POSTER CATALOGUE
WASP Winter Conference 2020

Poster Session 3
Wednesday 15 January 09.30-10.30

Poster no.	First name	Last name	Title of poster
3	Abdelrahman	Eldesokey	Uncertainty of Sparse Depth Data in CNNs
7	Alexander	Nilsson	Software-based Side-Channel Attacks on Crypto
11	Alfred	Åkesson	ComPOS: Composing Oblivious Services
15	Cesar	Soto Valéro	A Study of Bloated Dependencies in the Maven Ecosystem *
19	Christian	Rosdahl	Efficient Learning of Dynamical Systems
23	Dhasarathy	Parthasarathy	Explainable robustness for self-driving vehicles
27	-	-	-
31	Gustaf	Waldemarson	Photon Mapping Superluminal Particles
35	Gustav	Häger	Vision for robots *
39	Haorui	Peng	Massive MIMO Pilots Scheduling on a MAC-PHY Split Architecture
43	-	-	-
47	Jens	Henriksson	Out-of-Distribution Robustness for DNN's
51	Joakim	Johnander	Recurrent Neural Networks for Perception
55	Joakim	Brorsson	Keeping central authorities out of your business
59	Joel	Scheuner	Performance-Optimized Cloud Applications
63	Johan	Ruuskanen	Statistical Inference for the Self-adaptive Cloud
67	Johan	Karlsson	Exact remodeling of optimization programs with applications to autonomous driving
71	John	Törnblom	Formal Verification of Learning-based Software in Safety-Critical Systems
75	Joris	van Rooij	STAMINA: Stream Processing in the AMI
79	Karin	Stacke	Detection of colon cancer metastases in lymph nodes through deep learning
79	Apostolia	Tsirikoglou	Detection of colon cancer metastases in lymph nodes through deep learning
83	-	-	-
87	Lissy	Pellaco	Wireless Link Adaptation with outdated CSI -- a hybrid data-driven and model-based approach
91	Lucas	Brynte	Pose Proposal Critic
95	Martin	Larsson	Upgrade Methods for Stratified Sensor Network Self-calibration
99	Martin	Lindvall	Designing support for lymph node tumor detection: Assisted search for rare phenomena
103	Md Sakib Nizam	Khan	SoK: Ambient Assisted Living Systems and their Privacy/Security Considerations
107	Mina	Ferizbegovic	Robust LQ-controllers using application oriented exploration
111	-	-	-
115	Vidit	Saxena	Online Learning with Linear Constraints
119	Olivier	Moliner	Multi-Camera Extrinsic Calibration from Human Pose
123	Pegah	Nikbakht Bideh	Developing Tools and Analyze Methods for Secure Software Update
127	Peter	Varnai	Prescribed Performance Control Guided Policy Improvement for Satisfying Signal Temporal Logic Tasks
131	Sarit	Khirirat	Compressed Gradient Methods for Hessian-Aided Error Compensation
135	Sólrún Halla	Einarsdóttir	Towards Big Theory Exploration *
139	Sule	Anjomshoae	Intelligible Explanations in Intelligent Systems
143	Xuechun	Xu	Explicit-duration Hidden Markov Model on DNA Base-calling

*) Poster not available in this catalogue

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

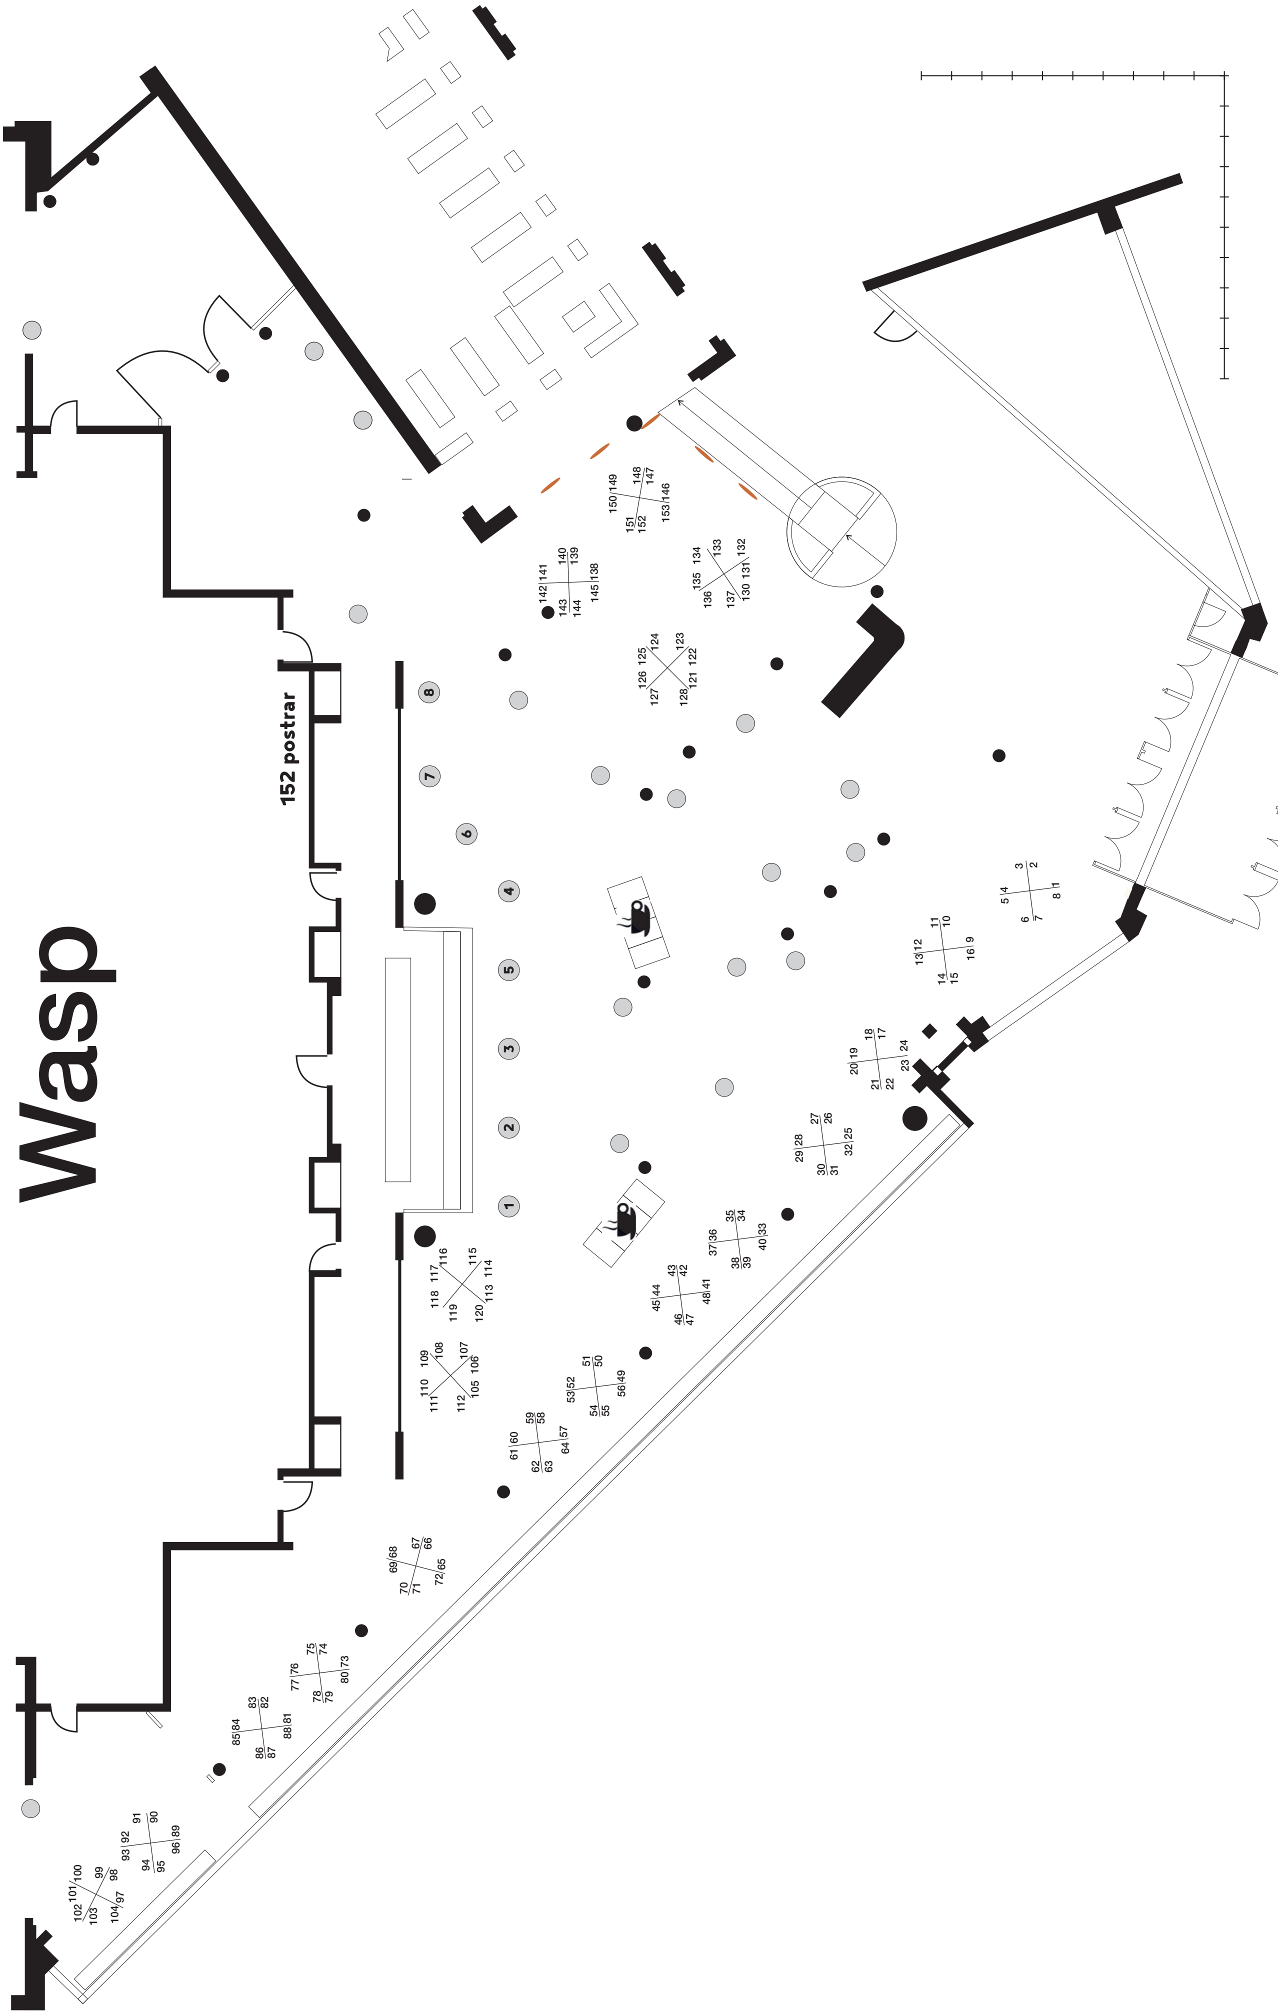

LINKÖPING
KONSERT &
KONGRESS

Uncertainty of Sparse Depth Data in CNNs

Abdelrahman Eldesokey, Linköping University

Computer Vision Laboratory with Michael Felsberg

LINKÖPING UNIVERSITY

Problem

- Depth sensors such as LiDARs are gaining popularity in robotics and autonomous driving applications.
- Data generated by LiDARs are very sparse due to the active nature of the sensor.
- A key challenge is how to efficiently handle this sparse data within CNNs and produce dense depth maps?

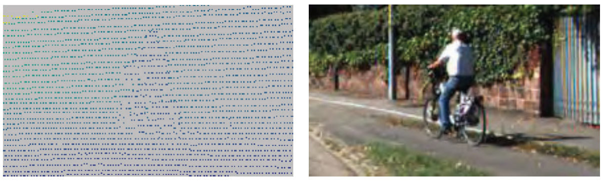

A projected depth map from a Velodyne LiDAR sensor to the left for the scene on the right.

Normalized CNNs

- We extended the classical normalized convolutional [1] that is very efficient in handling sparse data to a CNN layer in [2].
- The normalized CNN layer receives the sparse input + a pixel-wise input confidence and produces a dense output + an output confidence.

$$\mathbf{Z}^l_{i,j} = \frac{\sum_{m,n} \mathbf{Z}^{l-1}_{i+m,j+n} \mathbf{C}^{l-1}_{i+m,j+n} \Gamma(\mathbf{W}^l_{m,n})}{\sum_{m,n} \mathbf{C}^{l-1}_{i+m,j+n} \Gamma(\mathbf{W}^l_{m,n}) + \epsilon} + \mathbf{b}^l \ ,$$

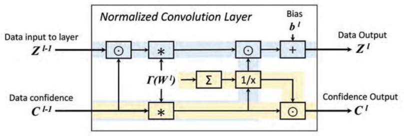

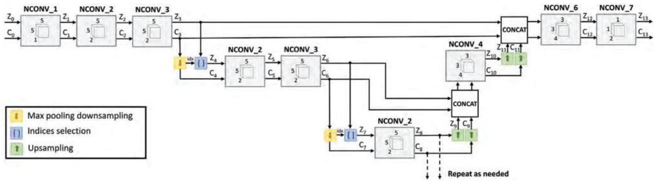

How to Fuse with RGB images?

- For autonomous driving setup, RGB images are also available.
- We investigate different fusion schemes in [3] and we design a very efficient fusion network with only 300k parameters.

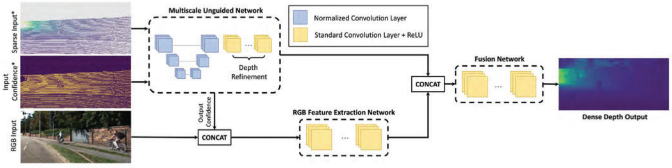
Sparse Input*
Input Confidence*
RGB Input
Multiscale Unguided Network
Depth Refinement
Normalized Convolution Layer
Standard Convolution Layer + ReLU
Output Confidence
CONCAT
RGB Feature Extraction Network
Fusion Network
Dense Depth Output

How to estimate input confidence?

- The input confidence is usually unknown for depth data and was assumed to be binary in the previous papers.
- We employ prediction error gradients to estimate the input confidence that minimizes the prediction error.

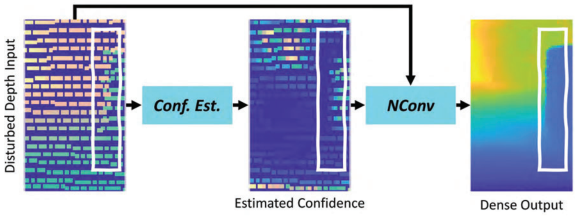

Does this approach generalize to other types of sparse data?

- To test the generalization capabilities of our approach to other types of sparse data, we evaluated it on sparse optical flow rectification problem.
- Our approach successfully rectifies the noisy flow and produce a denser flow that is very similar to the groundtruth.

