, $(\hat{\boldsymbol{v}}, \hat{\boldsymbol{u}})$ and the corresponding ground truth, $(\boldsymbol{v}, \boldsymbol{u})$

$$\boldsymbol{EPE} = \sqrt{(\hat{u} - u_{\mathrm{GT}})^2 + (\hat{v} - v_{\mathrm{GT}})^2}. \quad (2)$$

When ground truth values are unknown the quality of the estimated flow are measured by warping the first frame using the estimation and compute image similarity with the second frame. Such metrics are structural similarity index (SSIM) [6] and mean square error (MSE) [4].

Result

Figure 6-7 show the learning curve and EPE for both models.

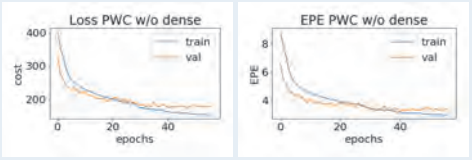

Figure 6: Learning curve and EPE for PWC-Net without DenseNet.

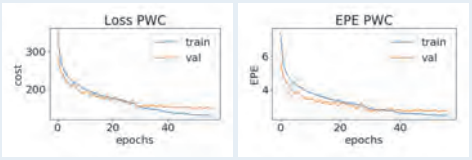

Figure 7: Learning curve and EPE for PWC-Net.

The model was tested on a medical head and neck MRI images set including the original image named reference image and an applied sine flow named moving image. The result was compared with a traditional optical flow algorithm called Farnebäck [2].

Model	SSIM	MSE	t(s)
PWC w/o DenseNet	0.9267	2.521	0.050
PWC with DenseNet	**0.9332**	**2.408**	0.045
Farnebäck	0.9317	2.5020	**0.034**

Table 1: Structural similarity index (SSIM), mean squared error (MSE) and estimation time (t(s)) for the different algorithms.

Illustrations of the estimated flow and warped images for both models and Farnebäck are shown in Figure 8

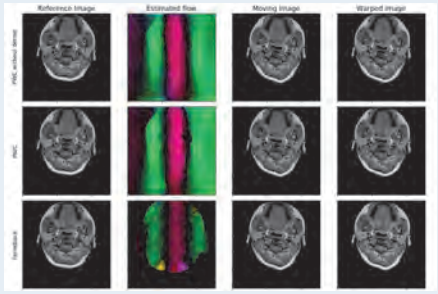

Figure 8: The reference image, estimated flow, moving image and warped image for the different algorithms.

Conclusions and further work

The result looks promising. With no information about the medical image domain the model estimates flow with similar performance as a well known optical flow method (Farnebäck). However the evaluation was performed on one synthetic case and needs to be evaluated further. It would be of interest to fine-tune the model with natural or synthetic medical images and evaluate the result.

References

[1] A. Dosovitskiy, P. Fischer, E. Ilg, P. Häusser, C. Hazırbas, V. Golkov, P. v.d. Smagt, D. Cremers, and T. Brox. Flownet: Learning optical flow with convolutional networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.

[2] G. Farnebäck. Two-frame motion estimation based on polynomial expansion. In *Scandinavian conference on Image analysis*, pages 363–370. Springer, 2003.

[3] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.

[4] E. L. Lehmann and G. Casella. *Theory of point estimation*. Springer Science & Business Media, 2006.

[5] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. 2018.

[6] Z. Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

niklas.gunnarsson@it.uu.se, niklas.gunnarsson@elekta.com, thomas.schon@it.uu.se, jens.sjolund@elekta.com

Robust Route Planning for Electric Vehicles

Niklas Åkerblom

Chalmers University of Technology

Description

The purpose of this project is to develop computationally efficient, accurate and robust algorithms for electric vehicle route planning and navigation. This combines combinatorial optimization with machine learning methods for modelling uncertainties in e.g. travel time, energy consumption and charging station performance.

Motivation

A common concern that people have about fully electric vehicles is the so-called "range anxiety". Due to the historically high cost of batteries, the driving range of electric vehicles has generally been much shorter than that of conventional vehicles, which has resulted in a fear of being stranded when the battery is depleted. It is possible to partly alleviate these concerns by providing improved navigation algorithms and route planning systems.

While batteries in newer electric vehicles often have higher capacity than in the past, a key difference which still exists between conventional combustion engine vehicles and battery electric vehicles is that charging a battery is much slower than refueling, even with fast charging technology. When charging is necessary during a trip, the charging time then becomes a significant part of the total travel time.

This motivates the inclusion of charging locations in route planning algorithms for electric vehicles. Since battery energy is a finite resource which is continuously depleted during a trip, it is natural to model the problem of finding the most efficient feasible path between two points as a shortest path problem with resource constraints. An additional factor is that while there are often many possible choices of charging stations for a complete trip, the density of charging station locations may be sparse in certain regions, which means that choosing an occupied or otherwise unavailable charging station may also incur a significant cost in travel time.

The project is coordinated by Volvo Car Corporation and is conducted in cooperation with Chalmers University of Technology.

CHALMERS UNIVERSITY OF TECHNOLOGY

VOLVO

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

With funding from the Strategic Vehicle Research and Innovation Programme (FFI):

FFI

Research Questions

- How can route decision making be supported for individual electric vehicles or autonomous fleets of electric vehicles?
- How do we handle the relatively long charging times of electric vehicles and the severe consequences of unavailable charging stations in a robust manner?
- How do we integrate accurate models for electric vehicle energy consumption and charging into an efficient approximate solution for a constrained shortest path problem, while avoiding unacceptable tradeoffs in terms of computational complexity?
- How can we incorporate cost effective exploration and data collection for such models into our route planning methods, without making significant compromises to the travel time and energy costs of individual consumers?

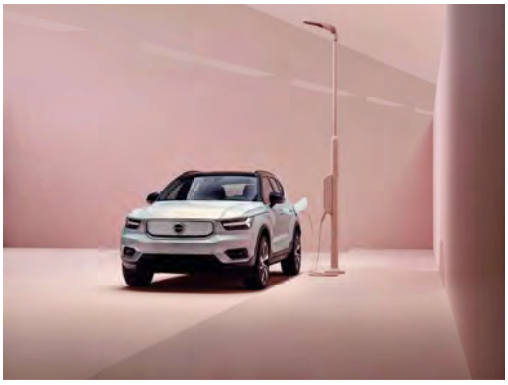

Approach & Status

Several promising avenues of research have been identified and are being pursued. For example, recent state-of-the-art papers show that road network graphs can be preprocessed with contraction hierarchies and depletion profile information to speed up constrained shortest path queries. [1]

We plan to extend these deterministic methods with probabilistic models for e.g. energy consumption, and we are currently evaluating efficient ways of updating such models with new information.

References

[1] Moritz Baum, Jonas Sauer, Dorothea Wagner, and Tobias Zündorf. Consumption Profiles in Route Planning for Electric Vehicles: Theory and Applications.
In Costas S. Iliopoulos, Solon P. Pissis, Simon J. Puglisi, and Rajeev Raman, editors, *16th International Symposium on Experimental Algorithms (SEA 2017)*, volume 75 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 19:1--19:18, Dagstuhl, Germany, 2017. Schloss Dagstuhl--Leibniz-Zentrum fuer Informatik.

Learning with noisy labels

Olof Zetterqvist
olofze@chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY

Goal of project

In the paradigms of supervised, semi-supervised or active learning with machine learning algorithms in general and deep neural nets in particular, the default models assume that the given annotations of training data are correct. However, this is not always the case. With a lower requirement on the quality of the data, the higher amount of data is available at the same price. The goal of our project is to get a better understanding of how label noise affects supervised learning. Some of the questions that we are interested in are:

- How to identify misclassified samples efficiently?
- How to make a model more robust against misclassified data?
- Does the robustness depend on the distribution of the noise?
- How do different types of regularisation techniques affect the robustness?

A new regularisation technique

- We extend the typical cost function for a neural network and put weights on the training points:

$$\min_{\theta,\omega:\,\omega>=0} \sum_i \omega_i L(x_i,\theta) + \frac{\alpha}{2}(\omega_i - 1)^2 + \frac{\lambda}{2}||\theta||^2, s.t.\ \bar{\omega} = 1$$

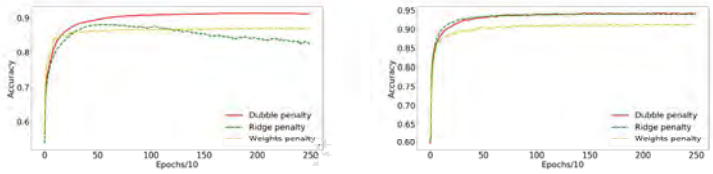

Figure 1: The accuracy on a test set during training for different penalty's on the CIFAR10 dataset. Left: With 20% label noise on training data. Right: With no label noise on training data.

- Advantages with the new cost function: Both experimental and some theoretical results seam to show that the training will become more robust against label noise.
  The model can converge to an optimum which will give a more theoretical understanding of our model.
- Disadvantages with the new cost function: The training will be slower since we have introduced more parameters.
  We will get at least one extra hyperparameter that needs to be optimized.

Future work

- How well does this generalize to a multiclass problem?
- How should ω_i be optimized?
- By having different "weight center-points" for different classes, could we get a more fair representation of class sizes?
- There is nothing that says that the penalty $(\omega_i - 1)^2$ is optimal. In a practical sense it's not since it requires the constraint $\omega_i \geq 0$ in the optimization. Is there a better penalty that is both numerical stable and has a good mathematical justification to be used.
- How sensitive is the method to variations in α and λ? Is it important to find a true optimal value of these parameters?

References

[1] Koh, Pang Wei, and Percy Liang. "Understanding black-box predictions via influence functions." *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017.

[2] Li, Mingchen, Mahdi Soltanolkotabi, and Samet Oymak. "Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks." *arXiv preprint arXiv:1903.11680* (2019).

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning Safe Decision-Making for Autonomous Robots

Olov Andersson, Linköping University
Artificial Intelligence and Integrated Computer Systems
Main advisor: Patrick Doherty

Motivation & Research Goals

Robots are increasingly expected to go beyond controlled environments in laboratories and factories, to enter real-world public spaces and homes. However, robot behavior is still usually engineered for narrowly defined scenarios. To manually encode robot behavior that works within complex real world environments, such as busy streets or work places, can be a daunting task. The aim of this research is to examine efficient methods for automatically **learning robot behavior under uncertainty**, lowering the costs of deploying robots to the real world. A key focus is satisfying the **safety requirements** and the **resource constraints** imposed by autonomous robots.

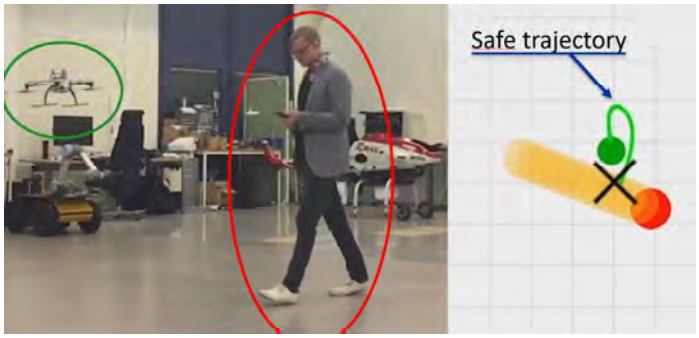

Methods

Formally, we seek general-purpose approximations to **planning and control** under **uncertainty**

$$\arg\min_{\pi(x)} \quad \mathbb{E}\left[c(\tau_{t\ldots t+H})\right]$$
$$\text{subject to}$$
$$\Pr(g(\tau_{t\ldots t+H}) \geqslant 0) > p$$

…with the **computational limitations** *and* **safety constraints** of real robot platforms

These are typically intractable. We instead leverage both **machine learning** techniques, and engineering techniques from **robotics** and **control**, to compute **approximations that satisfy safety constraints**. These have certain **robustness advantages** over deep reinforcement learning approaches **[5]**.

In particular we draw upon:
- Bayesian Learning & Bayesian Optimization
- Deep Learning & Deep Reinforcement Learning
- Trajectory optimization & Model-Predictive Control (MPC)

References

1. Andersson, O., Wzorek, M., Rudol, P., & Doherty, P. *Model-Predictive Control with Stochastic Collision Avoidance using Bayesian Policy Optimization*, Int. Conf. on Robotics and Automation **(ICRA)**, 2016.
2. Andersson, O., Wzorek, M., & Doherty, P. Deep learning quadcopter control via risk-aware active learning. AAAI Conference on Artificial Intelligence **(AAAI)**, 2017.
3. Andersson, O. (2017). *Methods for Scalable and Safe Robot Learning*, Licentiate Thesis, Linköping University.
4. Andersson, O., Ljungqvist, O., Tiger, M., Axehill, D., & Heintz, F. Receding-horizon lattice-based motion planning with dynamic obstacle avoidance, **CDC**, 2018.
5. Andersson, O., Doherty, P. DeepRL for Autonomous Robots – Limitations and Safety Challenges, **ICML'18 RML Workshop and ESANN'19**.
6. Andersson, O., Sidén, P., Dahlin, J., Doherty, P., & Villani, M. Real-Time Robotic Search using Hierarchical Spatial Point Processes. Conference on Uncertainty in Artificial Intelligence ***(UAI)**, 2019.*

Selected Results

Collision avoidance in **mixed human-robot** environments is one example of a behavior that is difficult to manually engineer
- There is considerable **uncertainty** due to inexact models, sensors, and especially difficult-to-predict human motion
- Need to consider **dynamics constraints** to find safe trajectories through busy workplaces or streets

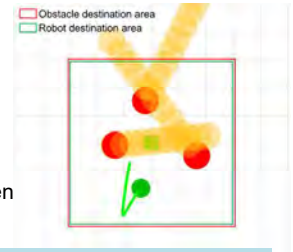

Example: Warehouse scenario
- Humans and UAV in small workspace
- UAV wants to pick up green packages
- 3 **non-cooperative** moving obstacles given destinations randomly

Summary of contributions (see [1]):
- Novel ***constrained*** **Bayesian policy optimization** to find deterministic **MPC approximations** that satisfy the **safety constraints under uncertainty**
- Demonstrated **real-time** MPC solution **to non-cooperative collision avoidance** under **uncertainty** and **dynamics** for flights with a **real quadcopter**

MPC with safe collision avoidance possible in real-time, but still requires capable on-board CPU.
- Want to synthesize behavior for **smaller robots and embedded systems**.


7 cm Bitcraze Crazyflie

Summary of contributions (see [2]):
- Learn fast **deep neural network** approximations for problems with **safety constraints**
- Demonstrated by embedding neural network for collision avoidance scenario **on-board nano-quadcopter** microcontroller

In a WASP collaboration we also considered lattice approximations to motion planning in such complex dynamic environments.

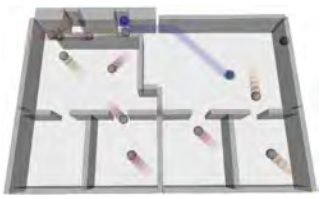

Summary of contributions (see [4]):
- Real-time **3D motion-planning in time** for moving obstacles
- Unified optimization-based graph planning & control architecture

Another interdisciplinary collaboration considered automating aerial drone search for victims after disasters (e.g. earthquakes).

Summary of contributions (see [6]):
- **Real-time learning and inference** in structural spatial point process model
- Minimized expected victim harm via **real-time Monte-Carlo tree search**

Motion planning and sensing-aware model predictive control of complex tractor-trailer combinations

Oskar Ljungqvist and Daniel Axehill

Motivation

The design of reliable path-following controllers is a key ingredient for successful deployment of autonomous tractor-trailer vehicles. It is challenging since the vehicle is unstable in backward motion and because the tractor has steering limitations. Additionally, optical sensors with a limited field of view (FOV) has been proposed to solve the joint-angle estimation problem online, which introduce restrictions on which states that can be reliably estimated. Moreover, in recent years there has been an increased interest for long tractor-trailer combinations to meet efficiency demands related to transportation. To improve these long vehicles ability to maneuver in confined environments, some trailers can be equipped with steerable wheels.

Sensing-aware model-predictive path-following control

Problem: Given a nominal path $(\boldsymbol{x}_r(\cdot), \boldsymbol{u}_r(\cdot))$ for a reversing general 2-trailer (G2T) vehicle, design a path-following controller that minimizes the path-following error $\tilde{\boldsymbol{x}}(t) = \boldsymbol{x}(t) - \boldsymbol{x}_r(s(t))$ while satisfying the vehicle's sensing $(\beta_3, \tilde{\beta}_2) \in \mathbb{P}$ and input constraints $\boldsymbol{u} \in \mathcal{U}$.

Solution: A controller based on model predictive control (MPC) that uses sensing and jack-knife preventing constraints:

- Stability region for computed using closed-loop simulations
- Based on the vehicle's parameters and the sensor's FOV

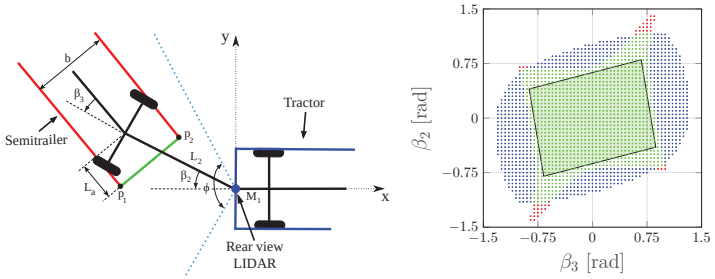

Illustration of sensing geometry (left) and modeled joint-angle constraint (right).

The MPC controller is $\boldsymbol{u}(t) = \boldsymbol{u}_r(t) + \tilde{\boldsymbol{u}}_0$, where $\tilde{\boldsymbol{u}}_0$ is computed by solving the following optimization problem online

$$\begin{aligned} \underset{\tilde{\boldsymbol{u}}_0,\ldots,\tilde{\boldsymbol{u}}_{N-1}}{\text{minimize}} \quad & \tilde{\boldsymbol{x}}_N^T P_N \tilde{\boldsymbol{x}}_N + \sum_{k=0}^{N-1} \left(\tilde{\boldsymbol{x}}_k^T Q \tilde{\boldsymbol{x}}_k + \tilde{\boldsymbol{u}}_k^2\right) \\ \text{subject to} \quad & \tilde{\boldsymbol{x}}_{k+1} = F_k \tilde{\boldsymbol{x}}_k + G_k \tilde{\boldsymbol{u}}_k, \\ & (\tilde{\beta}_{3,k}, \tilde{\beta}_{2,k}) \in \tilde{\mathbb{P}}_k, \quad \tilde{\boldsymbol{u}}_k \in \tilde{\mathbb{U}}_k, \quad k = 0, 1, \ldots, N-1, \\ & \tilde{\boldsymbol{x}}_0 = \tilde{\boldsymbol{x}}(t). \end{aligned} \tag{1}$$

Field experiments using proposed MPC controller

- Implemented on full-scale test vehicle together with Scania CV
- Benchmarked and shown to outperform an LQ controller

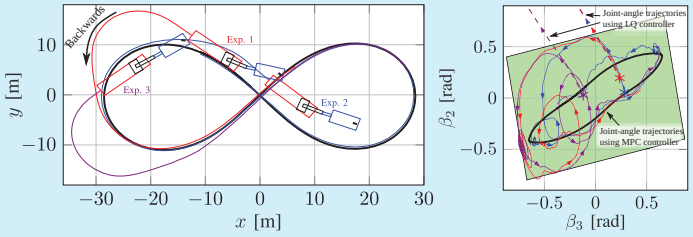

Trajectory planning for multi-steered N-trailer vehicles

Problem: Develop a trajectory planner for multi-steered N-trailer (MSNT) vehicles, composed of a car-like tractor and N trailers with fixed or steerable wheels.

The trajectory planning problem is given by

$$\begin{aligned} \underset{\boldsymbol{u}(\cdot),\, t_G}{\text{minimize}} \quad & J = t_G + \int_0^{t_G} l(\boldsymbol{x}(t), \boldsymbol{u}(t))\, \mathrm{d}t \\ \text{subject to} \quad & \dot{\boldsymbol{x}}(t) = \boldsymbol{f}(\boldsymbol{x}(t), \boldsymbol{u}(t)), \\ & \boldsymbol{x}(0) = \boldsymbol{x}_I, \quad \boldsymbol{x}(t_G) = \boldsymbol{x}_G, \\ & \boldsymbol{x}(t) \in \mathcal{X}_{\text{free}}, \quad \boldsymbol{u}(t) \in \mathcal{U}. \end{aligned} \tag{2}$$

Solution: Combine a lattice-based trajectory planner with an optimal control problem (OCP) step:

- Lattice planner solves combinatorial aspects
- OCP step refines the solution to local optimality

Results for a multi-steered 3-trailer vehicle

- Last trailer is steerable, making MS3T vehicle (blue paths) more flexible than single-steered 3-trailer (SS3T) vehicle (red paths)
- Objective value after OCP step ($\bar{J}_H$) is reduced ($\bar{r}_{\text{imp}}$) compared to state-lattice initialization ($\bar{J}_D$)

Table 1: Results from parallel parking planning scenario.

Vehicle	$\bar{J}_D$	$\bar{J}_H$	$\bar{r}_{\text{imp}}$
SS3T	270.6	122.4	-55%
MS3T	207.9	98.1	-52%

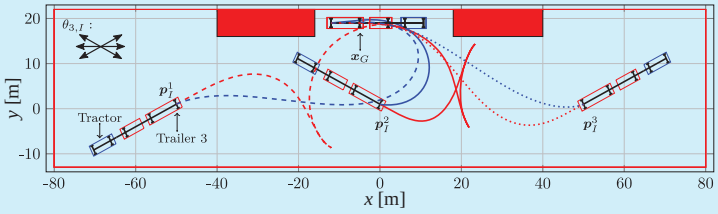

An MPC controller for multi-steered N-trailer vehicles

Problem: Execute a nominal path $(\boldsymbol{x}_r(\cdot), \boldsymbol{u}_r(\cdot))$ for an MSNT vehicle with a small path-following error $\tilde{\boldsymbol{x}}(t) = \boldsymbol{x}(t) - \boldsymbol{x}_r(s(t))$ while satisfying constraints on states $\boldsymbol{x}(\cdot) \in \mathcal{X}$ and control input $\boldsymbol{u}(\cdot) \in \mathcal{U}$.

Solution: An MPC controller similar to (1).

Results for a multi-steered 2-trailer vehicle

- Last trailer is steerable, benchmarked with a single-steered 2-trailer vehicle and with an LQ controller
- MS2T-MPC (blue) has a larger stability region compared to SS2T-MPC (red) and MS2T-LQ (green)
- The trailer steering drastically reduces the transient response of all path-following error states

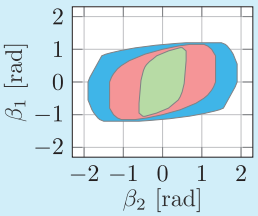

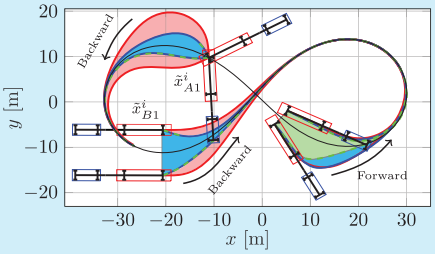

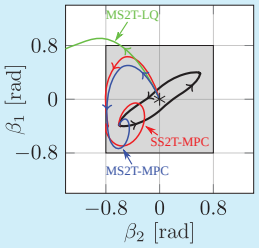

LINKÖPING UNIVERSITY

Autonomous-Vehicle Maneuver Planning Using Segmentation

Pavel Anistratov
Division of Vehicular Systems, Linköping University
pavel.anistratov@liu.se

Introduction & Background

- **Motion planning** is important for autonomous vehicles to allow **safe and reliable operation** under various conditions
- Motion-planning problems could be formulated and subsequently solved as an **optimization problem**
- Computational challenges are often associated with solving optimization problems in **online scenarios**
- Paper [2] shows a method for improving the computational performance by **splitting the full maneuver** into several smaller **maneuver segments**
- Alternating Augmented Lagrangian Method [3] is adapted to allow **parallel computation** of the maneuver segments

Idea of the method

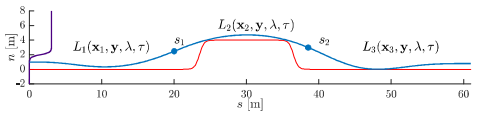

- The method is illustrated for a **double lane-change maneuver** (divided into three segments; can be extended to any number of segments)
- Compact form of the original motion-planning problem:

$$\text{min. } F(\mathbf{x}) \quad \text{s. t. } \mathbf{x} \in S,$$

where $\mathbf{x}$ is the vector of all variables and S is the set of points $\mathbf{x}$ that satisfy the constraints in the formulation

- Augmented Lagrangian function:

$$L(\mathbf{x}, \mathbf{y}, \lambda, \tau) = F(\mathbf{x}) + \lambda^T g(\mathbf{x}, \mathbf{y}) + \frac{\tau}{2} ||g(\mathbf{x}, \mathbf{y})||^2$$

where $g(\mathbf{x}, \mathbf{y})$ are coupling constraints

- Smaller coordinated subproblems (one subproblem corresponds to one maneuver segment)

$$\begin{aligned} &\underset{\mathbf{x}_1 \in S_1}{\text{minimize}} \quad L_1(\mathbf{x}_1, \mathbf{y}, \lambda, \tau) \\ &\underset{\mathbf{x}_2 \in S_2}{\text{minimize}} \quad L_2(\mathbf{x}_2, \mathbf{y}, \lambda, \tau) \\ &\underset{\mathbf{x}_3 \in S_3}{\text{minimize}} \quad L_3(\mathbf{x}_3, \mathbf{y}, \lambda, \tau) \end{aligned}$$

where $\mathbf{x}$ is split into three subvectors

References

[1] P. Anistratov, B. Olofsson, O. Burdakov, and L. Nielsen. Autonomous-vehicle maneuver planning using segmentation and the alternating augmented Lagrangian method. *Submitted manuscript*, 2019.

[2] P. Anistratov, B. Olofsson, and L. Nielsen. Segmentation and merging of autonomous at-the-limit maneuvers for ground vehicles. In *Proceedings of the 14th International Symposium on Advanced Vehicle Control*, Beijing, China, 2018.

[3] G. Galvan, M. Lapucci, T. Levato, and M. Sciandrone. An alternating augmented Lagrangian method for constrained nonconvex optimization. *Optimization Methods and Software*, 2019. doi: 10.1080/10556788.2019.1576177.

Initialization and Segmentation

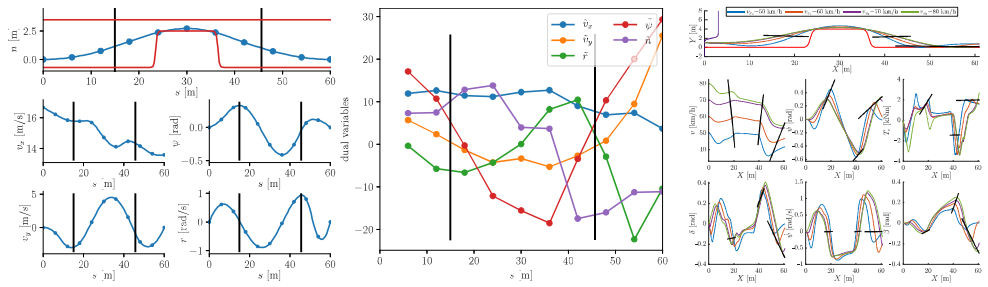

- For rapid convergence, an initialization strategy is used
- Solution for a loose tolerance bound and a small number of discretization points
- Segmentation of the maneuver (black lines) from a vehicle dynamics perspective [2]
- Scaling of the vehicle orientation ψ for $\tilde{\psi}$ to be in the same range as the other dual variables
- Interpolation techniques to initialize the variables $\mathbf{x}$, $\mathbf{y}$, and λ

Computational Approach

- $\mathbf{x}^{k+1}$ is obtained by solving a number of subproblems in parallel
- $\mathbf{y}^{k+1}$ is obtained from minimization of the augmented Lagrangian function, analytical solution is shown in [1]
- Update rules for λ and τ are from [3]
- Proof-of-concept implementation in Python 3.7 using CasADi

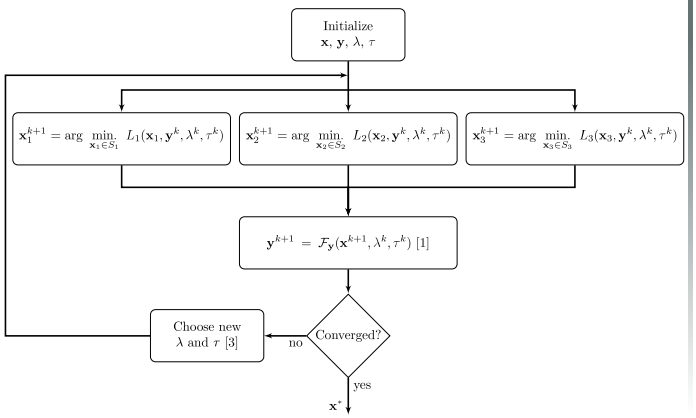

Results

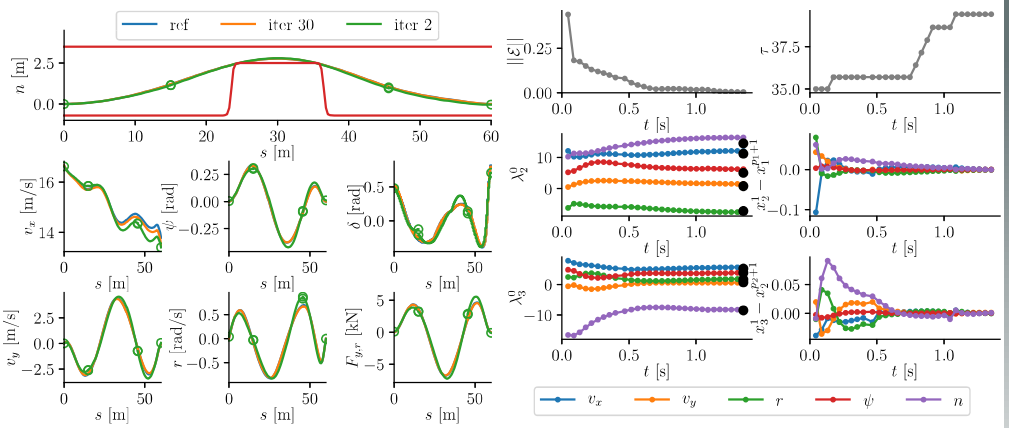

- **All dynamics constraints** at the segment junctions **are relaxed**, $x = \{v_x, v_y, r, \psi, n\}$
- **Segmented–merged maneuver** after 30 alternating iterations (iter 30) is almost **overlaying the full maneuver** (ref)
- Initialization and two alternating iterations (iter 2) take 0.08 s (0.13 s to solve the full problem)
- Less time to declare the solver for maneuver segments in CasADi
- Increasing number of segments to 10 ⇒ 0.023 s solution time for two alternating iterations