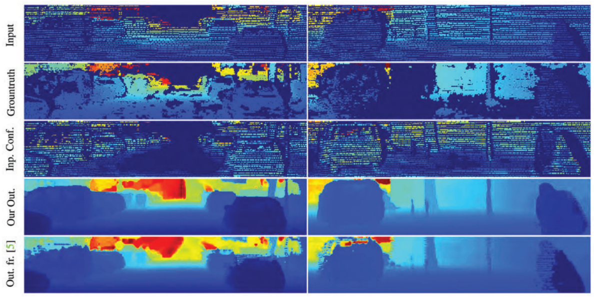

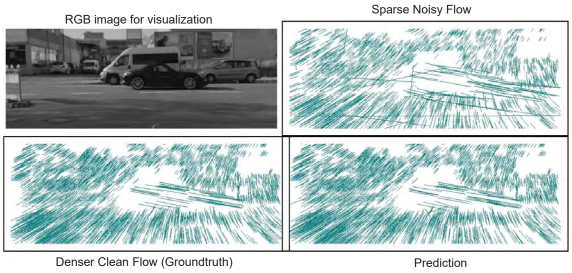

References

1. Knutsson, Hans, and C-F. Westin. "Normalized and differential convolution." *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 1993.
2. Eldesokey, Abdelrahman. "Propagating Confidences through CNNs for Sparse Data Regression." *The British Machine Vision Conference (BMVC)*. 2018.
3. Eldesokey, Abdelrahman, Michael Felsberg, and Fahad Shahbaz Khan. "Confidence propagation through CNNs for guided sparse depth regression." *IEEE transactions on pattern analysis and machine intelligence* (2019).

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Software-based Side-Channel Attacks on Crypto

Alexander Nilsson, Lund University / Advenica AB
Electrical and Information Technology Department (EIT)

Attacks and Defenses against Micro Architectural Side Channels

This project is covering both the attack and the design aspects of side-channel attacks in software in crypto, which is critical if they are to be employed to large scale **autonomous systems**. Although quite a lot is known on different attack strategies, there is still a massive interest in investigating more sophisticated software-based methods as well as improved theoretical approaches to process measured side-channel information.
So far we have investigated how timing information can be used to attack cryptographic systems based on error correcting codes. Next step is taking this information and using it to create new decoders which are not susceptible to timing attacks.

Published research:

Error Amplification in Code-Based Cryptography

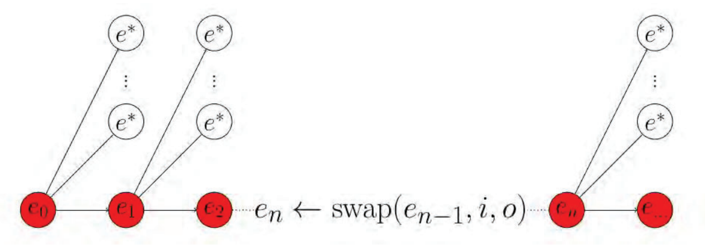

In [1] work we developed a “chaining method” in order to amplify the rate of decryption failures and the number of difficult to decode patterns.
We used these failures and difficult patterns (**timing** and **reaction** information) to statistically infer information about the so-called distance spectrum (se below) of the secret key.
This information can then be used to make a total reconstruction of the secret key.

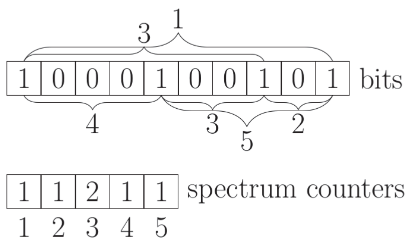

Publications

1. Nilsson, A., Johansson, T., Stankovski, P.: Error amplification in code-based cryp-tography. IACR Trans. Cryptogr. Hardw. Embed. Syst. 2019(1), 238–258 (2019),https://doi.org/10.13154/tches.v2019.i1.238-258
2. Guo, Q., Johansson, T., Nilsson, A.: A generic attack on lattice-based schemesusing decryption errors with application to ss-ntru-pke. Cryptology ePrint Archive,Report 2019/043 (2019),https://eprint.iacr.org/2019/043
3. D’Anvers, Jan-Pieter, et al. "Decryption Failure Attacks on IND-CCA Secure Lattice-Based Schemes." *IACR International Workshop on Public Key Cryptography*. Springer, Cham, 2019.

Ongoing research:

New Constant-Time Decoder

One strategy to prevent attacks such as the one presented to the left is to lower the risk of decryption failure. To do this the algorithm designers need to tweak the parameters. This can have detrimental effects on sizes of the ciphertexts and keys (higher memory usage) as well as on the performance of key generation, encryption and decryption.
However if the designers have access to good decoder implementations that have by themselves lower failure rates it is possible the tweaks can be made smaller which would increase the overall efficiency of the encryption scheme.
Our ongoing work is targeted towards this exact goal and we are looking towards utilizing the well established body of work regarding soft-input decoders in order to establish if they can be used to develop hybrid decoders which are both reasonably fast and with low decoding failure rate. We are also looking toward other techniques for accomplishing the same goal.

Future research:

Side Channels Attacks in Next Generation Asymmetric Crypto

Another consideration to take into account when trying to prevent attacks such as the one published by us in [1] (and other attacks) is that it is imperative that all cryptographic primitives are implemented in **constant time**. Constant time means that no matter what the input to the function is no information can be gained through observation of the amount of time it takes for the algorithm to complete.
Many techniques exists for doing this kind of implementations but they can sometimes be hard to get right. The goal of each technique is the removal of secret dependent conditionals. The removal of secret dependent memory addressing is a technique used make implementations more secure against cache attacks and other micro architectural attacks.
Unfortunately these techniques are not practical or trivial to implement and further research is needed to make the implementations of the next generation of cryptographic algorithms more secure.
Our focus will be towards the asymmetric algorithms currently being evaluated for standardization.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

LUNDS UNIVERSITET
Lunds Tekniska Högskola

ComPOS: Composing Oblivious Services

Alfred Åkesson, LTH

Description

Future internet-of-things systems need to be able to combine heterogeneous services and support weak connectivity. In this poster, we show ComPOS, a domain-specific language for composing services in IoT systems. We show how Maria, a bird watcher, can use ComPOS to build a system that allows her to spy on birds in the garden while she is not at home. We demonstrate how ComPOS handles the unpredictable nature of IoT system by analysing in what cases Maria's system is still useful when some devices are unavailable.

System overview

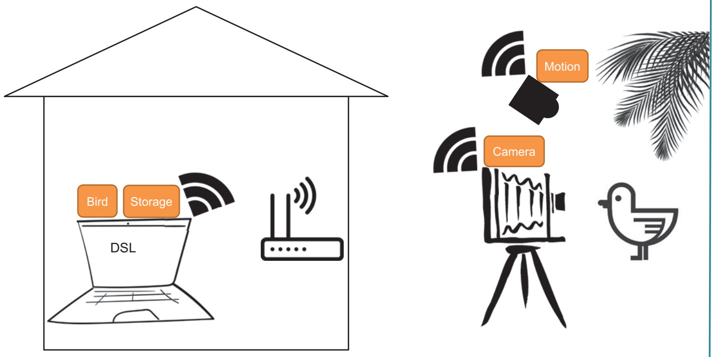

Palcom

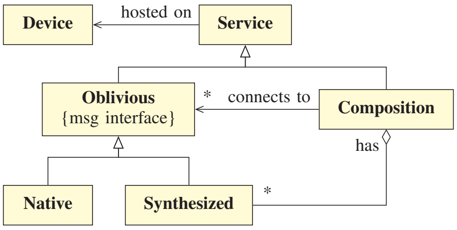

Palcom is a research middleware for creating pervasive systems.
In Palcom, a **device** has a unique id and the device can run services.
A **oblivious service** exposes an interface in the form of input- and output commands.
An **composition** is a script that connects oblivious services to create new systems.
A oblivious service does not know who it talks to – this knowledge is in the composition.

Extended system

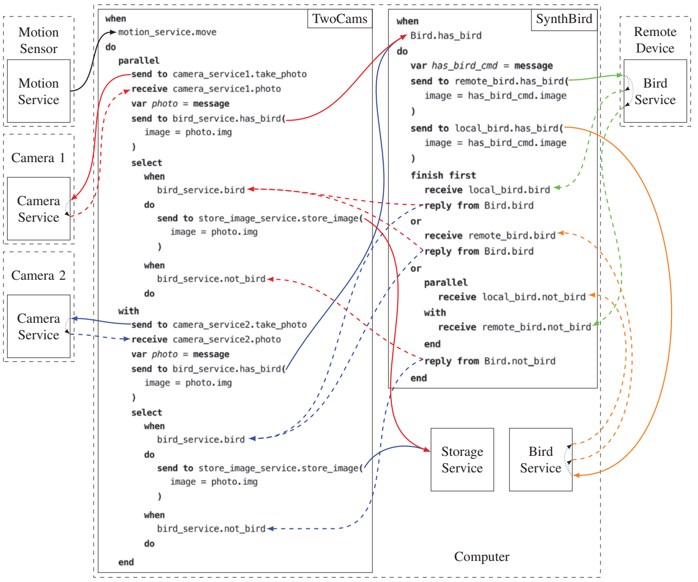

```
when
  motion_service.move
do
  parallel
    send to camera_service1.take_photo
    receive camera_service1.photo
    var photo = message
    send to bird_service.has_bird(
      image = photo.img
    )
    select
      when
        bird_service.bird
      do
        send to store_image_service.store_image(
          image = photo.img
        )

      when
        bird_service.not_bird
      do
  with
    send to camera_service2.take_photo
    receive camera_service2.photo
    var photo = message
    send to bird_service.has_bird(
      image = photo.img
    )
    select
      when
        bird_service.bird
      do
        send to store_image_service.store_image(
          image = photo.img
        )

      when
        bird_service.not_bird
      do
  end
```

```
when
  Bird.has_bird
do
  var has_bird_cmd = message
  send to remote_bird.has_bird(
    image = has_bird_cmd.image
  )
  send to local_bird.has_bird(
    image = has_bird_cmd.image
  )
  finish first
    receive local_bird.bird
    reply from Bird.bird
  or
    receive remote_bird.bird
    reply from Bird.bird
  or
    parallel
      receive local_bird.not_bird
    with
      receive remote_bird.not_bird
    end
    reply from Bird.not_bird
  end
```

ComPOS

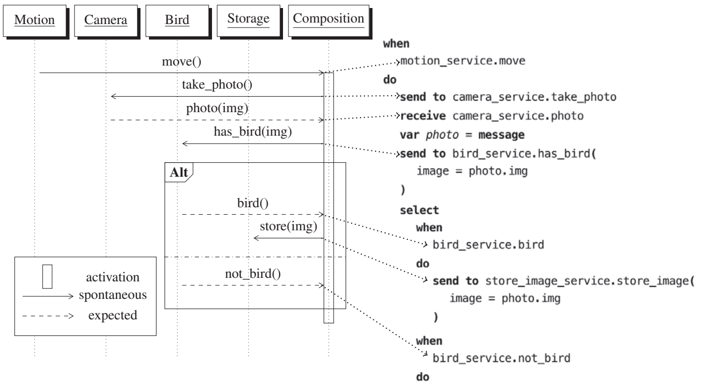

```
when
  motion_service.move
do
  send to camera_service.take_photo
  receive camera_service.photo
  var photo = message
  send to bird_service.has_bird(
    image = photo.img
  )
  select
    when
      bird_service.bird
    do
      send to store_image_service.store_image(
        image = photo.img
      )

    when
      bird_service.not_bird
    do
```

ComPOS features

- Asynchronous message passing
- Sequence of request and response
- Alternative depending on response
- Parallel and Finnish first
- Abstraction via synthesized services

ComPOS aborts the current reaction when a new message arrives and begins to execute the new reaction. ComPOS allows for multiple reactions to execute in parallel using reaction ids.

Utility analysis

- Analyses what happens to the usefulness of the system when devices come and go.
- Emulate what happens if a device is, for example, out of reach, out of battery, or has connection problems.
- The system is useless if no image can get stored.

Disconnected devices	Status	Reason
Computer	Useless	The system has nowhere to store images.
Motion sensor	Useless	No *move* messages arrive to start a reaction.
one Camera	Useful	The other camera can still take a photo and store it because the branch associated with that camera works as intended. For every new *move* message, the TwoCams composition creates a new activation and aborts the old one.
Remote Device	Useful	The synthesized service will never be able to send *not_bird*, but in the case the local bird service detects a bird the synthesized service will reply with *bird*.
both Camera	Useless	No camera to take the photo to be stored.
one Camera and Remote Device	Useful	If the local bird service detects a bird in a photo from the connected camera, that photo will be stored.
all other combinations	Useless	

A.Åkesson,G.Hedin,B.Magnusson,andM.Nordahl,"ComPOS:com- posing oblivious services," in *PerIoT'19 - Third International Workshop on Mobile and Pervasive Internet of Things (PerIoT'19)*, Kyoto, Japan, Mar. 2019.

WASP WALLENBERG AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Efficient Learning of Dynamical Systems

Christian Rosdahl, Bo Bernhardsson
Dept. of Automatic Control, Lund University
christian.rosdahl@control.lth.se, bo.bernhardsson@control.lth.se

Introduction

The project aims to develop methods that **learn complex dynamical models** through algorithms that actively explore the behavior of the system. The key challenge is to **trade the cost of "exploration"** (in terms of deteriorated momentary control) **with the future benefits from the resulting improved knowledge**. This area, sometimes named "dual control theory", is now vitalized from recent progress in machine learning and statistical estimation. The project will both develop new theory and algorithms, as well as work with applications.

Dual Control

Recent progress in the area of deep neural networks and machine learning have opened up for new approaches to the problem of controlling a system whose characteristics are unknown or partially unknown. The area has been studied under the name **dual control theory** since the 1960s. The name comes from the fact that the controller's objective is both **exploration**, to experiment with the system to learn its dynamics, and **exploitation**, to control the system as well as possible based on the achieved system knowledge. The problem is also studied in the area of **Bayesian reinforcement learning**, see [7] for a discussion.

The learning performance of the dual controller is judged by both its asymptotic performance: will optimal, or acceptable, control performance eventually be found? and also by the speed at which this is achieved. The latter is measured in the amount of "regret" of the controller, i.e. the amount of degraded performance before the optimal controller was found.

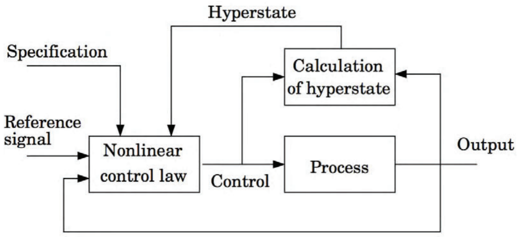

Figure 1: In dual control, the control signal should be chosen such that we learn sufficiently much about the partially unknown process (its hyperstate) in order to achieve optimal control in the long-term perspective, even if this would imply degraded momentary performance.

Problem

A solution to the dual control problem is in theory known. It is based on **dynamic programming** (DP) propagating the so-called **hyperstate** in a Hamilton-Jacobi-Bellman-Kolmogorov equation. The hyperstate consists of the probability distribution of the system's state and parameters. The hyperstate is unfortunately **infinite-dimensional**, except for special cases, and therefore this approach has been deemed unfeasible, except for smaller toy-examples, such as in [2].

Methods

The key issue in implementing the DP solution is that one needs to **find a structure where the cost of a certain hyperstate can be evaluated efficiently**. To choose control signal in the optimization step one needs to efficiently evaluate different candidate control signals' impact on the hyperstate and the resulting performance. Here we plan to use recent progress in **deep neural networks**, [5], and **machine learning** to learn approximations of this structure. Alternative approaches based on **Monte Carlo methods**, such as in [4] are also worth investigating and possibly combine with.

Information Geometry

To understand and learn the geometry of the manifold of the probability distributions of the system, new progress in **information geometry** seems to be useful, see [1]. The learning process is then viewed as an optimization problem defined over a statistical manifold, i.e., a set of probability distributions. Information geometry is a blend of statistics, differential geometry and information theory, where the Fisher information metric provides the Riemannian metric. For a discussion of its use in manifold learning, see [9]. We will study systems with nonlinear dynamics where the full state of the system is not measurable and where the dynamics is unknown or partially unknown. There are several applications of the theory, such as **intelligent robotics**, **autonomous vehicles**, a**daptive radio systems** and **intelligent process control**. One possible concrete new application is in **adaptive pilot assignment**, an idea with possible use for future radio communication systems, such as in [8]. Here, a controller chooses between using radio resource blocks for communication or for gaining knowledge about the evolution of the radio channel.

Feedback Particle Filters

The state of an observable linear system with Gaussian distributed disturbances can be optimally estimated with a Kalman filter. For more complicated systems, **particle filters** can sometimes be used for this purpose. These operate by sampling the state space and weighting the sample points, called particles, according to their probability given a process model as well as inputs and outputs of the process. The weights constitute an estimation of the probability distribution of the system state.

The number of particles needed to accurately approximate the probability distribution of the system state increases rapidly with the order of the system and makes the particle filter intractable for higher-order systems. It has been shown, see [10], that a new type of particle filter – the **feedback particle filter** – does not suffer from this problem. Therefore, we aim to investigate the possibilities of using and further developing this type of filter.

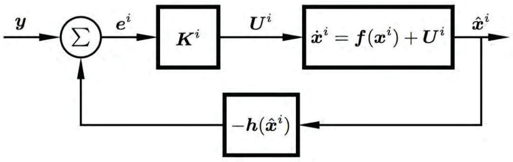
$\boldsymbol{y}$
Σ
$\boldsymbol{e}^i$
$\boldsymbol{K}^i$
$\boldsymbol{U}^i$
$\dot{\boldsymbol{x}}^i = \boldsymbol{f}(\boldsymbol{x}^i) + \boldsymbol{U}^i$
$\hat{\boldsymbol{x}}^i$
$-\boldsymbol{h}(\hat{\boldsymbol{x}}^i)$

Figure 2: In a feedback particle filter, the particles (sample points in state space) $\boldsymbol{x}^i$ are moved according to a feedback law, analogously to the Kalman filter. *Figure source:* [6]

References

[1] *Amari, S.*, "Information Geometry and Its Applications", Springer, 2016.
[2] *Bernhardsson, B.*, "Dual Control of a First-Order System with Two Possible Gains", International Journal of Adaptive Control and Signal Processing, 1989.
[3] *Berntorp, K.*, "Feedback Particle Filter: Application and Evaluation", 2015 18th International Conference on Information Fusion (Fusion), 2015.
[4] *Bayard DS, Schumitzky A.*, "Implicit dual control based on particle filtering and forward dynamic programming", International journal of adaptive control and signal processing. 2010;24(3):155-177.
[5] *Goodfellow, Bengio and Courville*, "Deep Learning", MIT Press, 2016.
[6] *Koller T., Berkenkamp F., Turchetta M., Krause A.*, "Learning-based Model Predictive Control for Safe Exploration", arXiv:1803.08287, 2018.
[7] *Klenske, Edgar D. and Hennig, Philipp*, "Dual Control for Approximate Bayesian Reinforcement Learning", J. Mach. Learn. Res., pp 4354-4384, 2016.
[8] *Marzetta, Larsson, Yang and Ngo*, Fundamentals of Massive MIMO, 2016.
[9] *Sun, Ke.*, "Information geometry and data manifold representations", Université de Genève, PhD Thesis, 2015.
[10] *Surace S. C., Kutschireiter A. and Pfister J.-P*, "How to avoid the curse of dimensionality: scalability of particle filters with and without importance weights", arXiv:1703.07879, 2017.

LUND UNIVERSITY

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Explainable robustness for self-driving vehicles

Dhasarathy Parthasarathy, Volvo Group

Understanding context in the traffic scene aids inference

Identifying vulnerable road users with a high degree of confidence is a major challenge for self-driving vehicles. Deep perception nets may have made remarkable progress in identifying critical elements in the vehicle's surroundings but their capabilities remain largely unknown, raising questions on overall system safety. This work focuses on examining properties of perception nets by testing its ability to identify entities in the scene that has definable geometry.

Can we make a self-driving vehicle safer by ensuring with high confidence that it identifies critical traffic signs?

Well-structured information in the traffic scene

Road transport authorities across the globe ensure that a driver is sufficiently informed about vulnerable road users. The most common way of doing this is to put up traffic signs. The structure of and content each sign is well-specified.

High likelihood of hazards – stay alert!

Given that a critical traffic sign is present, one must take any sign of a hazard seriously. Can we not apply the same principle to a self-driving vehicle?

If the system confidently identifies critical traffic signs, it can give more importance to related predictions of lower confidence.

Explaining deep net abilities in identifying structure

As proxies for traffic signs, two simple geometries were chosen – circles and squares. A simple network was trained to identify these shapes using the Quick-draw dataset [1]. This is the analogy for traffic signs in the 'wild'.

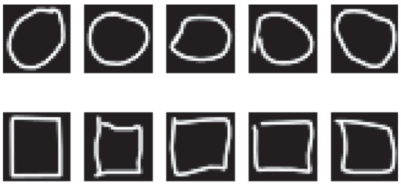

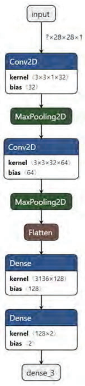

(Left) - Examples of real-world circles and squares

(Right) – 2-layer classifier trained on this dataset

We then test the net with a new ***specified*** dataset, which contains additional labels that describe the structure.

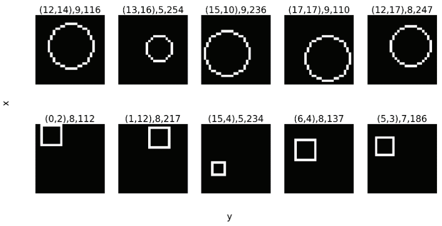

(Above) – Dataset with labeled structure properties
(Below) – Shapley values [2] indicating influence properties

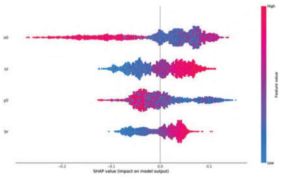

The result is an explanation of the deep nets capability in identifying structure.

This particular network is weak in identifying small squares in the bottom-right of the canvas. Either re-train it or mitigate the risk in some other way

References

1. https://github.com/googlecreativelab/quickdraw-dataset
2. Lundberg, Scott & Lee, Su-In. (2017). A Unified Approach to Interpreting Model Predictions.

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Photon Mapping Superluminal Particles

Gustaf Waldemarson
gustaf.waldemarson@arm.com

Michael Doggett
michael.doggett@cs.lth.se

Description

In order to understand how perception works for autonomous systems, a good understanding of the inverse problem is necessary: How does light propagate in a scene? This is the primary question posed in the field of light transport rendering and is commonly solved with ray tracing. In this work a new phenomenon is added to this field: Cherenkov radiation [2]. Light that comes from particles traveling faster than the speed of light for the current medium.

Cherenkov Radiation

As a particle moves through a medium it will excite atoms and cause them to emit electromagnetic waves spherically from the point of interaction.

When the particle exceeds the medium phase speed they constructively interfere, generating coherent photons known as Cherenkov radiation.

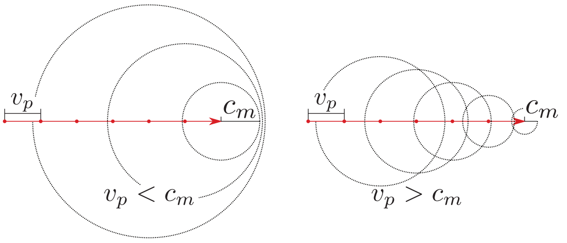

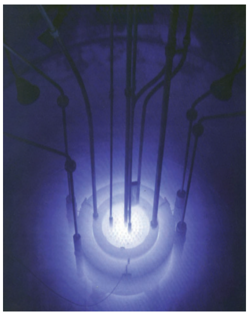

Frank-Tamm Equation

This equation describes how the photon energy gets distributed by a charged particles that travels faster than the speed of light in the medium.

$$\frac{d^2N}{dxd\lambda} = -\frac{2\pi\alpha\mu(\lambda)}{\lambda^2}\left(1 - \frac{c_0^2}{v^2n^2(\lambda)}\right)$$

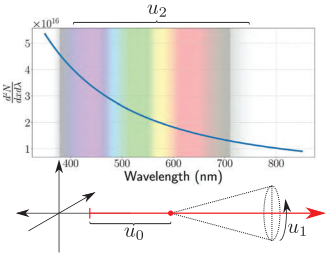

Photon Mapping

Many different ray tracing algorithms exist but one of the most flexible ones is the Stochastic Progressive Photon Mapping algorithm [1], where paths are traced from both light sources and the camera. Additionally, statistical tricks such as Russian roulette can be utilized to improve rendering performance.

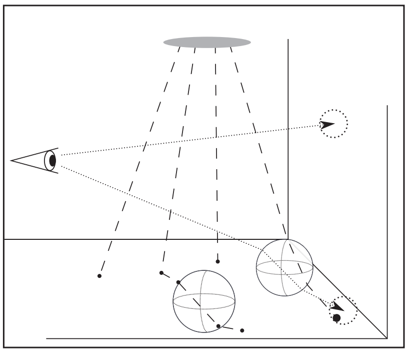

Photon Distributions

Statistical Russian roulette requires prior knowledge of various density distributions, typically denoted as $p(\mathrm{o}, \omega)$ for the photon origin and direction respectively. In this work, it is estimated as follows:

$$\begin{aligned} p(\mathrm{o}, \omega) &= p(\mathrm{o}) \cdot p(\omega) \\ p(\mathrm{o}) &= \frac{1}{\text{total particle length}} \\ p(S) &= \frac{\text{superluminal path lengths}}{\text{total path length}} \\ p(\omega) &= p(S)p(\omega_\mathrm{c}) + (1 - p(S))p(\omega_\mathrm{s}) \\ &= \frac{p(S)}{2\pi} + \frac{1-p(S)}{4\pi} = \frac{1+p(S)}{4\pi} \end{aligned}$$

Where the remaining quantities are:

o, ω Photon origin and direction.

$p(S)$ Probability of a particle being superluminal.

$p(\omega_\mathbf{c})$ density for a emitting in a known cone $(\frac{1}{2\pi})$.

$p(\omega_\mathbf{s})$ density for emitting in a sphere $(\frac{1}{4\pi})$.

Our Algorithm

1. Choose a point along the particle path
2. Find the refractive index at the location
3. If *superluminal* at the point

   Find the Cherenkov emission cone and choose a direction along that its surface to emit a photon towards
4. Otherwise

   Emit the photon in a random direction
5. Use the Frank-Tamm spectra for the particle as photon color
6. Trace the photon as in [1]

Rendering Results

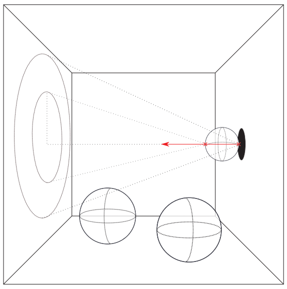

(a) Single particle and Cherenkov cone.

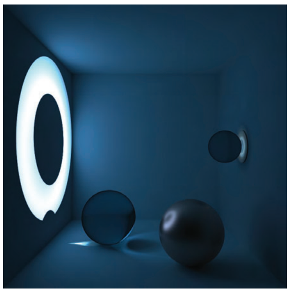

(b) Rendering of (a).

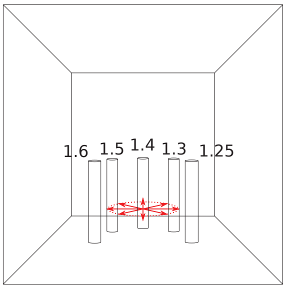

(c) Many particles and varying refractive index.

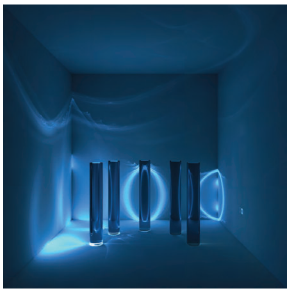

(d) Rendering of (c).

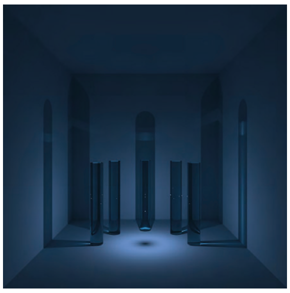

(e) Rendering (c) with a conventional area light.

References

[1] T. Hachisuka, S. Ogaki, and H. W. Jensen. Progressive photon mapping. *ACM Trans. Graph.*, 27(5):130:1–130:8, Dec. 2008.

[2] P. A. Čerenkov. Visible radiation produced by electrons moving in a medium with velocities exceeding that of light. *Phys. Rev.*, 52:378–379, Aug 1937.

Massive MIMO Pilots scheduling on a MAC-PHY split Architecture

Haorui Peng, Lund University
Department of Electrical and information Technology

Abstract

Cloud computing draws signification attention in the Telecommunication community in recent years as it enables to implement more flexible, cost-efficient and powerful mobile network [1][2]. Cloud Radio Access Network (Cloud-RAN) is a disruptive technology that introduces the centralized baseband processing via splits of the baseband units (BBU) from the radio remote head s(RRHs). Currently Cloud-RAN technologies considers only the provisioning of Digital Signal Processors (DSPs) at central entity, which is still costly for both hardware and software installation. This could benefit from push part of the network function into a cloud-native deployment. Another concerns about Cloud-RAN is that having centralized processing requires high network bandwidth between BBU and RRHs that only optical fibers can handle,. Thus it is still under discussion about the trade-off between the network bandwidth cost and functional scalability when having different split options in the signal processing chain. In this work, we consider the offload the MAC pilot scheduling function for Massive MIMO to the cloud and investigate the impact on the user experiences caused by the network dely between MAC layer and PHY.

System Architecture

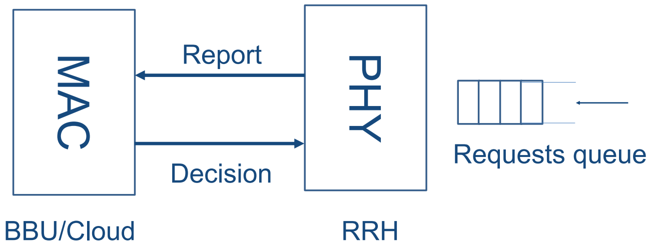

Report: {Current size of the requests queue}
Decision: {Number of requests to be scheduled each coherence interval}

Problem:
1. Reporting frequency != Coherence interval
→ The reported information highly depends on traffic variation
2. Delay between the entities
→ The report is dated when arrives at the MAC layer, the decision is also not up to date.