Conclusions

- After just a few alternating iterations, there is a **high correspondence** between the full and the segmented–merged maneuvers, obtained with **lower computation times**
- The **scaling** of certain state variables **is important** for achieving rapid convergence of the method

Unsupervised Geometric Deep Learning

Pavlo Melnyk[1], Michael Felsberg[1], Fredrik Kahl[2], Mårten Wadenbäck[1]


[1]Computer Vision Laboratory, Linköping University, Sweden
[2]Computer Vision Group, Chalmers University, Sweden


The Why

- Living in a 3D world, all we perceive visually are its **2D projections**
- Computer vision is to obtain the **underlying 3D** representation
- **Geometric invariance** injected into networks by data-augmentation has many limitations
- **Invariance** that cannot be expressed in the image plane is **unachievable**
- Increased **amount of computations during learning** without adding new real data

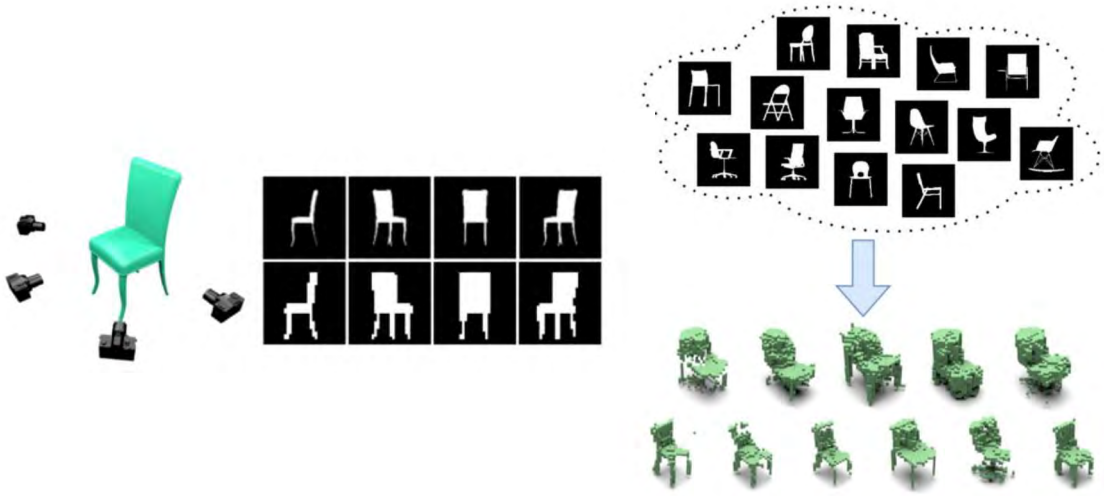

Gadelha, Matheus & Maji, Subhransu & Wang, Rui. (2016). 3D Shape Induction from 2D Views of Multiple Objects.

Szaboó, Attila & Favaro, Paolo. (2018). Unsupervised 3D Shape Learning from Image Collections in the Wild.

The What

- We want a model that is capable of **recovering the 3D** structure of an object given its 2D view(s)
- **Realistic**: textured objects with lighting

The How

- Disentanglement is to be carried out in an **unsupervised** mode
- By **injecting geometric** properties into the model construction
- Using **differentiable rendering**

Optimal Range and Beamwidth for Radar Tracking of Maneuvering Targets Using Nearly Constant Velocity Filters

Per Boström-Rost[⋆†], Daniel Axehill[⋆], William Dale Blair[‡], and Gustaf Hendeby[⋆]

[⋆]Linköping University, [†]Saab Aeronautics, [‡]Georgia Tech Research Institute

Introduction

The problem of tracking a maneuvering target using a mobile radar with limited beamwidth is considered.

- **Dilemma:** How to select the tracking range?
  - The risk of losing track of the target decreases with tracking range
  - The tracking error increases with tracking range
- **Contribution:** A method to optimize the tracking range

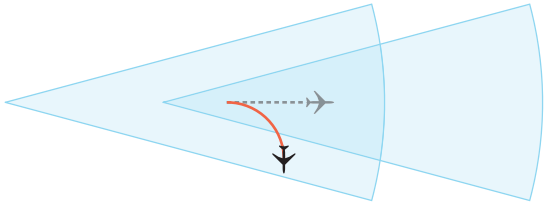

Targets that can turn sharply have to be tracked at a large range to be kept within field of view.

Performance prediction

The target state is estimated using an EKF, and an α-β-filter is used to predict the performance.

- α-β-filtering equations:

$$\hat{\mathrm{x}}_{k|k-1} = \hat{\mathrm{x}}_{k-1|k-1} + T\hat{\mathrm{v}}_{k-1|k-1}$$
$$\hat{\mathrm{v}}_{k|k-1} = \hat{\mathrm{v}}_{k-1|k-1}$$
$$\hat{\mathrm{x}}_{k|k} = \hat{\mathrm{x}}_{k|k-1} + \alpha(y_k - \hat{\mathrm{x}}_{k|k-1})$$
$$\hat{\mathrm{v}}_{k|k} = \hat{\mathrm{v}}_{k|k-1} + \tfrac{\beta}{T}(y_k - \hat{\mathrm{x}}_{k|k-1})$$

- Measurement equation:

$$y_k = h(x_k, s_k) + e_k$$
$$e_k \sim \mathcal{N}(0, \sigma_e)$$

- Performance measure: Maximum mean squared error

$$\mathrm{MMSE} = \sigma_e^2 \left[\frac{2\alpha^2 + \beta(2+\alpha)}{\alpha(4-2\alpha-\beta)} + \frac{1}{\beta^2} \frac{A_{\max}^2 T^4}{\sigma_e^2} \right]$$

Depends on:

- filter gains α, β
- measurement noise σ_e
- time between measurements T
- maximum acceleration of the target $A_{\max}$

Parameter selection

Select tracking parameters such that the MMSE is minimized under the constraint that the probability of losing track of the target is less than a user-defined value, γ.

Optimization problem:

$$\begin{aligned} &\text{minimize} && \mathrm{MMSE} \\ &\text{subject to} && p(\text{lost track}) \leq \gamma \end{aligned}$$

Optimization variables:

- Tracking range r
- Beamwidth θ_{BW}
- Process noise std dev σ_w

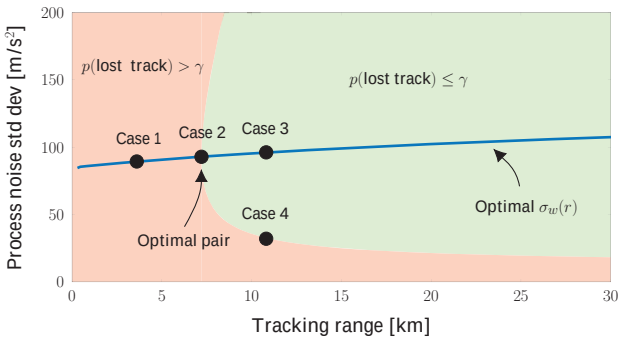

Optimal process noise standard deviation as a function of the tracking range for tracking a 6g-maneuvering target using a radar with five degrees beamwidth.

Experiments

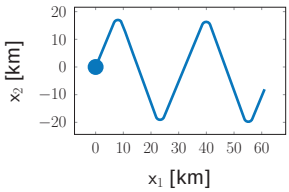

- Track the target using a mobile radar
  - Beamwidth $\theta_{\mathrm{BW}} = 5°$
  - Acceptable track loss probability $\gamma = 0.3\,\%$

- Evaluate pairs of tracking range and process noise standard deviation as indicated in the Parameter selection section.

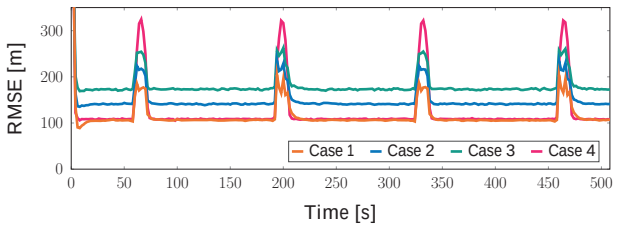

RMSE of the one-step predicted position estimates.

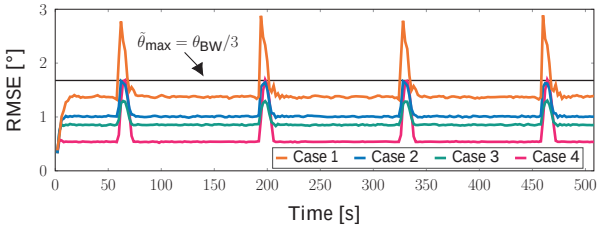

RMSE of the one-step predicted bearing estimates.

Conclusions

Using the proposed method, it is possible to optimize:

- the tracking range (for a given beamwidth)
- the beamwidth (for a given tracking range)

such that the MMSE is minimized, with an acceptable risk of losing track of the target.

P. Boström-Rost, W. D. B. Axehill, Daniel, and G. Hendeby. Optimal range and beamwidth for radar tracking of maneuvering targets using nearly constant velocity filters. To appear in *Proceedings of IEEE Aerospace Conference*, Big Sky, MT, USA, 2020.

Acknowledgments

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

LINKÖPINGS UNIVERSITET

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

SAAB

Automation of efficient, reliable, and elastic predictive control over the cloud

Per Skarin, Ind. PhD, Ericsson and Lund University
Dept. of Automatic Control, Autonomous Clouds and Networks
Supervisors: Karl-Erik Årzen (LU), Johan Eker (Ericsson) and Maria Kihl (LU)

LUND UNIVERSITY

ERICSSON

Motivation & Research Goals

Remote code execution has various uses within control; two of which are offloading heavy computations, and moving executions close to the data. Modern systems use efficient, cost effective, and scalable infrastructure clouds and networks which are becoming increasingly ubiqutous. This project identifies a research gap for critical systems which are to be seamlessly integrated in an era of fog/edge computing. The project studies automation of time sensitive systems over the cloud. The goals are to quantify the potentials and limitations, and to propose control and software strategies incorporating resiliency, reliability and resource management.

Methods

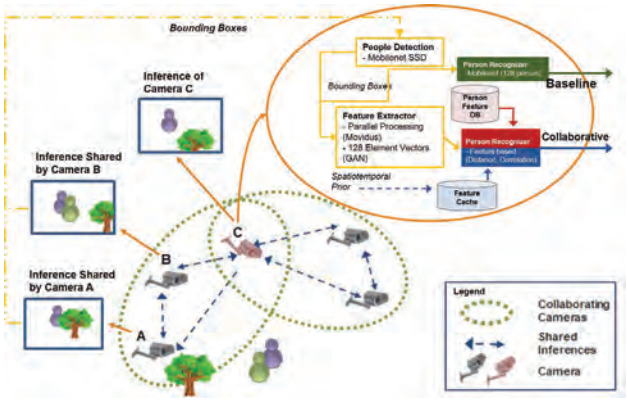

Clouds and fog computing are drivers in the current and future growth of large scale system automation. As industries, transports, and smart cities are increasing their presence in this domain we seek to investigate how real time operated cyber-physical-systems can make use of these ecosystems[1]. We use feedback control, public and private clouds, next generation wireless technology, and available software platforms to provide proof of concept infrastructure, and application benchmarks[2].

$$u^*(x) := argmin \quad \sum_{i=0}^{N-1} l(x_i, u_i) + V_f(X_N)$$
$$\text{s.t.} \quad x_{i+1} = f(x_i, u_i), \quad g(x_i, u_i) \leq 0$$
$$x_0 = x, \quad x_N \in \mathcal{T}$$

To introduce the concepts of elastic controllers we use Model Predictive Control (MPC)[3]. MPC provides flexibility in computational demands which can be traded for performance and formal guarantees. Several parameters can be tuned to allow execution in various contexts: rate of operation, prediction horizon, model complexity, constraints, and feasibility requirements. Extensions to robust control, ML, and large scale collaborative optimal control problems have a natural place in the cloud.

References

[1] Five Challenges in Cloud-Enabled Intelligence and Control
T. Abdelzaher, Y. Hao, K. Jayarajah, A. Misra, S. Yao, P. Skarin, D. Weerakoon and K. Årzén
ACM Transactions on Internet Technology, 2019

[2] Towards Mission-Critical Control at the Edge and Over 5G
Per Skarin, William Tärneberg, Karl-Erik Årzén, Maria Kihl
Best paper award, IEEE Services (EDGE), 2-7 July, 2018, San Fran., CA, USA

[3] Control Over the Edge Cloud - An MPC Example
Karl-Erik Årzén, Per Skarin, William Tärneberg, Maria Kihl
1st Int. Workshop on Trustworthy and Real-time Edge Computing for CPS, Nashville, USA

[4] Cloud-Assisted Model Predictive Control
Per Skarin, Johan Eker, Maria Kihl and Karl-Erik Årzén
2019 IEEE International Conference on Edge Computing (EDGE)

[5] Cloud-based model predictive control with variable horizon
Per Skarin, Johan Eker, Karl-Erik Årzén
Subm. to 21st World Congress of the International Federation of Automatic Control, 2020

Selected Results

A novel distributed cloud platform incorporating IoT and next generation wireless broadband has been developed. A define once, run anywhere software strategy is implemented using a stream programming Platform-as-a-Service framework. Software components can be relocated at any time while keeping closed loop control operational.

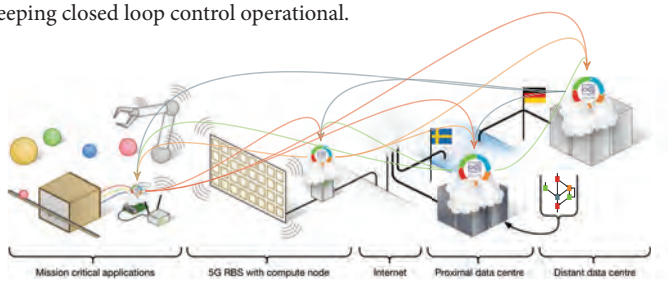

Demonstrated dynamic deployment of time sensitive closed loop control at 20 Hz over the distributed cloud with massive MIMO wireless technology, and IoT components[2][3].