Simulation with FaaS

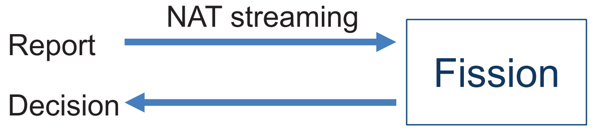

References

[1] 5GPP, "From Webscale to Telco , the Cloud Native Journey," 2018.
[2] D. Wübben *et al.*, "Benefits and impact of cloud computing on 5g signal processing: Flexible centralization through cloud-RAN," *IEEE Signal Process. Mag.*, vol. 31, no. 6, pp. 35–44, 2014.

Impact of the traffic variance on the system performance

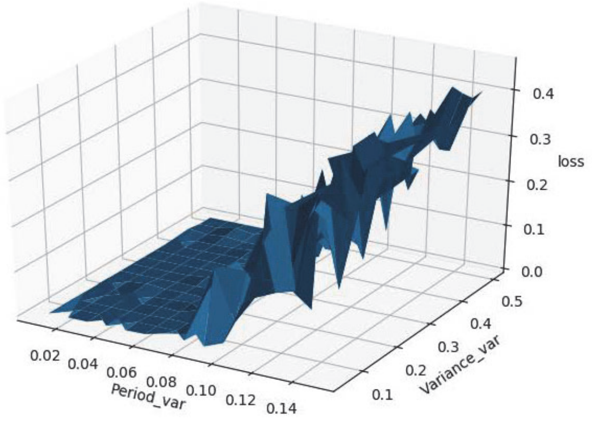

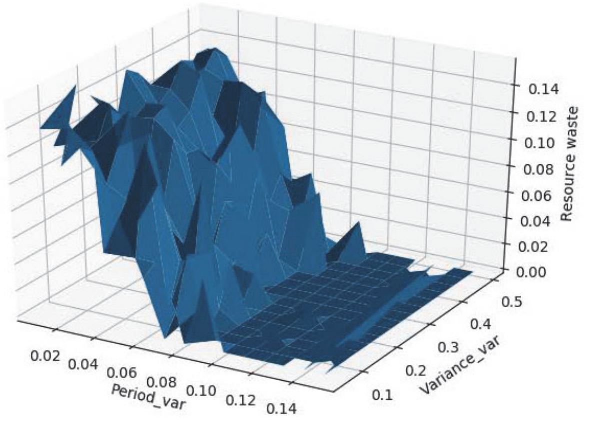

Trade-offs between packet loss and resource utilization
→ Over-scheduling : Resource waste
→ Under-scheduling: User-deadline missed

Solution under investigation:
Prediction on the user traffic and network delay in the fronthaul

Out-of-Distribution Robustness for DNN's

Jens Henriksson, Chalmers University of Technology
Department of Computer Science and Engineering
Main supervisors: Christian Berger, Stig Ursing

CHALMERS

semcon

Motivation & Research Goals

Deep Learning (DL) models have shown promising results for object segmentation and classification tasks past few years, thus making the method desired for inclusion in various products. However, before including DL models in safety critical applications, they need to show robustness by rigorous testing, which is still unclear how to properly conduct. The aim if this thesis is to investigate the effect of **out-of-distribution samples and how these can be identified for deep neural networks.** Detecting these outlier events allows the system design to incorporate fallbacks which reduces the risk of harmful events to occur, as well as can be used to argue for safety and that the model is operating within it's functional limits.

Methods

Comparing Out-of-distribution (OOD) methods is done by varying a threshold ϵ, to see how the true positive (TPR) and false positive rate (FPR) varies.

```
Require: threshold $\epsilon$
Require: Out-of-distribution method $F(\cdot)$
  Compute the output vector $\mathbf{v}(\mathbf{x})$ for input sample $\mathbf{x}$ with
  the Out-of-distribution method $F(\cdot)$
  Let $P$ be the Softmax of output vector $\mathbf{v}$ for all classes
  $j = 1, .., N$
      $P(y = j|\mathbf{x}) = \frac{e^{\mathbf{v}_j(\mathbf{x})}}{\sum_{i=1}^{N} e^{\mathbf{v}_i(\mathbf{x})}}$    (1)
  Let the anomaly score be $A(\mathbf{x})$ = 1 - argmax $P(y = j|\mathbf{x})$
  where $argmax$ is the function argument that maximizes
  function value.
  Let discriminator $D\{0, 1\}$ be defined as
      $D = \begin{cases} 1 & if\, A(\mathbf{x}) < \epsilon \\ 0 & otherwise \end{cases}$    (2)
```

We compare the TPR and FPR in the classical ROC-curve. The curve shows the probability of correct detection in contrast to the probability of a false alarm.

Additionally, we compare it to Risk/Coverage. The R/C-plot indicates how restrictive the OOD-method has to be to reduce the risk of false alarms. Typically, a requirement could be "a maximum if X false alarm per Y samples".

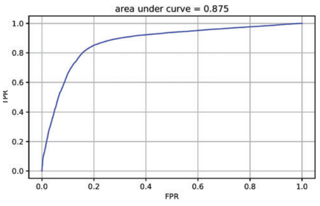

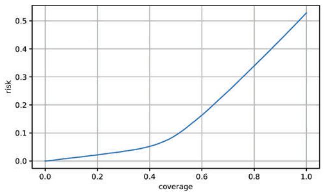

References

1. Henriksson et al. *Performance Analysis of Out-of-Distribution Detection on Various Trained Neural Networks,* SEAA 2019
2. Henriksson et al. *Towards Structured Evaluation of Deep Neural Network Supervisors*, IEEE AI-Testing 2019
3. Henriksson et al. *Automotive safety and machine learning: Initial results from a study on how to adapt the ISO 26262 safety standard,* ICSE-SEFAIAS 2018 (Workshop)

Experimental Results

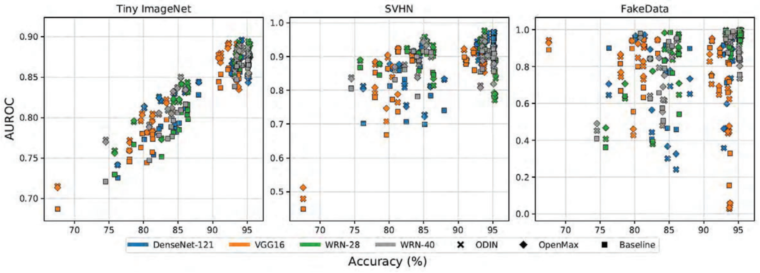

As the performance of the model increases, it simplifies the task of detecting out-of-distribution samples. However, small variations in the model, which doesn't change the predictive performance, can vary the out-of-distrubtion detection by a large amount.

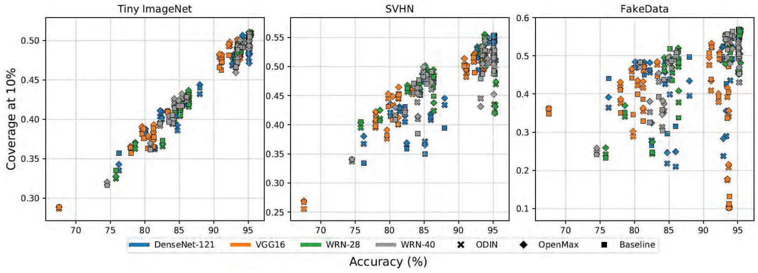

When operating towards a given accuracy (i.e a requirement of maximum 10% error rate), the coverage also increases. This indicates that well performing models will exclude out-of-distribution samples better, as well as cover more of the input domain.

Future directions: To validate deep neural networks before deploying them in safety-critical applications, it is important to have structured tests of robustness. Future work needs to investigate how to properly do testing of DNN's, and what testing methods are needed. Out-of-distribution detection is one testing method. This method can be further improved.

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Recurrent Neural Networks for Perception

Joakim Johnander[1,2], supervised by *Michael Felsberg*[2], *Martin Danelljan*[3], *Nicholas Wickström*[1]

[1] Zenuity

[2] Linköping University

[3] ETH, Zürich

Idea

- Humans maintain a state of the perceived scene, we know:
  - what we have seen
  - what we expect to see
- Additionally, we easily detect moving objects
- Should computers not work in a similar fashion?

Aim

This project investigates the use of Recurrent Neural Networks, a Deep Learning strategy designed for temporal and sequential data

- What information need we propagate?
- How do we represent this information?
- How do we actually learn to propagate anything useful (and avoid overfitting)?
- How do we utilize sparse error signals?
- How do we deal with the high correlation of the data?

Approach

Thus far, focus has been on image-plane tracking. The problem is suitable for the study of recurrent neural networks as it cannot be solved without propagating information temporally. Aim is to proceed to additional problems where spatio-temporal models may help (s.a. object detection, semantic segmentation, or depth estimation).

Work on **Video Object Segmentation**

Oral at CVPR2019

Joakim Johnander, Martin Danelljan, Emil Brissman, Fahad Shahbaz Khan, Michael Felsberg

How can we train a neural network to learn the appearance of a generic object, in order to track and segment it in a video stream?

Our Approach:

- **Idea:** Probabilistically model the feature distribution of the target and background
- Feature vectors are explicitly classified via the probabilistic model
- Formulated as End-to-end Differentiable Neural Network Modules
- Inference through a simple forward-pass

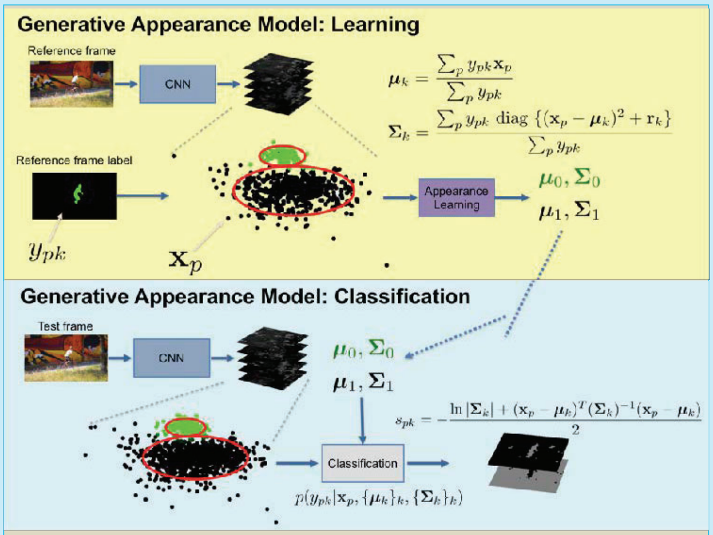

Generative Appearance Model: Learning

$$\mu_k = \frac{\sum_p y_{pk} \mathbf{x}_p}{\sum_p y_{pk}}$$

$$\Sigma_k = \frac{\sum_p y_{pk} \operatorname{diag}\{(\mathbf{x}_p - \mu_k)^2 + \mathbf{r}_k\}}{\sum_p y_{pk}}$$

Generative Appearance Model: Classification

$$s_{pk} = -\frac{\ln|\Sigma_k| + (\mathbf{x}_p - \mu_k)^T (\Sigma_k)^{-1} (\mathbf{x}_p - \mu_k)}{2}$$

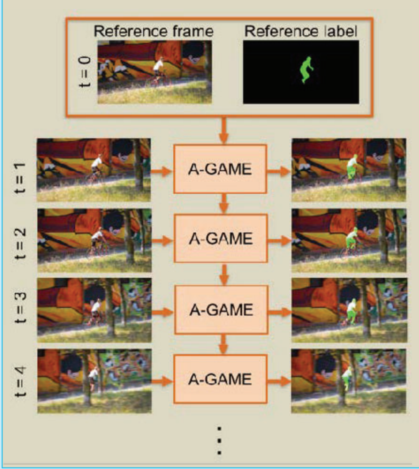

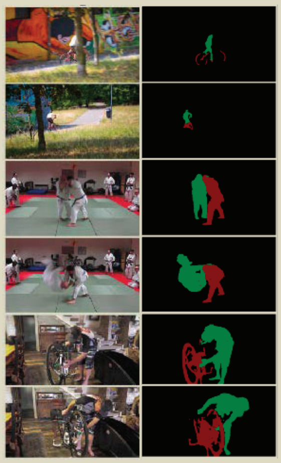

Work on **Semi-automatic Annotation of Objects in Visual-Thermal Video**

Workshop at ICCV2019

Amanda Berg, Joakim Johnander, Flavie Durand de Gevigney, Michael Felsberg

Dense annotation of video is expensive. This work aims to propagate a small set of annotations throughout video, to obtain a densely annotated video.

- Here we work with segmentations
- We detect automatic annotation failure via
  - Temporal consistency, if we track an object forward from the start, and backward from the end, do the tracks agree?
  - Modal consistency, does the tracking in different modalities agree? (here Visual and Thermal)

Keeping central authorities out of your business

Joakim Brorsson[1,2,3], Paul Stankovski[1], and Martin Hell[1]

[1]Department of Electrical and Information Technology, Lund University, Sweden
[2]Wallenberg AI, Autonomous Systems and Software Program (WASP)
[3]Combitech AB, Sweden

Motivation

As computing becomes more ubiqutous and important for vital society functions, it becomes more important to develop technical methods for guaranteeing the correct behaviour of these systems. That is, systems with access to sensitive data should be under audit for correct handling of this data, and systems should not be able to access more data that they have the right to. The research in this project investigates technical methods for restricting access to sensitive data and guaranteeing correct use of data when access is granted.

Research Directions

Transparent Authorities

This line of research is concerned with ensuring correct behaviour from authorities such as Certificate Authorities.

The aim is to, by technical means, force transparency and thereby gain accountability for central parties with a high degree of power. This way we hope to create a detection mechanism for misbehaviour and thereby both an incentive to avoid such behaviour and a basis for discussing what powers are reasonable to give such authorities.

Published work is a paper called ”Guarding the Guards: Accountable Authorities in VANETs” [1] addressing the paradoxical situation for Vehicular Ad-Hoc Networks (VANETs) where there are requirements for both privacy from authorities and accountability towards authorities for users.

Current work in this area is an efficient and realistic version of the published protocol aimed at general use instead of just for VANETs.

Secure Enclaves for a Trustable Cloud

While cloud computing is a positive development for scalability, price and performance, it does have consequences for security and privacy. Data residing in a cloud service is necessarily accessible by the cloud provider and potentially retrievable from co-hosted services by malicious attacks. To address this, research has been done on whether Intel's Software Guard Extensions (SGX) technique can guard the data from these new threats without restricting the benefits of the cloud, such as vitualizability and migratability.

End-to-end Security for IoT

A significant part of IoT devices have limited computing and power resources. Since traditional communication security protocols such as TLS are not optimized for these performance limitations, they are not feasible to use in IoT scenarios. For this reason, research [2] has been done evaluating the performance of OSCORE, a new IETF protocol for end-to-end security aimed at IoT scenarios.

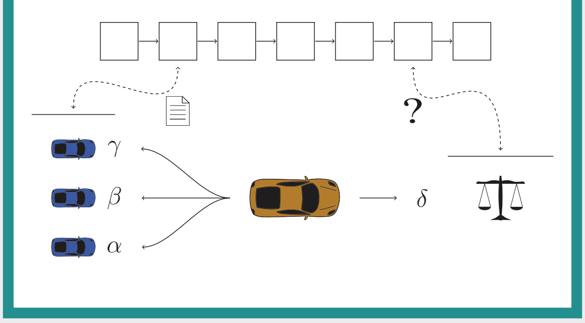

Future Work

Future work will focus on applying the research results on anonymous credential systems to get accountable authorities. These systems, in turn, will be applied to authentication mechanisms in Vehicular Networks and Payment Systems with requirements for conditional pseudonymity.

Published Work

[1] Brorsson, Joakim, Paul Stankovski Wagner, and Martin Hell. ”Guarding the Guards: Accountable Authorities in VANETs.” 2018 IEEE Vehicular Networking Conference (VNC). IEEE, 2018.

[2] Gunnarsson, Martin, Brorsson, Joakim, Palombini, Francesca, Seitz, Ludwig and Tiloca, Marco. ”Object Security for the Internet of Things.” Ad Hoc Networks XX (2020).

Performance-Optimized Cloud Applications

Joel Scheuner, Chalmers | University of Gothenburg

Computer Science and Engineering | Software Engineering Division

CHALMERS UNIVERSITY OF TECHNOLOGY | UNIVERSITY OF GOTHENBURG | ICET LAB

Empirical Performance Evaluation of IaaS Clouds

How can we systematically execute micro- and application-benchmarks in unstable [6] cloud environments?

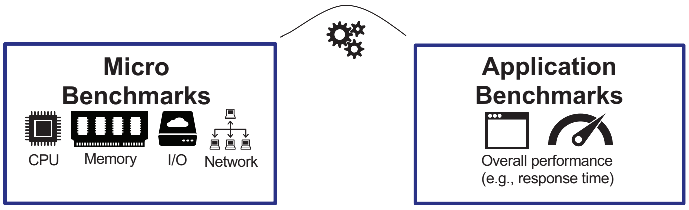

Generic	Specific
Artificial	Real-World
Resource-specific	Resource-heterogeneous

→ [1] A Cloud Benchmark Suite Combining Micro and Application Benchmarks

Performance Benchmarking
Supports finding optimal cloud services

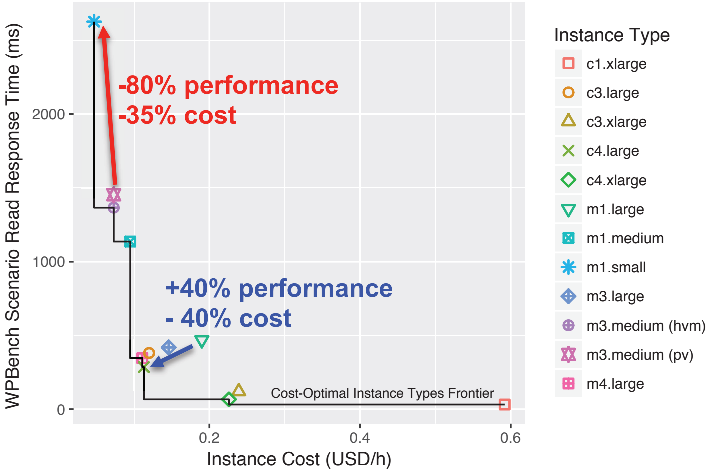

PRE – Performance Variability
Does the performance of equally configured cloud instances vary relevantly?

RQ1 – Estimation Accuracy
How accurate can a set of micro benchmarks estimate application performance?

RQ2 – Micro Benchmark Selection
Which subset of micro benchmarks estimates application performance most accurately?

→ [2] Estimating Cloud Application Performance based on Micro-Benchmark Profiling

Implications

- ✓ Suitability of *selected* micro benchmarks to estimate application performance
- Benchmarks cannot be used interchangeable → Configuration is important
- Baseline metrics vCPU and ECU are insufficient

Tutorials

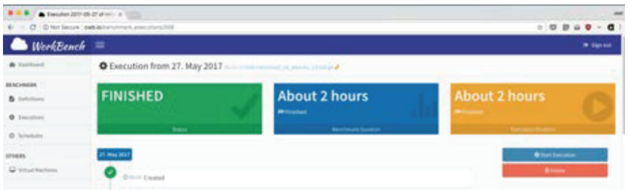
WorkBench

Execution from 27. May 2017

FINISHED

About 2 hours

About 2 hours

→ [3,4] Performance Benchmarking of Infrastructure-as-a-Service (IaaS) Clouds with Cloud WorkBench

Serverless Applications

RQ: How can we map (existing) single-machine code into applications composed of scalable cloud functions?

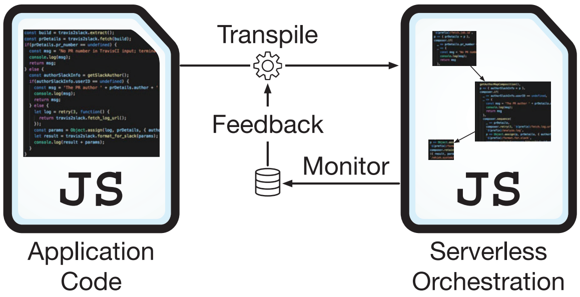

→ [5] Transpiling Applications into Optimized Serverless Orchestrations

References

[1] J. Scheuner, P. Leitner (2018). *A Cloud Benchmark Suite Combining Micro and Applications Benchmarks*. Companion of the 2018 ACM/SPEC Int. Conf. on Performance Engineering.

[2] J. Scheuner, P. Leitner (2018). *Estimating Cloud Application Performance Based on Micro-Benchmark Profiling*. 2018 IEEE 11th Int. Conf. on Cloud Computing (CLOUD).

[3] J. Scheuner, P. Leitner (2019). *Performance Benchmarking of Infrastructure-as-a-Service (IaaS) Clouds with Cloud WorkBench*. Companion of the 2019 ACM/SPEC Int. Conf. on Performance Engineering.

[4] J. Scheuner, P. Leitner (2019). *Tutorial – Performance Benchmarking of Infrastructure-as-a-Service (IaaS) Clouds with Cloud WorkBench*. 2019 IEEE 4th Int. Workshops on Foundations and Applications of Self* Systems (FAS*W).

[5] J. Scheuner, P. Leitner (2019). *Transpiling Applications into Optimized Serverless Orchestrations*. 2019 IEEE 4th Int. Workshops on Foundations and Applications of Self* Systems (FAS*W).

[6] C. Laaber, J. Scheuner, P. Leitner (2019). *Software microbenchmarking in the cloud. How bad is it really?*. Empirical Software Engineering.

scheuner@chalmers.se
joelscheuner.com
joe4dev
@joe4dev

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Statistical Inference for the Self-adaptive Cloud

Johan Ruuskanen & Anton Cervin

E-mail: {johan.ruuskanen, anton.cervin}@control.lth.se
Department of Automatic Control, Lund University, Sweden.

Introduction

Successful self-adaptive resource provisioning in the cloud relies on accurate tracking and prediction of workload variations and timely detection of changes in the infrastructure. However, the estimation problems in this context becomes challenging due to the massive number of measurements, heterogeneous behavior of tasks and time changing resource requirements. On this poster we present ongoing work in tackling these issues.

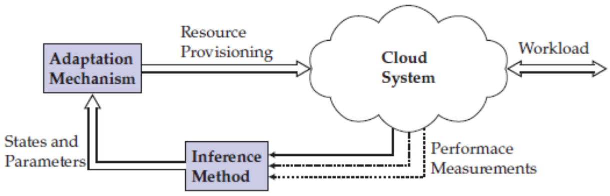

Event-based Particle Filtering

Server systems are commonly modeled using queuing theory. These models are often neither linear, continuous nor Gaussian, which makes the particle filter good candidate estimator.

Tracking States in Queue Models[1]

As a demonstrative problem, consider the simple yet relevant M/M/1 queue with arrival/service rates of λ and μ.

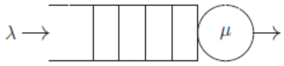

The M/M/1 queue.

From measurements of request response times, internal states such as the queue length and service rates can be estimated using the following state space model;

$$
\begin{aligned}
x_k &= \begin{cases} \mu_k = \mu_{k-1} + \Delta t_k w_k \\ q_k = \mathcal{P}(q_{k-1}, \mu_{k-1}, \lambda_{k-1}, \Delta t_k) \end{cases} & w_k &\sim N\left(0, \sigma^2\right), \\
y_k &= v_k & v_k &\sim \Gamma\left(q_k, \mu_k\right).
\end{aligned}
$$

The measurements are only available at request departures from the server, which happens in the form of events, making the system inherently event-based.

The Auxiliary filter for Event-based Sampling[2]

Event-based sampling gives rise to broad likelihood functions over the state space, which contracts when a new event is obtained. For a naive filter, the particles can become too spread out to give a good approximation of the posterior as illustrated below.

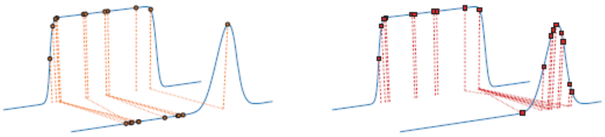

Comparison of the particle behavior between the bootstrap and auxiliary filter at a new event.

Instead, the *auxiliary particle filter* can be used to resample the particles conditioned on the likelihood of generating the observed event.

Predicting Resource Usage in Tasks using Gaussian Processes

Knowledge of future resource requirements of tasks running in a cloud can yield vital information to decision making. However, resource usage of tasks are wildly heterogeneous and thus prediction methods needs to be general and fairly computationally cheap. One such method is the *Gaussian Process*. Using data from google[3], we can demonstrate the method by fitting Gaussian processes to the average CPU rate of two task.

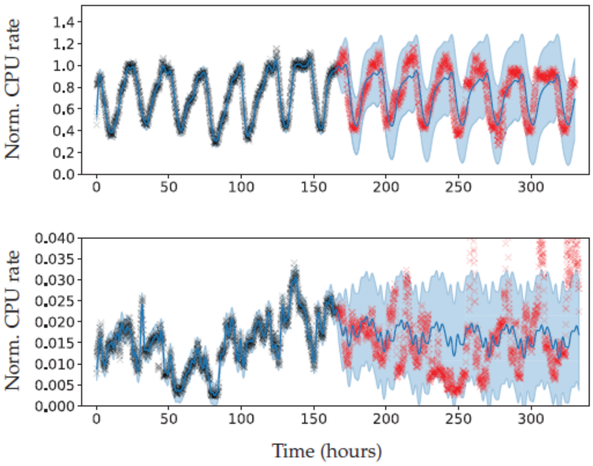

Prediction with 90 % confidence bounds from periodic data. Black dots represent seen data, red dots future data points.

Tailoring good kernel functions $k_\theta(x_1, x_2)$ to tasks that accurately captures the covariance is of importance, and how to detect changes which warrants retraining of the model.

Future Challenges

- Identifying problems where accurate estimations can provide a real impact.
- Identifying valid models of cloud systems from which states can be inferred.
- Access to data and test beds for evaluations on real systems.

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

[1] J. Ruuskanen and A. Cervin. *Internal Server State Estimation Using Event-based Particle Filtering.* 4th International Conference on Event-Based Control, Communication, and Signal Processing (EBCCSP). 2018.
[2] J. Ruuskanen and A. Cervin. *Event-Based State Estimation Using the Auxiliary Particle Filter.* European Control Conference (ECC), 2019
[3] https://github.com/google/cluster-data/blob/master/ClusterData2011_2.md

Exact remodeling of optimization programs with applications to autonomous driving

Johan Karlsson, Chalmers university

Electrical engineering

Motivation & Research goals

1. The goal of this research is to find efficient formulation of the optimal control program in MPC for trajectory generation in autonomous vehicles.
2. One problem is to implement collision avoidance due to the nonconvexity of this type of constraints.

Example: At what samples do the blue car reach the red rectangle?

This depends on the position and velocity of the blue car which is typically an optimization variable along with the position.

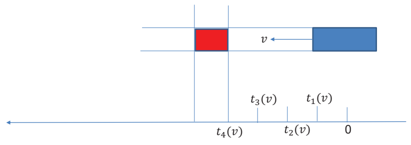

Licentiate

I. Remodeling longitudinal dynamics: Sample in longitudinal position instead of in time. Thus, if we have the vehicle model:

States/inputs: $\boldsymbol{x}(t) = [x(t)\ v_x(t)]^T$, $\boldsymbol{u}(t) = [F_x(t)]^T$

Linear dynamics: $\dot{\boldsymbol{x}}(t) = A\boldsymbol{x}(t) + B\boldsymbol{u}(t)$

Remodeling is done using the following steps:

- Change derivatives: x'=
- Replace longitudinal position with travel time in the state vector: $t' = \frac{1}{v_x} = z_x$

Remodeled states: $\boldsymbol{x}(s) = [\tilde{t}\,(s)\ z_x(s)]^T$**,** $\boldsymbol{u}(s) = -F_x(s) z_x^3(s)$**,**

Linear dynamics: $\boldsymbol{x}'(s) = A\boldsymbol{x}(s) + B\boldsymbol{u}(s)$

II. Remodeling with lateral dynamics: The remodeling can also be done using longitudinal and lateral dynamics:

Remodeling is done using the following steps:

1. Change of derivative derivatives.

States/input $\boldsymbol{x}(t) = \left[x(t) v_x(t)\ y(t)\ v_y(t)\right]^T$, $\boldsymbol{u}(t) = \left[F_x(t) F_y(t)\right]^T$

Linear dynamics: $\dot{\boldsymbol{x}}(t) = A\boldsymbol{x}(t) + B\boldsymbol{u}(t)$

Artificial slip: $v_y(t) \in [s_{\min}, s_{\max}] v_x(t)$

States/input: $\boldsymbol{x}(s) = \left[t(s)\ z_x(s)\ y(s)\ \tilde{v}_y(s)\right]^T$, $\boldsymbol{u}(s) = \left[u_x(s), u_y(s)\right]^T$

State space: $\boldsymbol{x}'(s) = A\boldsymbol{x}(s) + B\boldsymbol{u}(s)$

Artificial slip: $y'(s) \in [s_{\min}, s_{\max}]$

III: Relative velocity: If the object we want to measure the distance to is relative, it might be beneficial to switch to relative velocity: $\tilde{v}_x(x) = v_x(x) - v_L$, $\tilde{s} = s - v_l t(s)$, $\tilde{z} = 1/\tilde{v}_x$

States/input: $\boldsymbol{x}(s) = \left[\tilde{t}(s)\ \tilde{z}_x(s)\ y(s)\ \tilde{v}_y(s)\right]^T$, $\boldsymbol{u}(s) = \left[u_x(s), u_y(s)\right]^T$

State space: $\boldsymbol{x}'(\tilde{s}) = A\boldsymbol{x}(\tilde{s}) + B\boldsymbol{u}(\tilde{s})$

Artificial slip: $y'(\tilde{s}) \in [s_{\min}, s_{\max}](1 + v_L)\tilde{z}\,(\tilde{s})$

Applications/Result

Application to intersection problem: Remodeling longitudinal dynamics

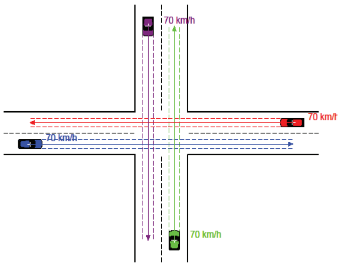

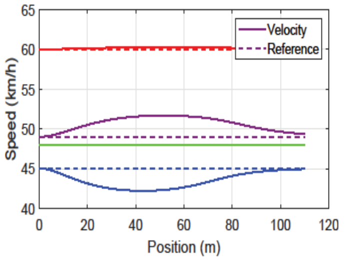

Application to the overtaking problem: Sampling in relative velocity**.**

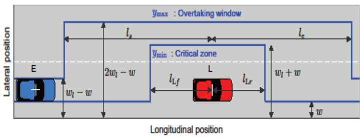

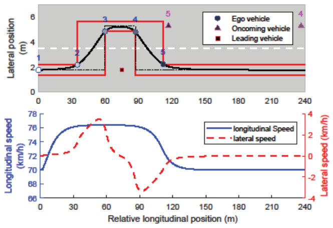

Current work

1. More advanced vehicle models
2. Crossing combinations:
3. Robust model predictive control

References

1. Karlsson Johan, "Computationally efficient exact remodeling of optimization programs with applications to autonomous driving". Licentiate thesis. Chalmers university, 2019

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Formal Verification of Learning-based Software in Safety-Critical Systems

LINKÖPING UNIVERSITY

John Törnblom
john.tornblom@liu.se

SAAB TECHNOLOGIES SAAB

Description

AI advances are now being applied in safety-critical systems where software defects may cause harm to humans and the environment. Machine learning models with large sets of parameters are difficult to interpret, and thus manufacturers of safety-critical systems are currently unable to provide convincing arguments that such models have been tested sufficiently before being deployed. As a complement to testing, formal verification methods are now being suggested and evaluated by the research community.