Push the boundaries of real-time control over the cloud. Propose flexible optimal control with variable execution characteristics and quality-elasticity in a frequency range of 10-100Hz.

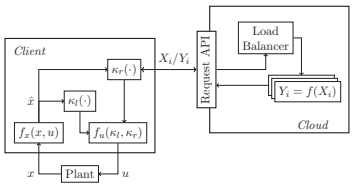

Quality-elastic MPC using *infinite* cloud compute is limited by communication and compute latencies. Graceful degradation and selective execution at the edge manages uncertainty[4].

We propose a formal elastic controller around dual-mode MPC; robust to response loss, and capable of handling an expanding state space.

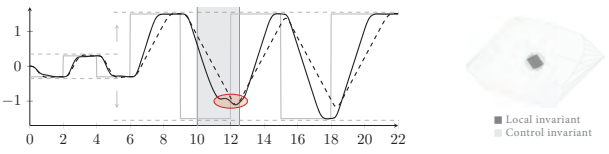

$$f_u(\kappa_l, \kappa_r) = \alpha_i(\vec{u}_k(k-i) - K\hat{x}_k^{lim}|x_{i-1}) + K\hat{x}_k^{lim}|x_{k-1}$$

The variable horizon dual-mode controller provides a formalization of reliable elastic control. Optimal control is obtained using the cloud. Locally, limited fallback performance is available through set-point shaping. Invariant sets are imposed for recursive feasibility, stability and robustness[5].

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

POLYTOPES FOR GRAPHICAL MODELS

S. Linusson and P. Restadh and L. Solus
KTH - Royal institute of technology, Sweden
Email address: linusson@math.kth.se and petterre@kth.se and solus@kth.se

Graphical Models

Graphical Models are used to encode conditional independancy (CI) relations entailed by data-generating distributions.

For a directed acyclic graph (DAG) $G = ([n], E)$, we say that a distribution $\mathbb{P}$ over $X_1, \dots X_n$ is **Markov** to G if $A \perp\!\!\!\perp B | C$ holds in $\mathbb{P}$ whenever A and B are d**-separated** given C.

Example 1. Three DAGs and some of the CI statements they encode and some they do not.

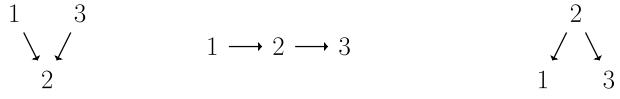

$1 \perp\!\!\!\perp 3 | \emptyset$, $1 \not\perp\!\!\!\perp 3 | 2$ | $1 \perp\!\!\!\perp 3 | 2$, $1 \not\perp\!\!\!\perp 3 | \emptyset$ | $1 \perp\!\!\!\perp 3 | 2$, $1 \not\perp\!\!\!\perp 3 | \emptyset$

If two DAGs encode the same CI statements, we call them **Markov equivalent**, and say they belong to the same **Markov equivalence class** (MEC).

Based on observational data the goal of **DAG model discovery** is the following:

Goal 1. Given a set of observed CI relations $\mathcal{C}$, find the MEC of DAGs that best encode $\mathcal{C}$.

Polytopes

A polytope is the convex hull of a finite set of points in $\mathbb{R}^d$, say $P := \text{conv}\{p_0, \dots, p_k\} \subseteq \mathbb{R}^d$. A **face** of a polytope is a subset maximizing a linear function.

Example 2. Let us consider $\text{CIM}(I_5)$.

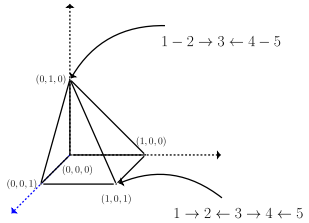

Characteristic Imset Polytope

Definition 1 (Characteristic Imset)**.** Let $D = ([n], E)$ be a DAG. Then we define the characteristic imset $c_D \colon \mathcal{P}([n]) \to \mathbb{R}$ to be

$$c_D(A) := \begin{cases} 1 & |A| \geq 2 \text{ there exists a node } a \in A \\ & \text{with } b \to a \text{ for all } b \in A \backslash \{a\} \\ 0 & \text{Otherwise} \end{cases}$$

Theorem 1. [1] Let D and D' be two DAGs. Then $c_D = c_{D'}$ if and only if D and D' are Markov equivalent.

Then we can consider each imset as a vector in $\mathbb{R}^{2^n - n - 1}$. Let the **characteristic imset polytope** be

$$\text{CIM}(n) := \text{conv}\left\{c_D \colon D \text{ DAG on } n \text{ nodes}\right\} \subseteq \mathbb{R}^{2^n - n - 1}.$$

Let G be an undirected graph and let $\text{AO}(G)$ be the set of acyclic orientations of G. The **characteristic imset polytope over** G is

$$\text{CIM}(G) := \text{conv}\left\{c_D \colon D \in \text{AO}(G)\right\}.$$

If G has n nodes, then $\text{CIM}(G) \subseteq \text{CIM}(n)$. Moreover $\text{CIM}(G)$ is a face of $\text{CIM}(n)$.

Score Based Methods

Some of the best preforming algorithms for DAG model discovery first learn the undirected skeleton and then orient the edges to find the best MEC[2, 3].

This corresponds to first learning a face $\text{CIM}(G)$ of $\text{CIM}(n)$ and then identifying an optimal vertex of $\text{CIM}(G)$.

Goal 2. For a fixed undirected graph G, describe the combinatorial geometry of $\text{CIM}(G)$.

Through a deeper understanding of the geometry of $\text{CIM}(G)$ we hope to device new methodology in DAG model discovery.

*

References
[1] Raymond Hemmecke, Silvia Lindner, and Milan Studený. Characteristic imsets for learning bayesian network structure. International Journal of Approximate Reasoning, 53(9):1336 – 1349, 2012. Fifth European Workshop on Probabilistic Graphical Models (PGM-2010).
[2] Markus Kalisch and Peter Bühlmann. Causal structure learning and inference: A selective review. Quality Technology and Quantitative Management, 11:3–21, 03 2014.
[3] Ioannis Tsamardinos, Laura E. Brown, and Constantin F. Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. Machine Learning, 65:31–78, 2006.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Engineering Trustworthy Self-Adaptive Autonomous Systems

Piergiuseppe Mallozzi
Chalmers University of Technology

Autonomous Systems (AS) are becoming ubiquitous in our society. Some examples are autonomous vehicles, unmanned aerial vehicles (UAV), autonomous trading systems, self-managing Telecom networks and smart factories. Autonomous Systems are based on a continuous interaction with the environment in which they are deployed, and more often than not this environment can be dynamic and partially unknown. AS must be able to take decisions autonomously at run-time also in presence of uncertainty.
It is essential to guarantee that AS act safely in the environment where they are deployed. In other words, the software must work in a reliable manner and must be safe for humans as life may depend on it. We refer to building trustworthy self-adaptive systems meaning to provide evidence that important aspects of the AS are correct. We investigate ways to provide assurance to the software system from the requirements to the actual system. We combine design-time modelling and verification, such as model checking, and runtime assurances techniques such as runtime verification.

Background

Goal models have been used over the years as an intuitive and effective means to capture the designer's intents and their hierarchical structure, mostly following an optative approach oriented to the design objectives. **Contracts**, on the other hand, enable formal requirement analysis in a modular way, rooted in sound representations of the system semantics and decomposition architecture, and offering an indicative approach oriented to the system components and their interactions, which are less explicit in goal models.

Reinforcement learning (RL) algorithms discover policies that can lead AS to achieve their goals in a dynamic and unknown environment. The developer does not specify anymore how the system should act in each possible situation but rather the RL algorithm can achieve an optimal behavior by trial and error. Once trained, the AS will be capable of taking decisions and performing actions autonomously while still learning from the environment.

Trustworthiness: AS act safely in the environment where they are deployed.

Research Directions Overview

Analyzing the trustworthiness of autonomous systems from different angles.

Leveraging on ***formal methods*** *techniques*, analyze
system requirements,
system model,
and system execution traces.

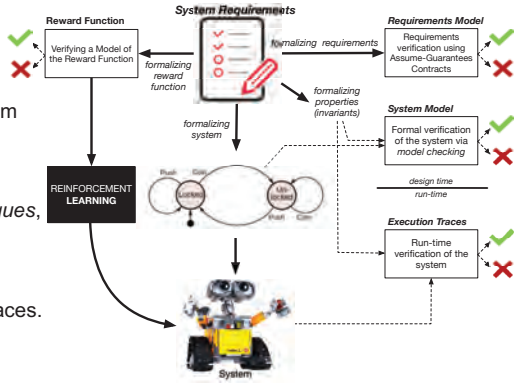

Preliminary Results

1. **CoGoMo**: a requirement engineering framework combining the goal-oriented and contract-based paradigms. We propose a new formal model, termed *contract-based goal tree* that represents goal models in terms of hierarchies of contracts. Based on this model, we propose algorithms that leverage contract operations and relations to check goal consistency, completeness, and perform incremental and hierarchical refinement of goals by mapping goals to promising component architectures. [1]

2. **WiseML**: an approach that combines machine learning with runtime monitoring to detect violations of system invariants in the actions execution policies. [2][6]

3. **MoVeMo**: a technique to formally express reward functions in a structured way; this stimulates a proper reward function design and as well enables the formal verification of it. [4]

Selected Publications

[1] P. Mallozzi, I. Romero, P. Nuzzo, A. Sangiovanni-Vincentelli, P. Pelliccione **"CoGoMo: Incremental Refinement of Goal Models with Contracts"**, in submission, 2020

[2] P. Mallozzi, E. Castellano, P. Pelliccione, G. Shneider, K. Tei **"A Runtime Monitoring Framework to Enforce Invariants on Reinforcement Learning Agents Exploring Complex Environments"**, Robotics Software Engineering (RoSE'19)

[3] P. Mallozzi, P. Pelliccione, A. Knauss, C. Berger, and N. Mohammadiha, **"Autonomous vehicles: state of the art and state of practice"** in Automotive Software Engineering: State of the Art and Future Trends, 2017.

[4] P. Mallozzi, R. Pardo, V. Duplessis, P. Pelliccione, and G. Schneider, **"MoVEMo - A structured approach for engineering reward functions"** in IEEE International Conference on Robotic Computing, 2018.

[5] P. Mallozzi, M. Sciancalepore, and P. Pelliccione, **"Formal verification of the on-the-fly vehicle platooning protocol"** in International Workshop on Software Engineering for Resilient Systems. Springer, 2016.

[6] P. Mallozzi, P. Pelliccione, and C. Menghi, **"Wiseml: Keeping intelligence under control"** in International Workshop on Software Engineering for Cognitive Services, 2018

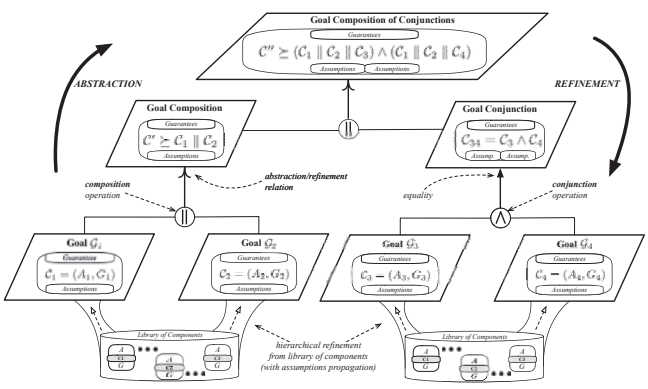

Figure 1. **CoGoMo** framework

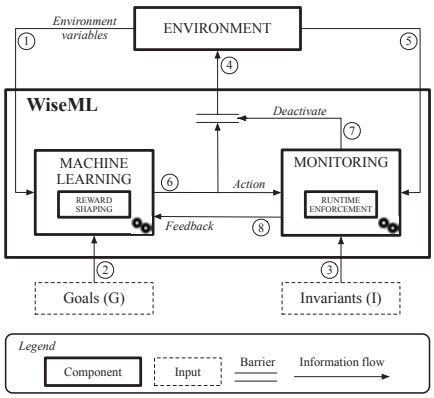

Figure 2.**WiseML**

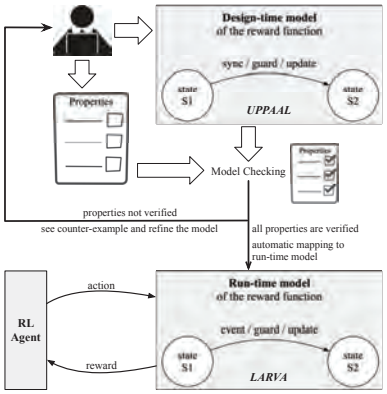

Figure 3. **MoVEMo** framework

Reinforcement Learning in Continuous Spaces with Interactively Acquired Knowledge-based Models

Quantao Yang, Johannes Andreas Stork, Todor Stoyanov
Autonomous Mobile Manipulation Lab, AASS, Örebro University, SE-70182 Sweden
quantao.yang@oru.se

Manipulation tasks that seem trivial to a human can be hard to learn for robots, especially from scratch without initial human demonstration, due to high sample complexity. This project aims to address the problem of robotic manipulation with reinforcement learning algorithms in **continuous** action space [1]. Specifically, the interaction between the semantic physical model and reinforcement learning methods will be investigated in depth.

Background & Motivation

Industrial robots are pre-programmed to execute specific tasks

- Human safety [2]
- Environments vary from one task to another
- Low efficiency due to fixed trajectory

Research Goal & Questions

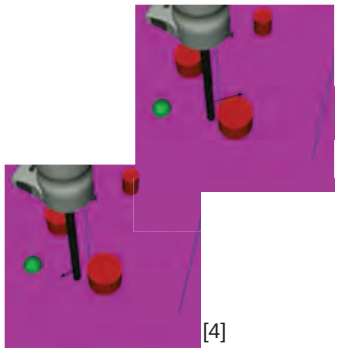

[4]

- How do we apply **constraints** to manipulation settings? [3]
- In what way can a sequence of constraints **encode** semantic information?
- How do we construct an **interaction model**?
- How do we improve **sample efficiency**?

Methods & Preliminary Results

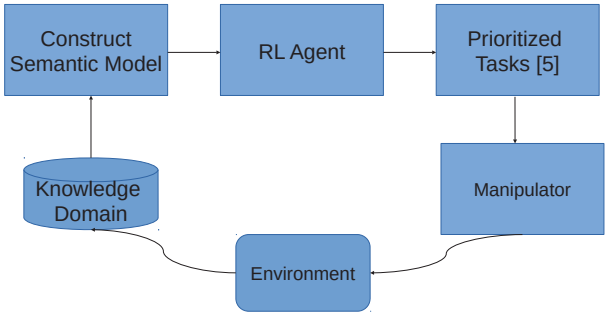

Reinforcement Learning

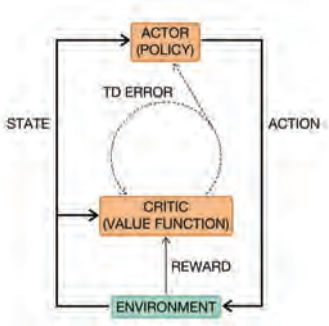

- **Symbolic/Hierarchical** RL
- **Actor-Critic** RL [6]
- Interaction between **semantic physical model** and RL agent
- Integrate prior knowledge of the **dynamic model** of a robot manipulator
- Transfer learning on new tasks

[1] Smruti Amarjyoti. Deep reinforcement learning for robotic manipulation-the state of the art. arXiv preprint arXiv:1701.08878, 2017.
[2] Gal Dalal, Krishnamurthy Dvijotham, Matej Vecerik, Todd Hester, CosminPaduraru, and Yuval Tassa. Safe exploration in continuous action spaces. arXiv preprint arXiv:1801.08757, 2018.
[3] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In Proceedings of the 34th International Conference on Machine Learning-Volume 70, pages 22–31. JMLR. Org, 2017.
[4] Björn Magnusson and Måns Forslund, Safe and efficient reinforcement learning, bachelor thesis, Örebro University.
[5] Adrien Escande, Nicolas Mansard, and Pierre-Brice Wieber. Hierarchical quadratic programming: Fast online humanoid-robot motion generation. The International Journal of Robotics Research, 33(7):1006–1028, 2014.
[6] Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil An-thony Bharath. A brief survey of deep reinforcement learning. arXivpreprint arXiv:1708.05866, 2017

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Continuous Optimization of Software

Rasmus Ros, Lund University
Computer Science
Supervisors: Per Runeson and Elizabeth Bjarnason

Motivation & Research goals

Modern service based software is developed in short cycles with continuous user feedback in controlled experiments (aka. A/B tests). My research interests are in building AI tools that support this process and empirical evaluations of their use. The tools use machine learning, combinatorial optimization, and domain specific languages to specify software configurations that should be optimized with user data. Evaluation should be performed in realistic environments involving real end user data at industry sites.

Research Output

The project was started with two exploratory studies. Firstly, a literature study of the published literature [1, 4] which identified important topics on A/B testing and optimization techniques. Secondly, the applicability of AI solutions to do A/B testing was studied with interviews at an e-commerce company [3].

These studies was used as input in designing the toolkit Combo that does online optimization with user data. Using it involves specifying a search space of the software configuration and constraints (logic, arithmetic constraints, etc.). Constraints allow the developers to filter bad variations and for personalization. This is exemplified to the right where the variable *Width* can be specified as *S* for mobile devices or *M* for tablets.

The optimization problem is formalized as combinatorial multi-armed bandit (CMAB). The problem is solved with probabilistic machine learning, such as generalized linear models, neural networks, or random forests, in combination with Thompson sampling. Evolutionary algorithms and CMAB hybrids were also explored [2].

The toolkit was built to handle use cases at the company Apptus that does e-commerce algorithms. Search auto complete widgets was selected as a first step to optimize on because it is an isolated component. The implementation was evaluated with simulations on real data and is yet to be put into production.

Future work include, further evaluation of the toolkit [5] with more use cases, better handling of changes in the search space, and further interviews on the applicability of the techniques in [3].

Combo Toolkit

```
val ac = model("Auto complete") {
    nominal("Width", "S", "M", "L")
    optionalNominal("Terms", 1, 3, 5, 10)
    val cards = optionalNominal("Cards", 1, 3, 5)
    impose { width["S"] equivalent !cards }
    model(cards) {
        val tc = bool("Two-column")
        val hz = bool("Horizontal")
        impose { cards[1] implies !hz }
        impose { width["L"] equivalent (tz or hz) }
    }
    impose { "Terms" or "Cards" }
}
val algo = linearNeuralBandit(ac)
```

Simplified model specification in the Combo DSL of the use case below.

The Combo [5] toolkit supports specifying a search space of potentially thousands of variables and constraints in an embedded DSL and has several online decision algorithms for optimizing the model with user data.

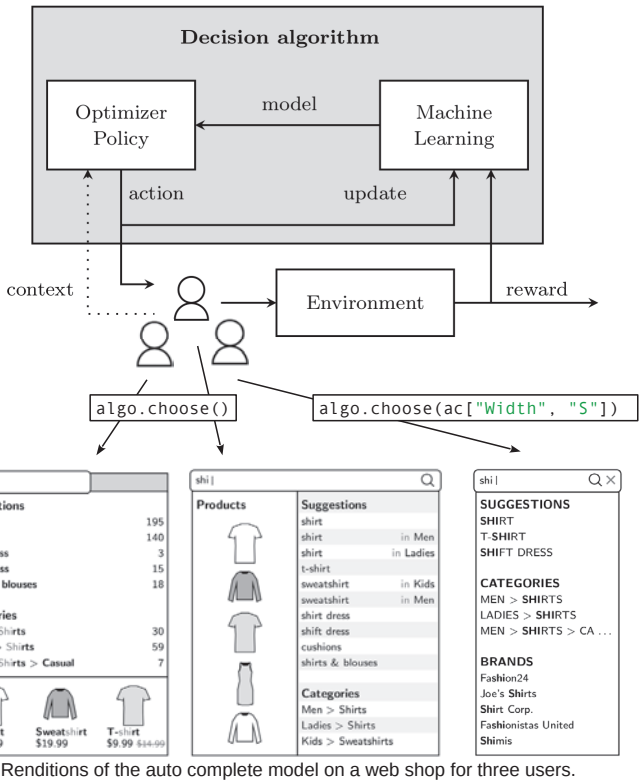

Renditions of the auto complete model on a web shop for three users.

References

1. Ros, R., & Runeson, P. 2018. Continuous Experimentation and A/B Testing: A Mapping Study. RCoSE.
2. Ros, R., Bjarnason, E., & Runeson, P. 2017. Automated Controlled Experimentation on Software by Evolutionary Bandit Optimization. SSBSE.
3. Ros, R., & Bjarnason, E. 2018. Continuous Experimentation Scenarios: A Case Study in e-Commerce. Euromicro SEAA.
4. F. Auer, R. Ros, L. Kaltenbrunner, P. Runeson, & M. Felderer. 2019. Status and Open Challenges in Continuous Experimentation. *In submission.*
5. R. Ros, M. Hammar. Combo: A Toolkit for Continuous Optimization of Software. 2019. *In submission.*

Reinforcement Learning for robotic manipulation of deformable objects

Rita Laezza

Developing foundational methods applicable to a wide-ranging class of objects, for deformation/shape control.

Background & Motivation

- Robotic manipulation of deformable objects has been largely overlooked when compared with rigid objects.
- Strategies developed for rigid objects cannot be directly applied to their non-rigid counterparts.
- There is a lack of general approaches that apply to different deformable object types.
- Sensing of deformable objects is challenging.
- Modeling of deformable objects is also challenging.

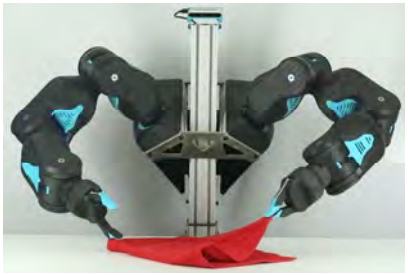

Research Goal & Questions

- The initial research focus will be limited to deformable linear objects (DLOs).
- For DLOs with large strain, it is interesting to study deformation control, where the goal is to achieve a target plastic deformation.
- Due to effects of elastic deformation, there may be a *springback* effect that determines the final configuration.

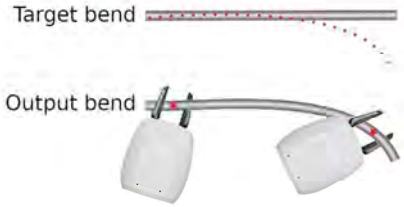

The main research questions to answer at this stage of the project are:

- Is it possible to learn a deformation trajectory?
- Can these trajectories account for the *springback* effect?
- Can the grasping points be included into the decision making?

Methods & Preliminary Results

- Initially, reinforcement learning (RL) algorithms will be trained in simulation environment.
- The robot will not be included in the simulation, instead the resulting policy will be expressed as a trajectory of one or two controlled points.
- Model-free and model-based RL methods will be explored.
- Results from simulation will be used as a starting point for learning with the real robot.

.

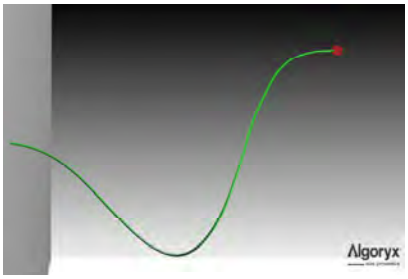

Roadmap & Milestones

There will be several stages with growing complexity, in order to progress faster into the implementation of RL algorithms.

Stage 1	Stage 2
• In simulation • 1 grasping point • Perfect information state	• In simulation • 2 grasping points • Perfect information state

Stage 3	Stage 4
• In simulation • 2 grasping points • Imperfect information state	• With real robot • 2 grasping points • Imperfect information state

At each stage, there will be two additional cases studied:

1. Choice of grasping points.
2. *Springback* prediction after release.

CHALMERS UNIVERSITY OF TECHNOLOGY

Algoryx POWERED BY AGX DYNAMICS

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Simulation and Manipulation of Deformable Objects using Deep Learning

Robert Gieselmann

robgie@kth.se

ROYAL INSTITUTE OF TECHNOLOGY

Why deformable objects?

Many real-world applications involve deformable objects. For instance:

- Manipulating textiles or food items
- Predicting soft tissue deformation in medical applications
- Soft robotics

However, the mechanical behavior of deformable objects is difficult to predict due to their infinite dimensional configuration space.

Differentiable Physics

Engineers have developed constitutive models to describe the deformation of a material under certain conditions. Solving the corresponding differential equations requires numerical methods like FEM, which are computationally demanding and usually not practical for real-time applications.

Deep neural networks are powerful nonlinear function approximators. We can employ a neural network to learn a physics engine for deformable objects from data. Due to their differentiability, we can use them in an inverse setting for control and manipulation.

References

[1] Maziar Raissi, Paris Perdikaris and George E. Karniadakis *Physics Informed Deep Learning (Part I): Data-driven Solutions of Nonlinear Partial Differential Equations* CoRR , 2017

[2] Michael Lutter, Christian Ritter and Jan Peters *Deep Lagrangian Networks: Using Physics as Model Prior for Deep Learning* ICLR , 2019

[3] Sam Greydanus, Misko Dzamba and Jason Yosinski *Hamiltonian Neural Networks* NeurIPS , 2019

[4] Alvaro Sanchez-Gonzalez, Nicolas Heess, Jost Tobias Springenberg, Josh Merel, Martin Riedmiller, Raia Hadsell and Peter Battaglia *Graph Networks as Learnable Physics Engines for Inference and Control* PMLR , 2018

[5] Angelina Wang, Thanard Kurutach, Kara Liu, Pieter-Abbeel and Aviv Tamar *Learning Robotic Manipulation through Visual Planning and Acting* RSS, 2019

What we know from Physics

There is a vast amount of physical knowledge that is currently not being used in deep learning models. Partial differential equations which describe the mechanical behavior of a material can be derived from continuum mechanics.

Example: A beam is attached between to walls and subject to a force applied on the surface.

Given a material model (e.g. Neo-Hookean) we obtain the following mixed boundary value problem for the elasto-static case.

$$\begin{aligned} \text{div} \cdot \sigma &= 0 \text{ in } \Omega \\ u &= 0 \text{ on } \Gamma_D \\ \sigma \cdot n &= \tau \text{ on } \Gamma_N \end{aligned}$$

Example: In continuum mechanics, the deformation gradient F is the derivative of the deformed state x w.r.t. the undeformed state X.

$$F = \frac{\delta x}{\delta X} = 1 + \frac{\delta u}{\delta X}$$

An infinitesimal volume cannot shrink to a point, i.e. zeros volume. In mathematical terms that means $det(F) > 0$.

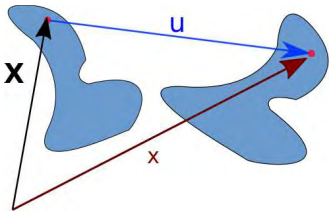

By incorporating physics knowledge in our neural network model, we seek to improve sample-efficiency, generalization and physical plausibility of the predictions.

How do we incorporate those physics priors?

One way to make the neural network respect the laws of physics is to add penalty terms to the loss function which minimize the model violation [1]. Automatic differentiation packages, as used in deep learning frameworks, allow us to compute the derivatives of the output w.r.t. to the input variables.

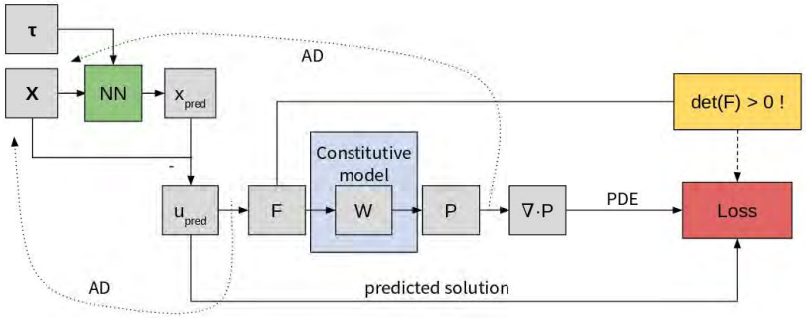

Another approach is to directly incorporate priors into the network architecture. The idea is to impose hard constraints that assure physical plausibility. This was done in [2] and [3] which drew inspiration from Lagrangian and Hamiltonian mechanics respectively.

A fully general method based on graph networks was presented in [4]. A structural inductive bias is introduced by representing objects with nodes and relations with edges. During training several update functions are learned which encode the body dynamics, interaction dynamics and global properties.

In a preliminary experiment we used graph networks to predict the deformation of a soft beam which is subject to an external force.

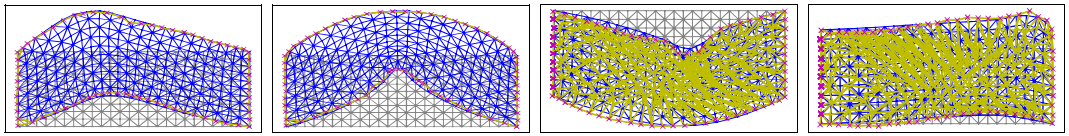

Blue: Ground truth deformation, yellow/purple: predicted deformation, gray: undeformed state

Future Work:

In future work we will investigate how to incorporate prior knowledge into deep generative models to enable physically plausible predictions. In particular, we would like to explore the potential of Generative Adversarial Networks for learning plannable representations for deformable objects (similar to [5]).

Acknowledgement

This work was supported by WASP, the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program.