Background & Motivation

By the year 2020–2025, cars are expected to be running autonomously on public roads. However, organizations in safety-critical domains are currently unable to provide convincing arguments that complex software based on machine learning algorithms are safe and correct.

Type	Deaths per billion		
	Journeys	Hours	km
Bus	4.3	11.1	0.4
Rail	20	30	0.6
Car	40	130	3.1
Foot	40	220	54.2
Air	117	30.8	0.05
Motorcycle	1640	4840	108.9

Table 1: Transportation statistics from the United Kingdom 1990–2000 (Wikipedia)

Research Goal & Questions

Our goal is to understand existing formal methods, and develop new methods specifically for machine learning models.

- How can existing formal methods be tailored for supervised machine learning models, and do these methods scale to realistically sized verification problems?
- What are the trade-offs between different machine learning models when verifiability is important?

Methods & Preliminary Results

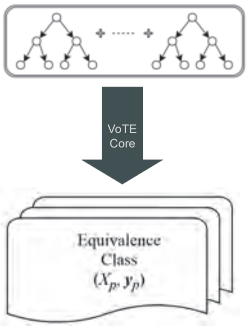

- Efficient method to partition input domain of decision trees, and explore all path combinations in tree-ensembles.
- An abstraction-refinement method that counteracts combinatorial explosion when exploring path combinations in tree-ensembles.
- VoTE, an implementation of the methods for random forests and gradient boosting machines.

Roadmap & Milestones

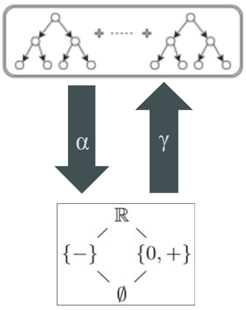

- A mathematical formalization of the methods in an abstract interpretation framework.
- Soundness and completeness proofs of the formalized algorithms.
- Application of the methods in a safety-critical case study relevant to industry.

Overview of VoTE

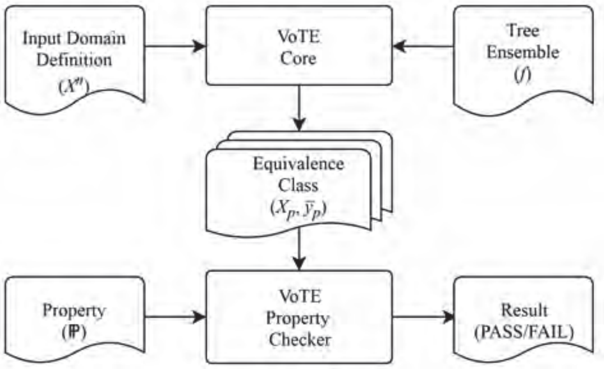

Discovered Counter Example

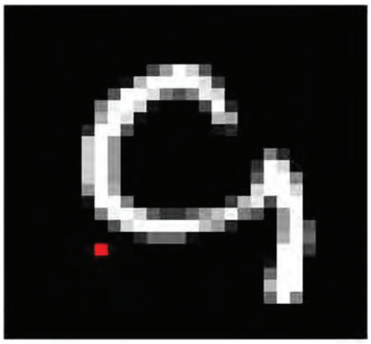

- One of many counter-examples discovered by VoTE while verifying robustness against noise in tree-ensembles trained on the MNIST dataset.
- Since the added noise is invisible to the naked eye, the noise (a single pixel incremented by one) is highlighted in red.

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

STAMINA: Stream Processing in the AMI

Joris van Rooij, Chalmers & Göteborg Energi

CHALMERS UNIVERSITY OF TECHNOLOGY

Distributed Computing and Systems
Chalmers university of technology

Göteborg Energi

Adding smart electricity meters to a power distribution network doesn't make the grid smart when they are only used for monthly readings. This project investigates how the hundreds of thousands of data streams from the smart meters in the Advanced Metering Infrastructure can be used to enable close to real time monitoring of the low voltage distribution grid and to make the Smart Grid smart. The project also investigates how the computing resources in the AMI can be utilized efficiently by using edge computing. Among the methods used are distributed and parallel stream processing and edge/fog computing.

Background & Motivation

The introduction of the Smart Meter greatly increased the amount of metering data. Yearly readings turned into hourly. High-resolution power quality data from smart meters can enable unprecedented control in the low voltage network. This control will be necessary when increasing amounts of electricity usage and renewable production transform the traditional power flow in the network.

Research Goal & Questions

Low Voltage
Supervisory
Control
And
Data
Acquisition

Design methods to:
- Validate readings in (close to) real time
- Detect unwanted situations such as outages, broken equipment or theft.

- Gain insights in how to provide near real time information for supply-demand matching.
- Where to process what: utilize cloud/edge computing

Methods & Preliminary Results

Flink

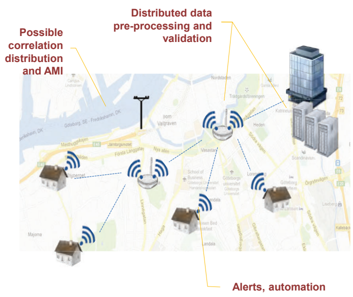
Distributed data pre-processing and validation
Possible correlation distribution and AMI
Alerts, automation

The main method we use to harness the large amount of data is distributed and parallel stream processing using Apache Flink.
A prototype validation engine showed good results for real time validation for 300.000 meters. The engine detected both single events (negative or too large consumption) as well as composite events where single events appear in specific patterns. (published at IEEE EnergyCon 2018)

Voltage readings for meters that are connected to the same transformer show high correlation. This can be leveraged to detect faulty meters at a very early stage by comparing readings with other meters connected to the same transformer.
The computations can run centrally or distributed at the data collectors. (published at DEBS2018)

Future Research

Advanced validation
Identifying negative or larger than possible readings is easy. More challenging to find suspicious readings.
Idea: Compare readings on different levels in the distribution grid

Total Control
Readings can be aggregated in time and space to get information that is useful for network control
Idea: Distributed aggregation by data collector, grid substation or by cable.

In Progress: Efficient Out of Order Stream Processing

- Ordered data is important to get consisted results in stream processing
- However late data is not uncommon for Smart Meters
- State of the art approach is to cache all data until all data has arrived → Large memory requirements
- Our work, TinTiN leverages knowledge of missing data to only store data that is relevant for late arrivals
  - Provides swift results for on-time data
  - Processes late arrivals as soon as they arrive
  - Minimizes memory requirements

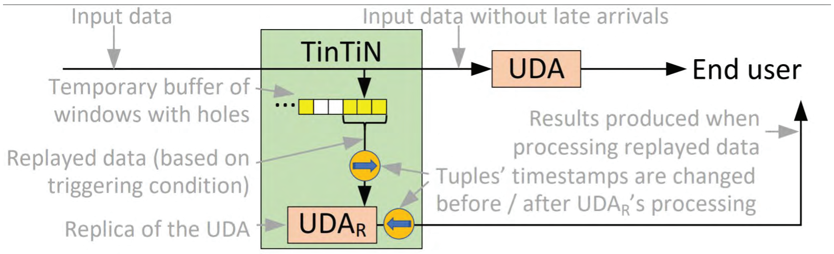
Input data
Input data without late arrivals
TinTiN
UDA
End user
Temporary buffer of windows with holes
Replayed data (based on triggering condition)
Replica of the UDA
UDA_R
Results produced when processing replayed data
Tuples' timestamps are changed before / after UDA_R's processing

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

IMAGES OF **BREAST TUMOR** HELPS DETECTION OF **COLON CANCER**

Detection of colon cancer metastases in lymph nodes through deep learning[✶]

Karin Stacke, Apostolia Tsirikoglou**

Intro

- Colon cancer is the **4th most common cancer** type in Sweden
- Finding metastases is hard and time consuming
- Digitalization of pathology images enables image analysis applications to assist
- Deep learning methods require hard to come-by annotated data

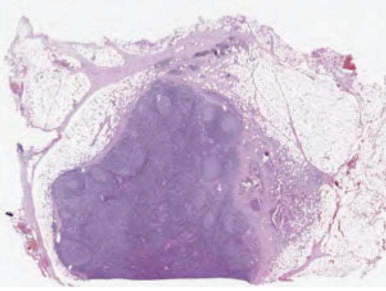

Colon lymph node tissue

Method

- **Cycle-GAN**[1]s was used to synthetically generate colon tumor images from breast tumor images (*breast2colon*)
- These images were used for augmenting original colon data[2] to train a DL classifier for tumor detection

Results

Accuracy on **colon tumor detection increased** when using *breast2colon* images together with original data, compared to a model trained only on limited amount of colon data.

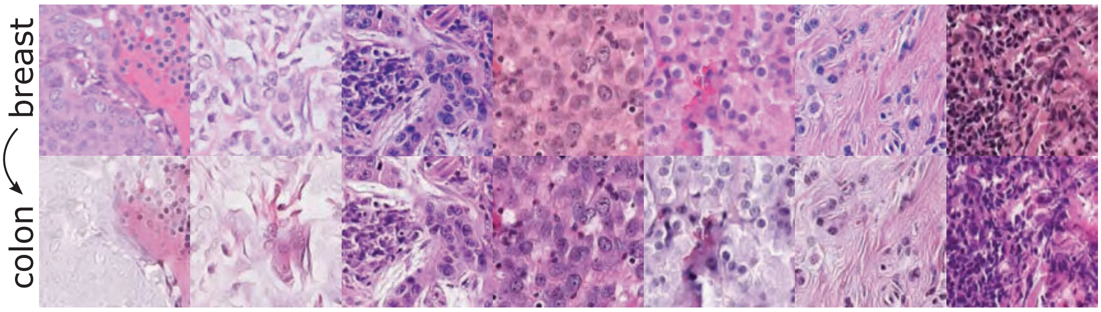

Example of Cycle-GAN transformed breast tumor images to colon tumor images

✶ WASP Project Course 2019, in collaboration with Martin Lindvall

* Equal contribution, Department of Science and Technology, Linköping University

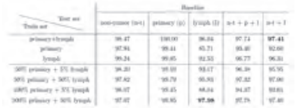

Test set / Train set	Baseline				
	non-tumor (n-t)	primary (p)	lymph (l)	n-t + p + l	n-t + l
primary+lymph	98.47	100.00	96.84	97.74	**97.41**
primary	97.94	99.41	65.71	95.40	92.60
lymph	99.24	99.85	92.53	96.77	96.31
50% primary + 5% lymph	98.20	99.59	93.17	96.39	95.95
50% primary + 50% lymph	97.82	99.79	95.93	97.32	97.00
100% primary + 5% lymph	98.07	99.45	88.84	94.37	93.81
100% primary + 50% lymph	97.07	98.85	**97.98**	97.78	97.40

Table 1: Patch-level accuracy of the non-tumor/tumor classifiers trained on full-size (*upper*) and fractions (*lower*) of the total dataset. Non-tumor notation is omitted from the train set column, as it is always present as one of the two classes.

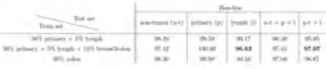

Test set / Train set	Baseline				
	non-tumor (n-t)	primary (p)	lymph (l)	n-t + p + l	n-t + l
50% primary + 5% lymph	98.20	99.59	93.17	96.39	95.95
50% primary + 5% lymph + 10% breast2colon	97.12	100.00	**96.62**	97.41	**97.07**
80% colon	98.30	99.50	94.56	97.10	96.42

Table 4: Patch-level accuracy for ***breast2colon*** experiments.

1. Zhu et al.,. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In: 2017 IEEE International Conference on Computer Vision (ICCV). 2017. p. 2242–51.

2. Maras G, Lindvall M, Lundstrom C. Regional lymph node metastasis in colon adenocarcinoma. 2019; Available from: doi:10.23698/aida/lnco

LINKÖPING UNIVERSITY

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Wireless Link Adaptation with outdated CSI – a hybrid data-driven and model-based approach

L. Pellaco, V. Saxena, M. Bengtsson, J. Jaldén

pellaco@kth.se

Motivation

- Wireless spectrum is a **scarce resource**
- Maximizing the **spectral efficiency** (information rate sent over a given bandwidth) is crucial
- **Transmission parameters** affect the spectral efficiency
- The optimal transmission parameters depend on the **instantaneous *channel state information* (CSI)**
- ***Link adaptation***: adjusting transmission parameters according to the CSI to maximize the spectral efficiency
- The instantaneous CSI is usually estimated through a **pilot sequence**

Challenges

- The optimal transmission parameters are **not a simple function** of the CSI
  We address it in a **DATA-DRIVEN** fashion
- The CSI at the base station might be **outdated** due to the time delay (**feedback delay**) between pilot transmission, CSI reception, and data transmission
  We address it in a **MODEL-BASED** fashion

System Model

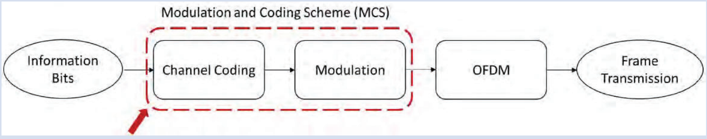

- We select the modulation and coding scheme (MCS) used for transmission
- The instantaneous CSI is the channel vector $\boldsymbol{h}$
- E_k represents the event of successful ($E_k = 0$) or unsuccessful ($E_k = 1$) frame decoding with the k^{th} MCS
- Our goal: maximize the spectral efficiency

$$\max_k \frac{T_k(1 - P_{E_k|\boldsymbol{H}}(E_k = 1|\boldsymbol{h}))}{B_k}$$

- k is the MCS index
- T_k is the number of information bits sent with the k^{th} MCS
- B_k is the bandwidth occupied with the k^{th} MCS
- The probability of unsuccessful decoding (= error event) with k^{th} MCS given the CSI at transmission time is: $p_{k|\boldsymbol{h}} = P_{E_k|\boldsymbol{H}}(E_k = 1|\boldsymbol{h})$

Proposed Hybrid link adaptation

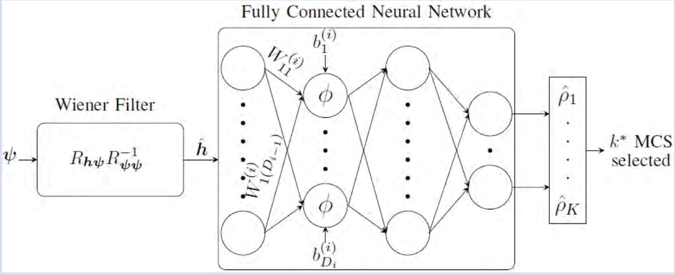

- We can realistically assume that the base station stores some channel history $\boldsymbol{\psi}$
- Use the Wiener filter (a linear minimum mean square error estimator) to give an estimate of $\boldsymbol{h}$, i.e., $\hat{\boldsymbol{h}}$
- $\boldsymbol{\psi}$ is a collection of past CSIs
- $R_{\boldsymbol{h}\boldsymbol{\psi}} = \mathbb{E}\{\boldsymbol{h}\boldsymbol{\psi}^T\}$, $R_{\boldsymbol{\psi}\boldsymbol{\psi}} = \mathbb{E}\{\boldsymbol{\psi}\boldsymbol{\psi}^T\}$, $\hat{\rho}_k = p_{k|\hat{\boldsymbol{h}}}$

End-to-end link adaptation

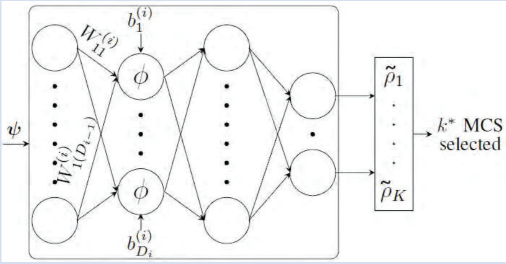

- $\boldsymbol{\psi}$ is a collection of past CSIs, $\tilde{\rho}_k = p_{k|\boldsymbol{\psi}}$
- Pitfalls: **increased dimensionality** and **lack of explainability**

Optimality of the Hybrid approach

If the channel evolves as a Gaussian random process, then $\hat{\boldsymbol{h}}$ is a sufficient statistic of $\boldsymbol{h}$ $\Rightarrow$ this two-part hybrid approach comes without loss of *optimality*

Numerical results

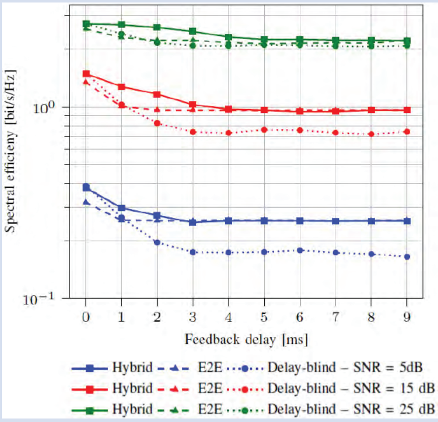

Pose Proposal Critic

Lucas Brynte, Fredrik Kahl
brynte@chalmers.se, fredrik.kahl@chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY

WASP
WALLENBERG
AUTONOMOUS SYSTEMS
AND SOFTWARE PROGRAM

Rigid Object Pose Estimation

- Estimate 6D pose θ for known rigid objects (position & orientation in 3D space).
- Input: 1 RGB image.
- Training data: Pose annotations & CAD models.

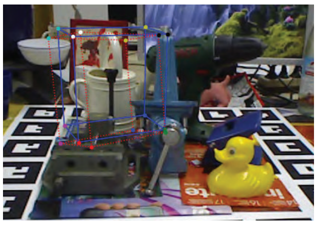

Estimated pose of a can

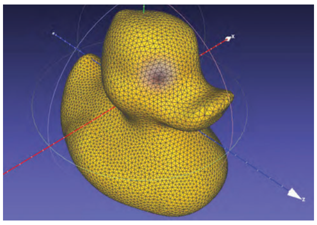

CAD model of a duck

Pose Refinement

- For RGB-D pose estimation, a refinement step using Iterative Closest Point (ICP) is common practice, aligning the observed point cloud with the CAD model.
- In the RGB-only case, it is less straight-forward how to refine the pose. Yet this can be badly needed, especially in occluded scenarios, where algorithms typically fail to reliably provide precise pose estimates.

Rendering-Based Refinement

- Recent work [2, 3] leverage on online rendering of synthetic images for pose refinement in the RGB-only case.
- A synthetic image of the CAD model is rendered, where the object is aligned according to some proposed pose estimate θ_k.
- A CNN is trained to estimate the relative pose $\Delta\theta$ between θ_k and the true underlying pose θ^* of the observed image. A refined pose estimate θ_{k+1} is then proposed by perturbing θ_k by $\Delta\theta$.

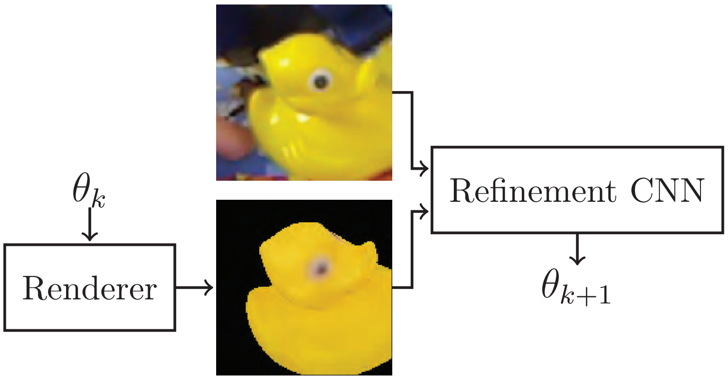

Proposed Pipeline

- A modified pipeline is proposed, with two main differences (marked in red in the figure):
  - The renderer is replaced by a differentiable renderer [1]. This will open up for new ways of refining the proposed pose estimate.
  - Instead of letting the CNN estimate the pose residual $\Delta\theta$, we train it to quantify its magnitude, i.e. estimating the value of some error function $J(\theta^*, \theta_k)$.
  - Finally, due to the whole pipeline being differentiable end-to-end, we may during test-time find our next pose estimate by searching for $\theta_{k+1} = \text{argmin}_{\theta^*} J(\theta^*, \theta_k)$ using gradient descent.

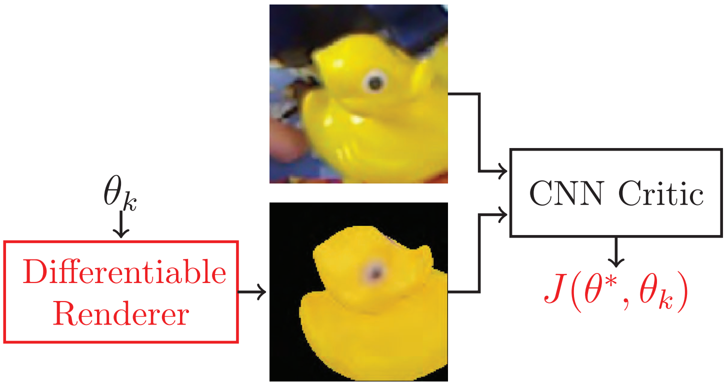

Features

- Great flexibility regarding what kind of error function J to estimate the values of.
  - Potential in letting J be a simple quantity highly related to the domain of observations, e.g. pixel-level errors such as the reprojection error.
- Robustness to generalization errors – a bias on J has zero impact, as long as its minimum remains unchanged.

Outlook

- Scrutinize the need for an initial proposed pose estimate. An efficient GPU-accelerated refinement pipeline could potentially defeat the purpose of such a proposal.

References

[1] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. "Neural 3D Mesh Renderer". In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2018.

[2] Yi Li et al. "DeepIM: Deep Iterative Matching for 6D Pose Estimation". In: *The European Conference on Computer Vision (ECCV)*. Sept. 2018.

[3] Fabian Manhardt et al. "Deep Model-Based 6D Pose Refinement in RGB". In: *The European Conference on Computer Vision (ECCV)*. Sept. 2018.

Upgrade Methods for Stratified Sensor Network Self-calibration

Martin Larsson, Lund University
Centre for Mathematical Sciences

Abstract

Estimating receiver and sender positions is often solved using a stratified, two-tiered approach. In the first step the problem is converted to a low-rank matrix estimation problem. The second step can be seen as an affine upgrade. This affine upgrade is the focus of this paper. In the paper new efficient algorithms for solving for the upgrade parameters using minimal data is presented. It is also shown how to combine such solvers as initial estimates, either directly or after a hypothesis and test step, in optimization of likelihood. The system is verified on both real and synthetic data.

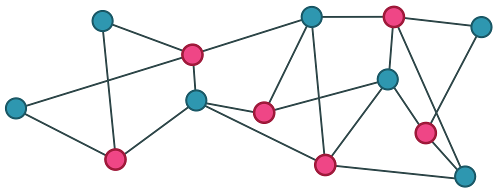

Problem

We wish to find the position of the receivers $\boldsymbol{r}_i$ and senders $\boldsymbol{s}_j$ given distance measurements $\boldsymbol{d}_{ij}$ such that

$$\boldsymbol{d}_{ij} = \|\boldsymbol{r}_i - \boldsymbol{s}_j\|.$$

This is the time of arrival (ToA) formulation of the problem but one could also consider similar problems such as time difference of arrival (TDoA).

What is common for these problems is that there is a relaxed problem

$$\boldsymbol{d}_{ij}^2 = -2\boldsymbol{u}_i^T\boldsymbol{v}_j + \boldsymbol{a}_j + \boldsymbol{b}_i,$$

where $\boldsymbol{a}_j$ and $\boldsymbol{b}_i$ only depend on the measurements. This is a low-rank approximation problem were the rank depends on the space $\boldsymbol{r}_i$ and $\boldsymbol{s}_j$ are embedded in.

A solution $(\boldsymbol{U},\boldsymbol{V})$ to the relaxed problem can be upgraded to a solution in $(\boldsymbol{R},\boldsymbol{S})$ using

$$\boldsymbol{R} = \boldsymbol{L}^{-T}\boldsymbol{U}, \qquad \boldsymbol{S} = \boldsymbol{L}(\boldsymbol{V} + \boldsymbol{q}),$$

where $\boldsymbol{L}$ and $\boldsymbol{q}$ are found by solving a system of polynomial equations [1]. Such systems can be solved using action matrix methods [2]. There are many possible systems that can be constructed to solve for $\boldsymbol{L}$ and $\boldsymbol{q}$ and consequently many different solvers can be constructed.

References

1. Y. Kuang, S. Burgess, A. Torstensson, and K. Åström, "A complete characterization and solution to the microphone position self-calibration problem," in International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2013.
2. H. Stewénius, Gröbner Basis Methods for Minimal Problems in Computer Vision, Ph.D. thesis, Lund University, APR 2005.

Results

Numerics for the solvers

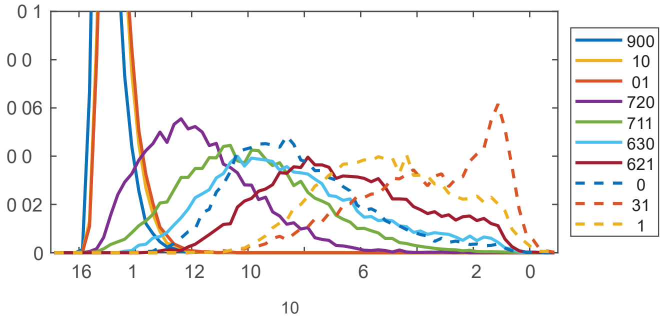

Numerics for the solvers when $\boldsymbol{r}_i$ and $\boldsymbol{s}_j$ are sampled from a sphere

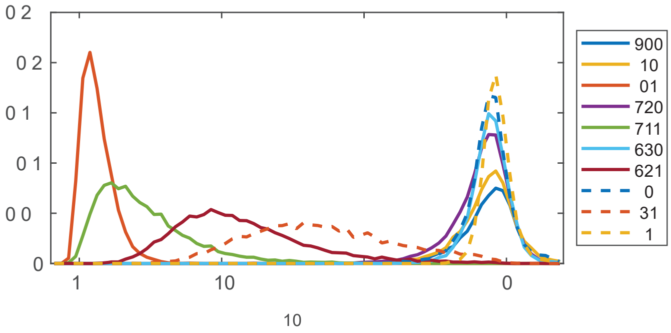

Solver	#solutions		Template size		Exec. time
	Sat.	No sat.	Sat.	No sat.	
900	1	-	-	-	39 μs
810	3	3	-	-	130 μs
801	4	4	-	-	130 μs
720	9	9	12 × 21	12 × 21	170 μs
711	12	12	16 × 28	16 × 28	210 μs
630	21	17	88 × 109	112 × 129	600 μs
621	30	26	122 × 152	156 × 182	1.2 ms
540	∞	21	-	310 × 331	5.3 ms
531	∞	38	-	493 × 531	19 ms
441	∞	42	-	817 × 859	72 ms

EFFICIENT CANCER DIAGNOSIS WITH A HUMAN IN THE LOOP

Designing support for lymph node tumor detection: Assisted search for rare phenomena

*Martin Lindvall**

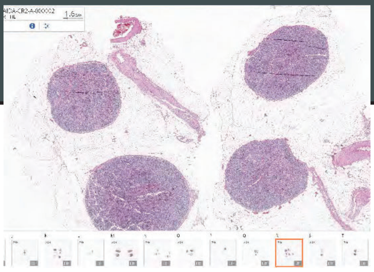

BACKGROUND: Deep neural networks have shown expert level accuracy in medical applications. The resulting predictive models have the potential to aid in decision-making if they can be integrated into clinical practice.

DIAGNOSTIC TASK: Pathologists visually assess around 40 lymph node sections to detect tumors. The number found impact the treatment of the patient

SOLUTION: AI-Assissted search with a human in the loop

1

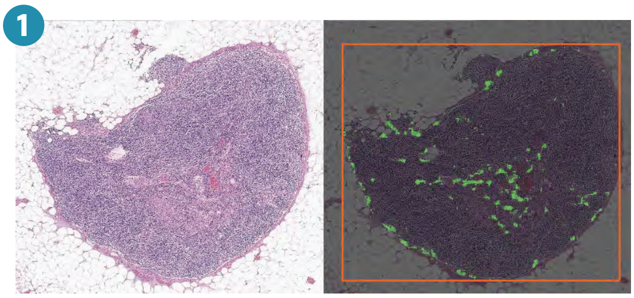

In preprocessing, sections are segmented and a tumor classification model, tuned to high sensitivity, is applied to each pixel. Candidate detection areas are segmented and scored according to probability of true positive.