Motion Planning for Self-Driving Long and Multi-Body Vehicles

Rui Oliveira, KTH
Division of Decision and Control Systems

Motivation and Research Goals

- Long vehicles (buses) and multi-body vehicles (tractor-trailers) require a significant different way of driving.
- The vast majority of motion planning research for self-driving vehicles focuses on passenger vehicles.
- We extend motion planning methods from passenger vehicles to buses and tractor-trailers.

Tractor-trailers

We study suitable modelling and driving objectives for articulated vehicles, such as tractor-trailers within motion planning approaches. The existence of multiple vehicle bodies introduces new challenges to the motion planning task. Current research does not successfully address this type of vehicles in urban driving scenarios.

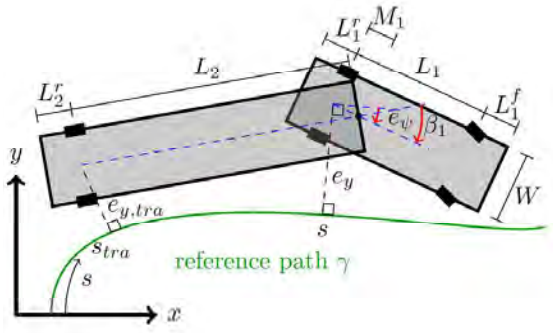

Intermediate results:

- Designed approximation techniques for the tractor-trailer road-aligned vehicle model
- Proposed and studied in detail a set of candidate optimization objectives, while showing that current approaches are not suitable for tractor-trailer vehicles
- Tested the planner in simulated challenging scenarios

Buses

We present a method that takes full advantage of the particular characteristics of buses, namely the overhangs, an elevated part of the vehicle chassis, that can sweep over curbs.

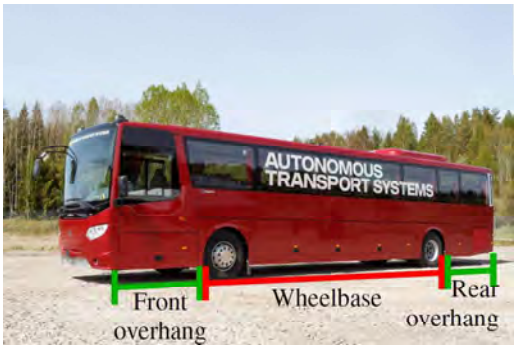

Unlike other motion planning approaches, our method exploits curbs and other sweepable regions, which a bus must often sweep over in order to manage certain maneuvers.

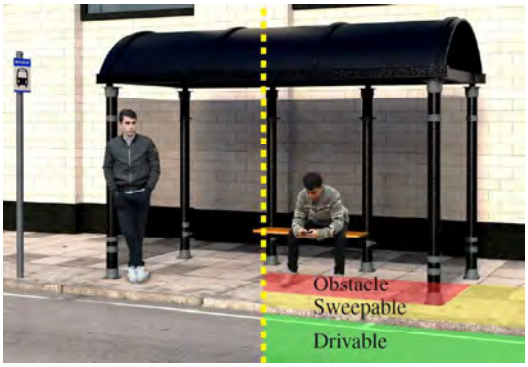

The approach is formulated as an optimal control problem, and solved using a Sequential Quadratic Programming approach.

Results:

- Study of new driving objectives (optimization metrics) suitable for bus driving
- Definition of new region classification schemes and vehicle body constraints
- Tackled the task of bus driving in urban environments by mimicking bus driver behavior

References

[Buses] Rui Oliveira, Pedro F. Lima, Gonçalo Collares Pereira, Jonas Mårtensson, and Bo Wahlberg. "Path Planning for Autonomous Bus Driving in Highly Constrained Environments." In 2019 IEEE Intelligent Transportation Systems Conference (ITSC)

[Tractor-trailers] Rui Oliveira, Oskar Ljungqvist, Pedro F. Lima,, and Bo Wahlberg. "Optimization-Based On-Road Path Planning for Articulated Vehicles." Submitted to 2020 21st IFAC World Congress

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning Contact-rich Manipulation Skills

Shahbaz A. Khader, KTH, ABB

Power and productivity for a better world™

Description:

Manipulation tasks that involve force interactions between the robotic manipulator and the manipulated object are perhaps the most difficult class of tasks that can be performed by a robot. Instead of manually solving the underlying control problem and coding it up in software, we look at how such control policies or skills can be learned by the robot autonomously.

Background & Motivation

Contact-rich manipulation tasks such as parts insertion, fitting etc., can be formulated as control policies or skills. A human worker develops such a skill after receiving formal instructions and subsequent practicing. So we ask the question: can a robot, provided with a high-level representation of the task, practice on its own and learn such a skill?

Research Goal

The goal is to find practical reinforcement learning (RL) methods for learning to control robotic manipulator performing contact-rich manipulation tasks.

Interesting aspects:

1. Model-based vs model-free methods
2. Model learning for contact tasks
3. Policy search vs online optimization
4. Role of simulation in RL
5. Safe learning

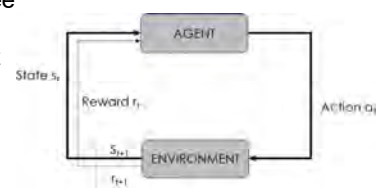

Probabilistic Model learning and Prediction

- ❑Learn a system of switching dynamics models (Gaussian Process Regression - GPR)
- ❑Probabilistic long-term prediction with discontinuous state evolution

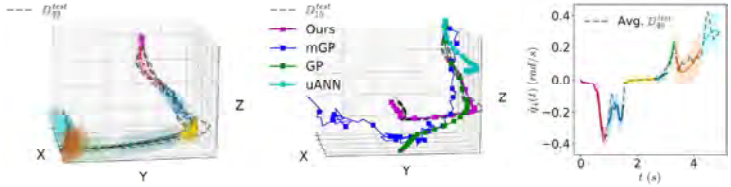

* Khader, S.A., Yin, H., Falco, P. and Kragic, D., 2019. Probabilistic Model Learning and Long-term Prediction for Contact-rich Manipulation Tasks. *arXiv preprint arXiv:1909.04915*. Submitted to RA-L/ICRA

mGP: Calandra, R., Peters, J., Rasmussen, C.E. and Deisenroth, M.P., 2016, July. Manifold Gaussian processes for regression. In *2016 International Joint Conference on Neural Networks (IJCNN)* (pp. 3338-3345). IEEE.

uANN: Chua, K., Calandra, R., McAllister, R. and Levine, S., 2018. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems* (pp. 4754-4765).

Stability Guaranteed Policy Learning

- ❑Learn Variable Impedance Control (VIC) policy

$$\tau_{IC} = \boldsymbol{f}(\boldsymbol{s}, \dot{\boldsymbol{s}}) = -\bar{\boldsymbol{S}}(\boldsymbol{s}, \dot{\boldsymbol{s}})\big(\boldsymbol{s} - \bar{\boldsymbol{\mu}}_{\boldsymbol{s}}(\boldsymbol{s}, \dot{\boldsymbol{s}})\big) + \\ - \bar{\boldsymbol{D}}(\boldsymbol{s}, \dot{\boldsymbol{s}})\big(\dot{\boldsymbol{s}} - \bar{\boldsymbol{\mu}}_{\dot{\boldsymbol{s}}}(\boldsymbol{s}, \dot{\boldsymbol{s}})\big) + \bar{\boldsymbol{f}}_{\boldsymbol{\tau}}(\boldsymbol{s}, \dot{\boldsymbol{s}})$$

- ❑Parameter space likelihood ratio

$$\nabla_\theta J(\theta) = \int_H p(h|\theta) \sum_{t=1}^{T} \nabla_\theta p(a_t|s_t, \theta) r(h) dh$$

$$\nabla_\rho J(\rho) = \int_\Theta \int_H p(h, \theta|\rho) \nabla_\rho \log p(h, \theta|\rho) r(h) dh d\theta$$

- ❑Stability guaranteed exploration

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Configurable observability for 5G orchestration

Simon Lindståhl

KTH Division of Decision and Control Systems, Ericsson Research

Main supervisors: Alexandre Proutiere (KTH), Andreas Johnsson (Ericsson)

Motivation

In modern networks, many different types of services, with different service requirements such as latency or reliability, are run in tandem. The practice of allocating network resources, physical or virtual, to different services in an effective way is known as *orchestration*. For orchestration systems to have full observability, many measurements are needed, but every measurement incurs cost to the network. This work, started in autumn 2019, is a collaboration between KTH and Ericsson with the purpose of investigating the tradeoff between performance gain, as provided by measurements, and measurement costs, when applied to 5G orchestration systems.

Measurements in 5G network slicing

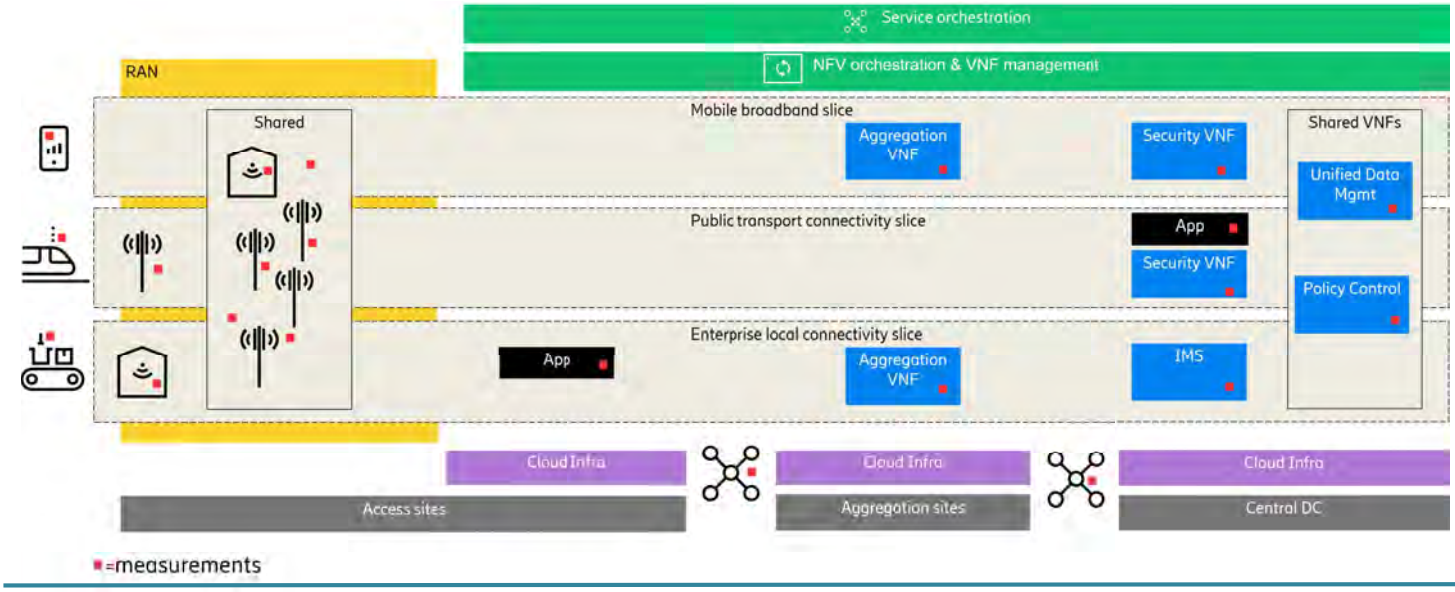

Challenges

Investigate how models, trained in environments with full observability, perform in environments with limited observability.

Ensuring 5G orchestration ML-model performance in operational environments without harmful exploration of the state space.

Study the impact of network slice admission control and reconfiguration on QoS, QoE and SLA.

Study automatic configuration of observability functions in the network to limit resource usage, while maintaining model performance.

Contact

Simon Lindståhl
Ericsson Research
simon.lindstahl@ericsson.com

Andreas Johnsson
Ericsson Research
andreas.a.johnsson@ericsson.com

Alexandre Proutiere
KTH, Decision and Control Systems
alepro@kth.se

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Equivalent G/G/1 Modeling for Server Systems with Redundant Requests

Tommi Nylander, Johan Ruuskanen, Karl-Erik Årzén, Martina Maggio
Department of Automatic Control, Lund University

Introduction

In cloud computing, cloning is used to speed up processing of requests. The basic idea is that instead of sending requests to only one server, the requests are cloned and sent to multiple servers simultaneously. The response to the request is the first completed one, and when this happens, the other clones are cancelled. Cloning can yield substantial improvements to the performance of data centers, as demonstrated in [1].

Significant performance modeling improvements were proposed by Joshi et. al. in [2]. In our work, we extend and further generalize these results.

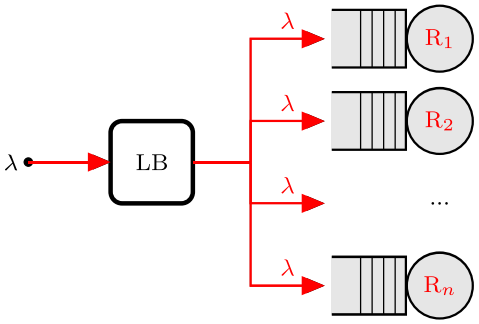

Our Contributions:

- *We generalize* the i.i.d. models of today to allow for *any* inter-arrival or service time distributions. We thus allow for both heterogeneity and dependencies under *any* queuing discipline.
- *We analyze* server systems under cloning and compute the optimal cloning factor.
- *For more complex systems*, we provide a co-design method for joint synthesis of cloning strategy and load-balancing policy.

Model

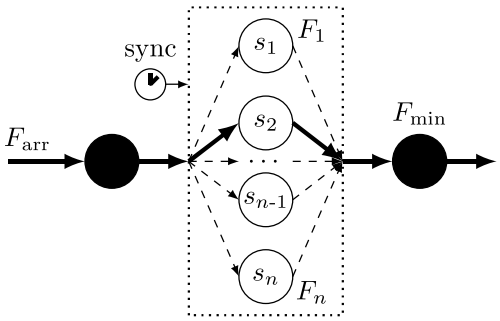

- **Synchronized service** assumes clones r^c to all original requests r^o to be sent to all n servers, and requests to be cancelled immediately.
- We *equivalently* model the cloned server system as **one** server, with inter-arrival and service time distributions F_{arr} and $F_{\min}$.
- $F_{\min}$ can be determined as

$$F_{\min}(x) = 1 - \prod_{i=1}^{n} \{1 - F_i(x)\} \tag{1}$$

Examples

- Heterogeneous exponentials with rates μ_i:

$$F_{\min}(x) = 1 - \prod_{i=1}^{n} \{1 - F_i(x)\} = 1 - \prod_{i=1}^{n} e^{-\mu_i x} = 1 - e^{-\sum_{i=1}^{n} \mu_i x} \tag{2}$$

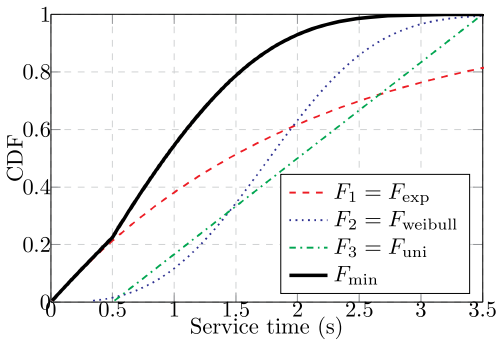

- Heterogeneous distributions: $F_{\min}(x)$ computed numerically.

Server Systems

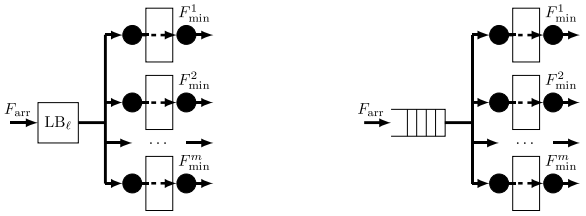

- Servers guaranteeing synchronized service can form **clusters**.
- Co-design of load-balancing policy ℓ and cloning factor c_f required.

Evaluation

- Simulations performed using service time distributions X^c from [3]:

$$X^c = Z^o \cdot S^c, \tag{3}$$

with Z^o as shared task size and S^c server slowdown.

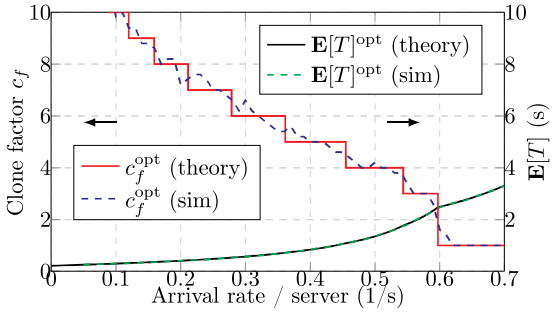

- **Clone-to-All**: Theoretical vs simulated optimal cloning factors.

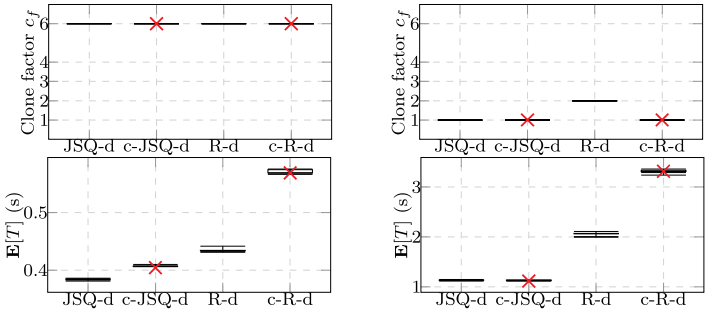

- **Co-designs**: Theoretical vs simulated optimal cloning factors.

References

[1] G. Ananthanarayanan, A. Ghodsi, S. Shenker and I. Stoica. "Effective Straggler Mitigation: Attack of the Clones". In: *NSDI* 2013.
[2] G. Joshi, E. Soljanin and G. Wornell. "Efficient replication of queued tasks for latency reduction in cloud systems". In: *Allerton* 2015.
[3] K. Gardner, M. Harchol-Balter and Alan Scheller-Wolf. "A Better Model for Job Redundancy: Decoupling Server Slowdown and Job Size". In: *MASCOTS* 2016.

Situational Learning and Decision Making for AVs

Truls Nyberg, supervised by
Laura Dal Col and Christoffer Norén, ATS Pre-Development & Research, Situational Awareness and Decision Making, Scania CV AB
Jana Tumová and Patric Jensfelt, Division of Robotics, Perception and Learning (RPL), KTH Royal Institute of Technology

Background & Motivation

To drive safely, a driver needs to perceive, predict and comprehend its environment. In [1], this ability is defined for human decision makers in dynamic systems as **Situation Awareness**. This project strives to establish situation awareness for the autonomous system in order to enable risk-aware and human-like decisions. We define the **risk** of a state as a combination of the **severity** of its consequences and its **likelihood** of occurrence.

Scania CV and KTH Royal Institute of Technology are currently conducting research in the fields of agent prediction, path planning, and control, see [2], [3] and [4]. This research aims at bridging the ongoing projects in agent prediction with motion planning.

Simulation Setup

Prototypes of algorithms will be developed and demonstrated in simulation environments such as the open-source **CARLA** or Scania's proprietary alternative with a full autonomous software stack, seen in Figure 2.

Figure 2: Predictions of future states in a Scania simulation environment.

Research Goals

The goal of this research project is to further bridge the autonomous vehicles' sensing and perception capabilities with its decision making. This will be done in the following three steps:

1. Create a **situation-aware** model of the environment and its future states.
2. Quantify **risks** and give probabilistic guarantees in the environment.
3. Quantify **desires** in the environment and demonstrate **real-time** short-term **decision making** in urban driving scenarios.

The items can be summarized into the following research question: *How can we form a situation-aware model of the environment with quantified risks and desires, in order to take safe and efficient real-time short-term decisions in urban autonomous driving?*

Scenarios

Consider the scenario represented in Figure 1. It captures a wide range of the challenges arising when driving in **urban** environments, e.g. blocked road lanes, agents in **occluded** areas, and agents possibly bending or breaking the **rules of the road**. Autonomous vehicles will need to **negotiate** risk and efficiency in these scenarios to be accepted in public traffic.

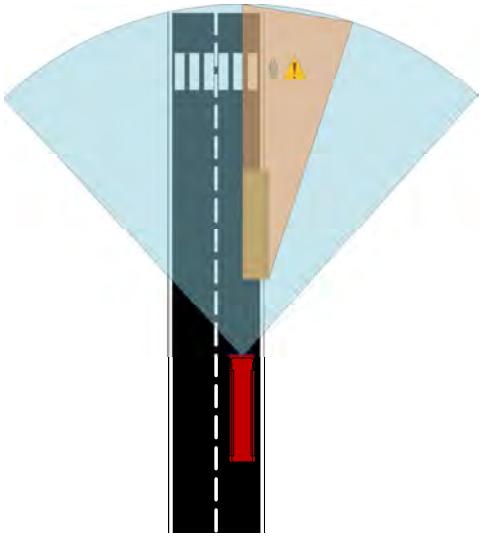

Figure 1: High-risk scenario where the ego vehicle in red needs to perform an overtake with part of the crosswalk occluded.

Experimental Setup

The platform and possible test track for demonstrations can be seen in Figure 3.

Figure 3: Scania's autonomous bus and the city area at the test track AstaZero Borås, Sweden.

References

[1] Mica R Endsley. Toward a Theory of Situation Awareness in Dynamic Systems. *HUMAN FACTORS*, page 33, 1995.

[2] Joonatan Mänttäri and John Folkesson. Incorporating Uncertainty in Predicting Vehicle Maneuvers at Intersections With Complex Interactions. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 2156–2163, June 2019. ISSN: 1931-0587.

[3] Rui Oliveira, Marcello Cirillo, Jonas Mårtensson, and Bo Wahlberg. Combining Lattice-Based Planning and Path Optimization in Autonomous Heavy Duty Vehicle Applications. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 2090–2097, June 2018. ISSN: 1931-0587.

[4] Gonçalo Collares Pereira, Pedro F. Lima, Bo Wahlberg, Henrik Pettersson, and Jonas Mårtensson. Linear Time-Varying Robust Model Predictive Control for Discrete-Time Nonlinear Systems*. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 2659–2666, December 2018. ISSN: 0743-1546.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM
KTH VETENSKAP OCH KONST
SCANIA

Shared Situational Awareness under Complex Traffic Scenarios

Vandana Narri, Scania & KTH

Christoffer Norén
Laura Dal Col
ATS Research, Situational Awareness and Decision Making, Scania CV AB

Karl Henrik Johansson
Jonas Mårtensson
Division of Decision and Control Systems, Electrical Engineering and Computer Science, KTH

Motivation & Research goals

- This project focuses on automating operations of heavy-duty vehicles efficiently and safely under collaborative, yet dynamically changing situations in highways and urban roads.
- The objective of this project is to model, formalize, and analyse a shared situational awareness framework for the extended vehicles, i.e., connected vehicles and infrastructure.
- 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.

V2V and V2X communication.

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 (V2X) communication based on 5G or IEEE 802.11p standards, can help gather more information about the environment, and address the shortcomings of CAV sensors.
- The main research areas are connectivity (enabled by V2X and V2V), cooperative driving, situational awareness and traffic flow optimization.

Problem Statement

- How should a CAV comprehend and anticipate the environment and various road users behavior in it?
  - How can a CAV obtain and fuse data from local and external sensors to improve situational awareness?
  - How can the framework adapt to dynamically changing scenarios?
  - How can a CAV take action, based on the varying degree of decisions?

Methods

Multi-agents sensing modeling
- Decentralized, centralized or distributed network.
- Coordination and control of multi-agent systems.
- Stability, robustness, and performance of multi-agent dynamic systems.

Algorithm for shared situational awareness
- Fusion of the data from local sensors and external sensors.
- Efficient integration of partial observations.
- Sensor data : using heterogenous/ diverse sensors

Interplay with decision-making and planning
- The control architecture includes multiple hierarchical layers.
- High comprehension of the environment should be able to improve the performance of layers.

Development Environment

- Matlab/Simulink - Learn & Prototype.
- Open-sourse simulator - Develop for Simulation Demo.
- Scania autonomous vehicle and the test facility AstaZero, Borås, Sweden.

Experimental Scenario

- In the proposed scenario, the CAV (yellow) approaches the four-way crossing and its sensor view (yellow shaded area) is affected by occlusions (orange shaded areas) by parked cars and buildings. A cyclist (orange circle) is approaching behind the corner from the right, which cannot be detected with the on-board sensors. Furthermore, a vehicle (purple) is approaching the intersection, but its position is outside the host vehicle's sensor range.
- The proposed framework should allow the CAV to overcome the sensing limitations.

The proposed experimental scenario.

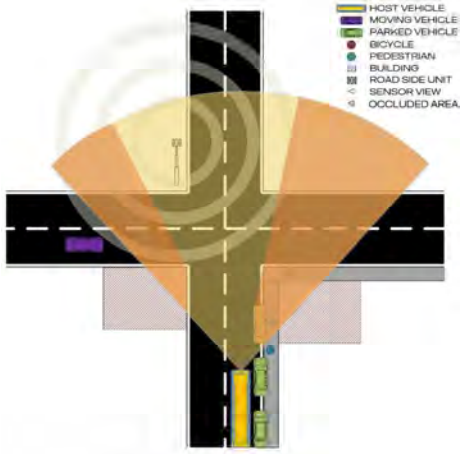

MATLAB simulation of a traffic scenario using "Driving Scenario Designer" toolbox.

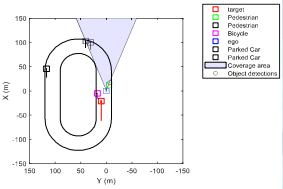

The test facility AstaZero, Borås, Sweden for experiments.

Example of an open-source simulator for autonomous driving scenario - CARLA simulator.

Scania research prototype vehicle.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Yaw-Moment Control At-the-Limit of Friction Using Individual Front-Wheel Steering and Four-Wheel Braking

Presented at the 9th IFAC Symposium on Advances in Automotive Control (AAC), Orléans, France, 2019 [1]

V. Fors[1], B. Olofsson[1,2], and L. Nielsen[1]

1: Linköping University, 2: Lund University

Introduction

- Control of autonomous vehicles in safety-critical situations by combined braking and steering action.
- Closed-loop controller for safety-critical maneuvers at-the-limit of friction developed in [2].
- Incorporate yaw control by use of the Modified Hamiltonian Algorithm (MHA) [3].

Combined Tire Forces At-the-Limit of Friction

- The steady-state longitudinal and lateral tire forces are both functions of the longitudinal and lateral tire slips.
- Combinations of different tire slip and many parameters make a full tire model difficult to use in real-time applications.
- Assuming the desired tire-forces are at the friction-ellipse limit, the modeling is simplified.

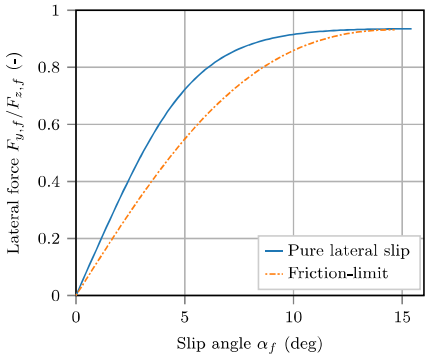

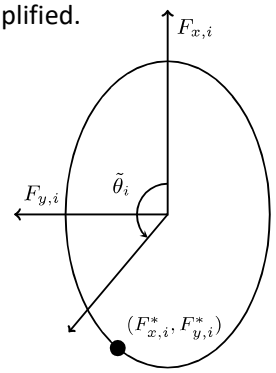

Friction-limit tire model and friction-ellipse assumption enable explicit computation of desired slip angles and braking forces.

Modified Hamiltonian Algorithm

Control allocation achieved by minimization of a Hamiltonian function H.

- To include yaw-moment control, the Hamiltonian to minimize is $H = p_x F_x + p_y F_y + \lambda M_z$
- λ is adapted online such that the yaw moment M_z tracks a desired yaw moment M_z^{d} as $\lambda \rightarrow \lambda + S(M_z - M_z^{\mathrm{d}})$
- Cascading the Hamiltonian H to individual wheels gives $H_i = \cos(\tilde{\theta}_i)\, F_{x,i} + \sin(\tilde{\theta}_i)\, F_{y,i}$

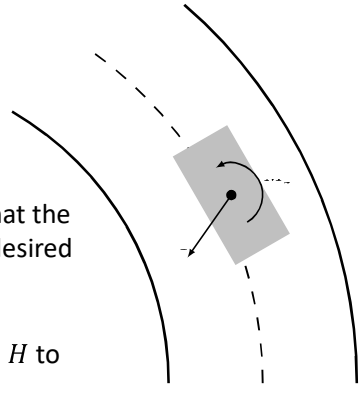

Control Layout

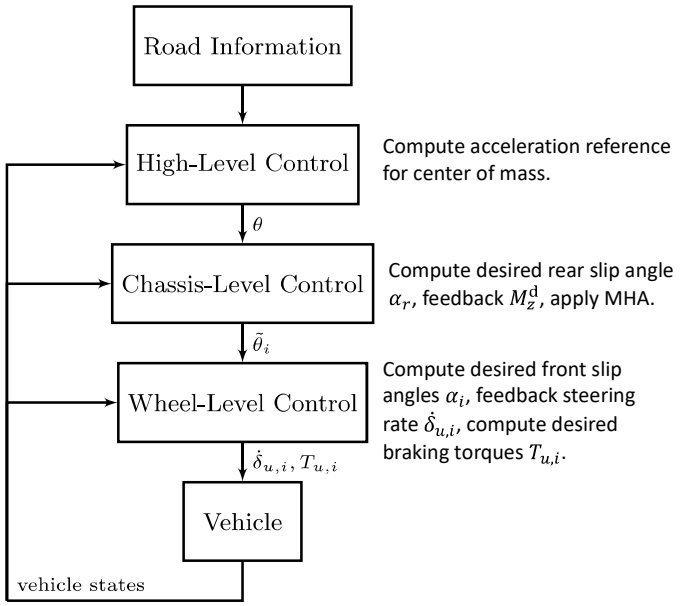

Results

Scenario: Minimize deviation from center of lane, driving at 110 km/h when the car enters a turn with radius 90 m.

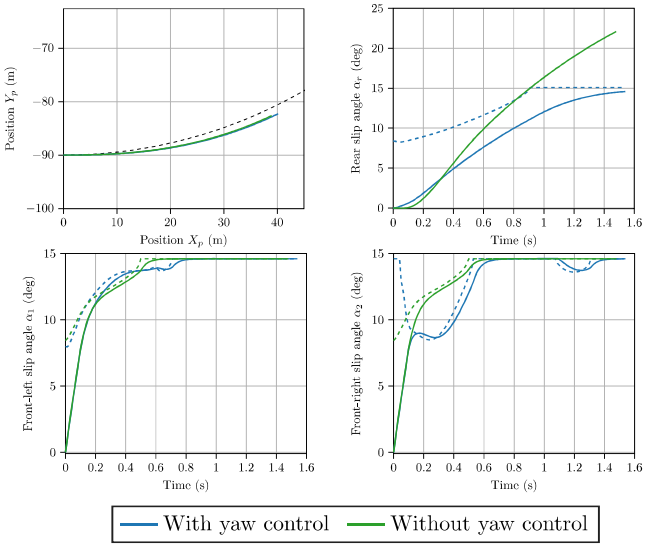

References

[1] V. Fors, B. Olofsson, and L. Nielsen, "Yaw-moment control at-the-limit of friction using individual front-wheel steering and four-wheel braking," in *9th IFAC Symposium on Advances in Automotive Control (AAC)*, vol. 52, no. 5, Orléans, France, 2019, pp. 458–464.

[2] V. Fors, B. Olofsson, and L. Nielsen, "Slip-angle feedback control for autonomous safety-critical maneuvers at-the-limit of friction," in *14th International Symposium on Advanced Vehicle Control (AVEC18)*, Beijing, China, 2018.

[3] Y. Gao, M. Lidberg, and T. J. Gordon, "Modified Hamiltonian algorithm for optimal lane change with application to collision avoidance." *MM Science J.*, no. MAR 2015, pp. 576–584, 2015.

Voronoi Boundary Classification: A High-Dimensional Geometric Approach via Weighted Monte Carlo Integration

Vladislav Polianskii, Florian T. Pokorny
{vpol, fpokorny}@kth.se

ROYAL INSTITUTE OF TECHNOLOGY

Method description

We propose a geometric approach to classification based on the idea of computing integrals of a weight function over Voronoi-Cell boundaries. We utilize a Monte Carlo based approach that enables us to compute these integrals even for very high-dimensional image data without the need to explicitly compute Voronoi Diagrams. The resulting classification method has the following main characteristics:

- An intuitive, nonparametric, geometrically inspired method
- Classification does not require training and new classes can be added on the fly
- Experiments indicate robustness to imbalanced data
- Classification results on par with random forrests
- Convergence to 1NN for sequence of increasingly peaked weight functions

Main idea

Voronoi Boundary Classification result is based on the highest *Voronoi boundary rank* r_i.
$r_i(t) = \sum_{j:c_j=i} I_j$, where
$I_j = \int_{V(t)\cap V(x_j)} w(x)\mathrm{dVol}$.

Here t is the test point, j spans indices of known data points, c_j are the labels, $V(\cdot)$ denotes a corresponding cell and w is a weight function.

Classification example in 2D

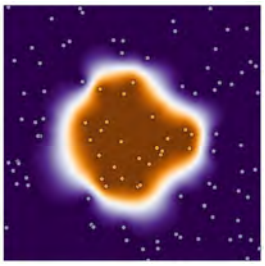

Classification qualities

Variation of σ parameter

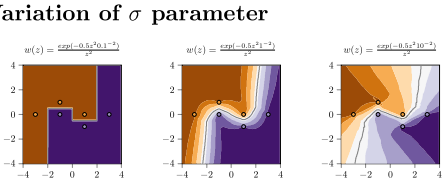

$\sigma = 0.1$ $\sigma = 1$ $\sigma = 10$

We show VBC provably converges to 1NN for small σ.

Invariance to data imbalance

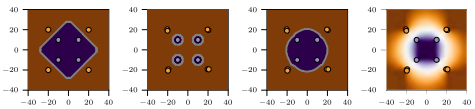

SVM (8p) SVM (12p) VC (8 or 12 points)

Comparison to SVM shows robustness to imbalanced data.

Voronoi rank computation

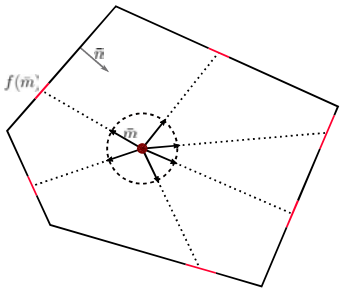

Spherical integral

$$I_j = \int_{f_j^{-1}(V(t)\cap V(x_j))} w(f_j(m)) \frac{\|f_j(m) - t\|^{d-1}}{|\langle m, n\rangle|} \mathrm{dVol}_{\mathbb{S}^{d-1}}$$

Monte Carlo integration

$$I_j = V_j \lim_{T\to\infty} \frac{1}{T} \sum_{i=1}^{T} w(f_j(m_i)) \frac{\|f_j(m_i) - t\|^{d-1}}{|\langle m_i, n\rangle|}$$

Evaluation

Accuracy

	Frogs	MNIST	CIFAR-10
VBC(10k)	**.986±.0002**	.969±.0003	**.494±.0011**
1NN	.982	.969	.354
SVM	.903	.944	.440
RF(1500)	.974±.0004	**.972±.0004**	**.495±.0016**

Our algorithm (VBC) utilizes additional geometric information about Voronoi Cell structure as compared to 1NN, has a fundamentally different approach compared to Random Forests (RF) and SVM and is training-free, but performs on par with RF.

Accuracy convergence

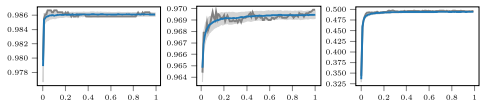

Frogs MNIST CIFAR-10

Confidence measurement

The considered confidence value is $\frac{r^{(1)}-r^{(2)}}{r^{(1)}+r^{(2)}}$, where $r^{(1)}$ and $r^{(2)}$ are top two Voronoi boundary ranks.

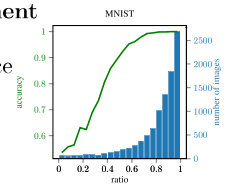

Links

Additional information can be found at `people.kth.se/~vpol` and `people.kth.se/~fpokorny`.

Source code available at `github.com/vlpolyansky/voronoi-boundary-classifier`

Acknowledgement

This work has been supported by the Knut and Alice Wallenberg Foundation.

Convergence of integrals

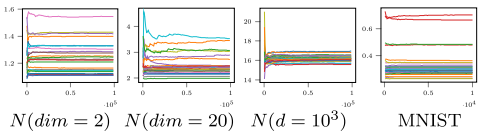

$N(dim = 2)$ $N(dim = 20)$ $N(d = 10^3)$ MNIST

We demonstrate rapid convergence of Monte Carlo integral approximation.

Single cell convergence with 100 trials

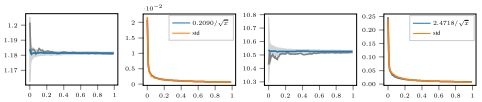

$N(dim = 2)$ MNIST

Plots 1 and 3: avg integral value with error bounds.
Plots 2 and 4: obtained convergence rate and $O(T^{-0.5})$ fit.

Extra topics

Families of decaying weight functions

Our approach is applicable for a class of functions decreasing with distance from the test data point, that are infinite at 0 and integrable over possibly infinite cell boundaries. Specifically, we consider:

$$\hat{w}(z) = \frac{e^{-0.5\sigma^{-2}z^2}}{z^d}$$

Convergence to one-nearest-neighbor

We prove that our classification method converges to one-nearest neighbor classification for a sequence of increasingly peaked functions:

Theorem. *Consider a labeled dataset* $D = \{(x_1, c_1), \ldots, (x_n, c_n)\}$ *of data points* $x_i \in \mathbb{R}^d$ *and corresponding class labels* $c_i \in \{1, 2, \ldots, k\}$ *and a test point* $t \in \mathbb{R}^d$, *where* $t \neq x_i$ *for all* $i \in \{1, \ldots, n\}$. *Assume that* t *only has a single nearest neighbor among* $\{x_1, \ldots, x_n\}$. *Consider a sequence of weight functions* $\{w_n\}_{n=1}^{\infty}$, $w_n : \mathbb{R}_+ \to \mathbb{R}_+$ *which are each monotonically decreasing and where*

$$\lim_{n\to\infty} \int_{z_2}^{+\infty} \frac{w_n(z)}{w_n(z_1)} z^{d-2}\, dz = 0 \quad \textit{for all} \quad 0 < z_1 < z_2$$

Then, for sufficiently large n, *the class* $VBC(t|D, w_n)$ *assigned by the Voronoi Boundary Classifier of* t *is equal to the class* $1NN(t|D)$ *assigned by the nearest neighbor classifier for* t.

Asymptotic Complexity

N: # of known points M: # of test points
T: # of iterations D: space dimensionality

Base version:
$O(MND + TMN + TMD + TND)$
simple matrix operations, GPU implementation

Approximate version:
$O(\mathrm{prep}(N, D) + TM \log \frac{\mathrm{diam}(X)}{\varepsilon_0}(\mathrm{nn}(N, D) + D))$
relies on approximate nearest neighbor implementation

Reinforcement Learning and Statistical Learning in Control

Yassir Jedra[1]

Linear Quadratic Regulator

Dynamics Consider a linear dynamical system defined for $t \geq 0$ as follows

$$x_{t+1} = Ax_t + Bu_t + w_t \tag{1}$$

with an initial state value $x_0 = 0$, and where the noise sequence $\{w_t\}_{t\geq 0}$ is a martingale difference process w.r.t a filtration $\{\mathcal{F}_t\}_{t\geq 0}$ (i.e. $\mathbb{E}\left[w_t|\mathcal{F}_{t-1}\right] = 0$). Additionally $u_t \in \mathbb{R}^p$ is the control vector, and depends on the sequence $x_0, u_0, \ldots, u_{t-1}, x_{t-1}$. $x_t \in \mathbb{R}^d$ is the state vector and is $\mathcal{F}_{t-1}$-measurable. The matrices (A, B) are unkwon and belong to some uncertainty set $\mathcal{D} \subseteq \mathbb{R}^{d\times d} \times \mathbb{R}^{d\times p}$.

Cost function Consider a quadratic cost function $c : \mathbb{R}^d \times \mathbb{R}^p \to \mathbb{R}_+$ defined by:

$$c(x, u) = x^\top Q x + u^\top R u,$$

where $Q \succeq 0$ and $R \succ 0$. The cumulative cost at each time step t is defined as

$$\mathcal{C}_t^\pi \triangleq \sum_{s=0}^{t} c_{s,\pi_s} = \sum_{s=0}^{t} x_s^\top Q x_s + u_s^\top R u_s,$$

where $\pi = \{\pi_t\}_{t\geq 0}$ is a sequence of $\mathcal{F}_t$-measurable functions, such that for all $t \geq 0$, $\pi_t(I_t) = u_t$, where I_t is the said history and is defined as follows $I_t = \{x_0, u_0, x_1, u_1, \ldots, x_t\}$. We denote the cost $c_{s,\pi_s} = c(x_t, u_t)$. Finally, we call π a control policy. The goal is to choose a control policy π while not knowing the true parameters (A, B).

Regret The notion of regret quantifies the cost of not knowing the true environment. Here we define this notion for a control policy in the LQR problem. We start by looking at the optimal solution of the problem when (A, B) are known. When we have an infinite horizon $t \to \infty$, the optimal control law is

$$\hat{u}_t^* = Kx_t,$$

where K is the gain matrix and is defined as

$$K = -(B^\top P B + R)^{-1} B^\top P A.$$

and the matrix P is the solution to the Algebraic Riccati Equation

$$P = A^\top P A - A^T P B (B^\top P B + R)^{-1} B^\top P A + Q.$$

Denote this optimal control policy by π^*. Now, the regret of a control policy π can be defined in the following manner

$$\mathcal{R}_T^\pi \triangleq \sum_{t=0}^{T} c_{t,\pi_t} - \sum_{t=0}^{T} c_{t,\pi^*}.$$

Note the aforementioned regret is stochastic, thus one needs to precise in what sense regret needs to be analysed. For a given policy π, there exists two types of analysis. First, one derives a ***regret upper bound with high probability***, which consists of finding a rate $\psi(.,.)$, such that for any $\delta \in (0, 1)$, for any T, we have:

$$\mathbb{P}\Big(\hat{\mathcal{R}}_T^\pi \leq C\psi(T, \delta)\Big) \geq 1 - \delta.$$

for some positive constant C. Second, one derives an ***expected regret upper bound***, which consists of finding a rate $\psi(.)$, such that for T, we have

$$\mathbb{E}\left[\hat{\mathcal{R}}_T^\pi\right] \leq C\psi(T),$$

for some positive constant C. Finally, one's objective is to find a control policy π that achieves an optimal rate $\psi(.)$.

Sample Complexity Lower Bounds of Linear Time Invariant Systems

Consider the linear dynamical system (1) where the noise sequence $w_0, w_1, \ldots$ is i.i.d. with a standard gaussian distribution.

Question: How many samples are required to identify the true parameters (A, B) ?

Definition. *((ϵ, δ)-locally-stable algorithms) An algorithm is (ϵ, δ)-locally-stable in (A, B) for some $\epsilon > 0$ and $\delta \in (0, 1)$ if there exists a finite time τ such that for all $t \geq \tau$ and all A' and B' such that $\|[A\ B] - [A'\ B']\|_F \leq 3\epsilon$,*

$$\mathbb{P}_{A',B'}(\|[\hat{A}_t\ \hat{B}_t] - [A'\ B']\|_F \leq \epsilon) \geq 1 - \delta.$$

Theorem (see [1]). *For any matrices A and B, for all $\epsilon > 0$ and $\delta \in (0, 1)$, the sample complexity $\tau_{A,B}$ of any (ϵ, δ)-locally-stable algorithm in (A, B) satisfies:*

$$\lambda_{\min}\Big(\mathbb{E}_{A,B}\Big[\sum_{t=0}^{\tau_{A,B}-1} \begin{bmatrix} x_t \\ u_t \end{bmatrix} \begin{bmatrix} x_t^\top & u_t^\top \end{bmatrix}\Big]\Big) \geq \frac{1}{2\varepsilon^2} \log\Big(\frac{1}{2.4\,\delta}\Big).$$

References

[1] Y. Jedra and A. Proutiere. Sample Complexity Lower Bounds for Linear System Identification. Arxiv, 2019.

[1] Division of Decision and Control Systems, School of Electrical Engineering and Computer Science, KTH, Stockholm, Sweden.