2

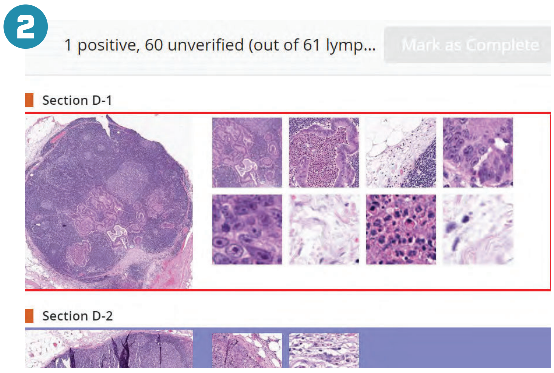

1 positive, 60 unverified (out of 61 lymp...

The user is presented with areas organized by section and ordered by tumor likelihood. Sensitivity can be tuned by requesting more or fewer areas.

3

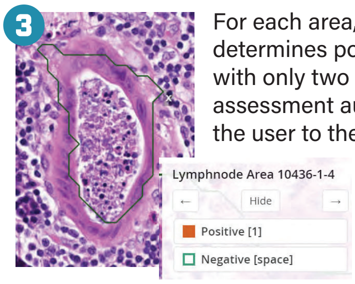

Lymphnode Area 10436-1-4

□ Positive [1]

□ Negative [space]

For each area, one rapidly determines positive or negative with only two keys. A positive assessment automatically moves the user to the next section.

EVALUATION: 60% time saved

Manual		AI-Assissted	
Time	Positive	Time	Positive
10m	8,3	4m	8,6

METHOD: Iterative design with data, model and user interface

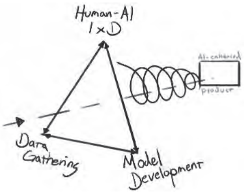

Slide collection, expert annotation, model development and interface design considerations were done iteratively. At each stage a prototype of suitable fidelity was developed and assessed qualitatively by pathologists.

In total we gathered 436 slides from 39 patients. We annotated 157 tumor sections and 270 normal sections. The model was trained using tensorflow with a CNN-architecture.

The latest prototype was quantitatively evaluated in a user study focusing on efficiency (time) and quality (accuracy). Ten cases were evaluated, with an average of 37 sections per case, by one user.

* ITN, LiU & Sectra AB. WASP Project course 2019, collaboration with Karin Stacke, Apostolia Tsirikoglou and Gordan Maras

LINKÖPING UNIVERSITY

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

SoK: Ambient Assisted Living Systems and their Privacy/Security Considerations

Md Sakib Nizam Khan[1] & Sonja Buchegger
email: (msnkhan[1], buc)@kth.se
KTH Royal Institute of Technology

Motivation & Research Goals

The confluence of several developments make privacy of Ambient Assisted Living (AAL) an increasingly important problem: aging population, scale and availability of sensors (IoT, health monitoring, smart home), and advances in data analysis and learning. Privacy research has not been in sync with these developments. Yet, privacy needs to be addressed to make the technology trustworthy to be adopted by a vulnerable population with sensitive health data. Numerous research efforts focus on AAL, but they are scattered. A clear understanding of the literature and commercial systems is needed. The aim of this research is to **systematize the knowledge** of **ambient assisted living systems** and their **privacy** and **security** situation to **identify** prominent **privacy concerns** in the domain.

Methods

Our Approach:

- ❖ Systematic literature review
  - ✓ Databases: IEEE Xplore & ACM
  - ✓ Search query:
    ((("health monitor*" OR "ambient assisted living" OR "e-health") AND (elder* OR senior OR old*) AND ("smart home" OR IoT)))
  - ✓ Performed on meta data (i.e. Title, Abstract, and Keywords)
  - ✓ Search Results
    - ▪ IEEE Xplore returned 69 articles
    - ▪ ACM returned 50 articles
  - ✓ Filtering
    - ▪ Title & abstract screening: 27 articles removed
    - ▪ Full content screening: 3 articles removed
    - ▪ Total 89 articles selected for thorough review
  - ✓ Questions for thorough review
    - ▪ What kind of frameworks or system architectures are proposed?
    - ▪ Does the proposed system or framework focus on any driving technology or component?
    - ▪ What data is collected (or which sensors are used) by the system?
    - ▪ Is there any security or privacy consideration?
- ❖ Survey on commercial systems
  - ✓ Performed on 8 different commercial systems

Systematization of Literature

- ❖ Based on the research focus, distinct categories identified during review
  - ✓ Literature is systematized across the categories

Research Focus	Number of Articles
Framework or Prototype for AAL	34
Sensors, Actuator, User Interface	13
Communication	9
Data Processing	16
Survey Review	12
Security	3
Business Model or AAL	2

Challenges in AAL

- ❖ Data collection
  - ✓ Collects both health and environmental data of the patient
  - ✓ A lot more things can be inferred from such data
- ❖ User interface
  - ✓ System intended for elderly people
  - ✓ Difficult to take consent from elderly people in such a setup
- ❖ Data disclosure
  - ✓ Data sanitization techniques apply random noise
  - ✓ Open research question: How to apply privacy-enhancing technologies on such data without hampering the utility provided by the data?

Findings

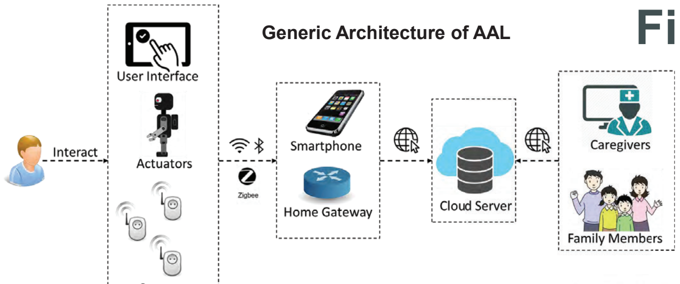

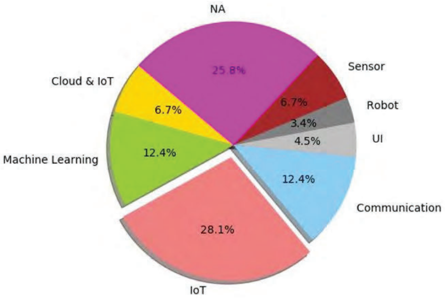

Technology Focus

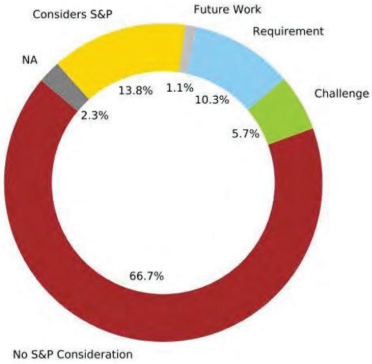

Security & Privacy Considerations

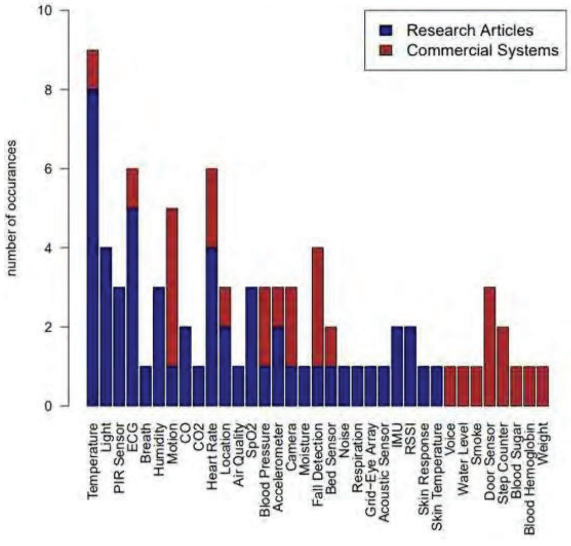

Sensors used in AAL Systems

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Learning robust LQ-controllers using application oriented exploration

Mina Ferizbegovic, Jack Umenberger, Håkan Hjalmarsson and Thomas B. Schön

UPPSALA UNIVERSITET

Summary and contributions

This work is concerned with the problem of minimizing the worst-case quadratic cost for an uncertain linear dynamical system. We propose:

- a robust controller to minimize the worst-case LQ cost, with probability $1-\delta$, which can be solved by a convex SDP,
- an approximate dual controller that simultaneously regulates the system and reduces model uncertainty, which can be solved by a convex SDP.

Problem statement

- Linear time-invariant dynamics

$$x_{t+1} = Ax_t + Bu_t + w_t \; w_t \sim \mathcal{N}\left(0, \sigma_w^2 I\right)$$

- System parameters (A, B) are unknown.
- Given $\mathcal{D} = \{x_t, u_t\}_{t=1}^{n}$, we can denote the model $\mathcal{M}(\mathcal{D}) = \{\hat{A}, \hat{B}, D\}$.
- D is a measure of uncertainty.
- The estimation error $\Delta^\top = \left[\hat{A} - A,\ \hat{B} - B\right]$.
- Static state-feedback law $u_t = Kx_t + e_t$, where $e_t \sim \mathcal{N}(0, \Sigma)$.
- Our decision variable is control policy $\mathcal{K} = \{K, \Sigma\}$.
- The worst-case cost

$$J(\mathcal{K}, \mathcal{M}) = \max_{A,B} \lim_{\tau \to \infty} \frac{1}{\tau}\mathbb{E}\left[\sum_{t=1}^{\tau} x_t^\top Q x_t + u_t^\top R u_t\right]$$
$$\text{s.t. } x_{t+1} = Ax_t + Bu_t + w_t,\ u_t = \mathcal{K}_i(x_t),\ \Delta^\top D \Delta \leq I$$

'Robust controller:'

Given initial data $\mathcal{D}_0$, find a policy $\mathcal{K}$ that minimizes the worst-case cost $\min_{\mathcal{K}} J(\mathcal{K}, \mathcal{M}(\mathcal{D}_0))$.

'Dual controller'

Two competing objectives:

1. minimize the worst-case cost (exploitation)
2. gather information about the system (exploration)

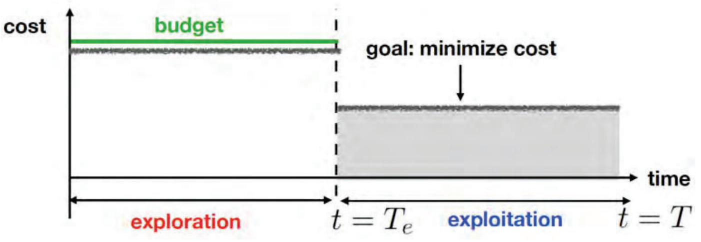

Minimize the worst-case cost of redesigned controller after exploration, subject to an exploration budget J_{exp}.

$$\min_{\mathcal{K}_0, \mathcal{K}_1} (T - T_e) \times J(\mathcal{K}_1, \mathcal{M}(\mathcal{D}_1))$$
$$\text{s.t. } T_e \times J(\mathcal{K}_0, \mathcal{M}(\mathcal{D}_0)) \leq J_{exp}$$
$$\mathcal{D}_1 = \{\mathcal{D}_0, \{x_t, u_t\}_{t=T_e}^{T}\}$$

Modeling uncertainty

We will work with **models** of the form $\mathcal{M}(\mathcal{D}) := \{\hat{A}, \hat{B}, D\}$ where D specifies the confidence region centered about $\{\hat{A},\ \hat{B}\}$:

$$\Theta_m(\mathcal{M}) = \{A,\ B\ :\ \Delta^\top D \Delta \preceq I,\ X = [\hat{A} - A,\ \hat{B} - B]^\top\}$$

For $D = \frac{1}{c}\sum_{t=1}^{n-1} \begin{bmatrix} x_t \\ u_t \end{bmatrix} \begin{bmatrix} x_t \\ u_t \end{bmatrix}^\top$, the true system $[A,\ B] \in \Theta_m(\mathcal{M})$ w.p. $1 - \delta$.

Designing a robust controller

For **known** A and B the covariance $W = \mathbb{E}\left[x_t x_t^\top\right]$ satisfies:

$$W \succeq [A\ B] \begin{bmatrix} W & WK^\top \\ KW & KWK^\top + \Sigma \end{bmatrix} [A\ B]^\top + \sigma_w^2 I. \qquad (1)$$

We introduce the **change of variables** $Z = WK^\top$ and $Y = KWK^\top + \Sigma$, collated in the variable $\Xi = \begin{bmatrix} W & Z \\ Z^\top & Y \end{bmatrix}$.To ensure that (1) holds for all $\{A, B\} \in \Theta_m(\mathcal{M})$ we use a theorem similar to S-procedure and can write (1) as an LMI $S(\lambda, \Xi, \hat{A}, \hat{B}, D) \geq 0$. $S(\lambda, \Xi, \hat{A}, \hat{B}, D)$ is **linear in** D (measure of uncertainty).

Theorem 1. *The problem* $\min_{\mathcal{K}}\ J(\mathcal{K}, \mathcal{M})$ *can be solved by a convex SDP.*

Designing a dual controller

Propagate uncertainty

Recall that: $D_1 = D_0 + \frac{1}{c}\sum_{t=1}^{T_e} \begin{bmatrix} x_t \\ u_t \end{bmatrix} \begin{bmatrix} x_t \\ u_t \end{bmatrix}^\top$. We use the approximation:

$$\bar{D}_{\text{dc}} = \sum_{t=1}^{T_e} \begin{bmatrix} x_t \\ u_t \end{bmatrix} \begin{bmatrix} x_t \\ u_t \end{bmatrix}^\top \approx T_e \begin{bmatrix} W_0 & Z_0 \\ Z_0^\top & Z_0^\top W_0^{-1} Z_0 + \Sigma_0 \end{bmatrix} \approx T_e \Xi_0$$

$S(\lambda_1, \Xi_1, \hat{A}_1, \hat{B}_1, D_1)$ is **linear in** D_1.

$$(\hat{A}_1, \hat{B}_1) = (\hat{A}_0, \hat{B}_0)$$
$$\min_{\mathcal{K}_0, \mathcal{K}_1}\ (T - T_e) \times \texttt{worst-case-cost}(\mathcal{K}_1, \mathcal{M}(\mathcal{D}_1))$$
$$\text{s.t.}\quad T_e \times \texttt{worst-case-cost}(\mathcal{K}_0, \mathcal{M}(\mathcal{D}_0)) \leq \text{budget}$$

Theorem 2. *The problem of a dual controller can be solved by a convex SDP, if we approximate* ***future*** *nominal parameter estimates with the* ***current*** *estimates.*

Numerical simulations

'Comparison of robust controller *(rc-prop)* with SLS robust controller *(rc-sls)* [1]'

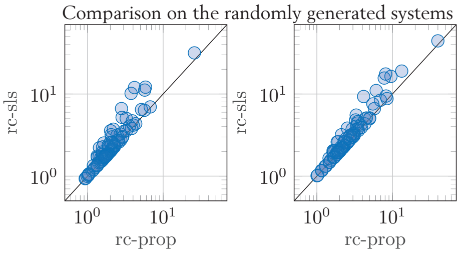

i) true system dynamics ii) the worst-case cost

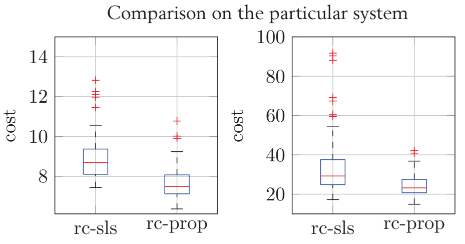

i) true system dynamics ii) the worst-case cost

'Comparison of dual controller *(dc)* with greedy exploration *(exp)*'

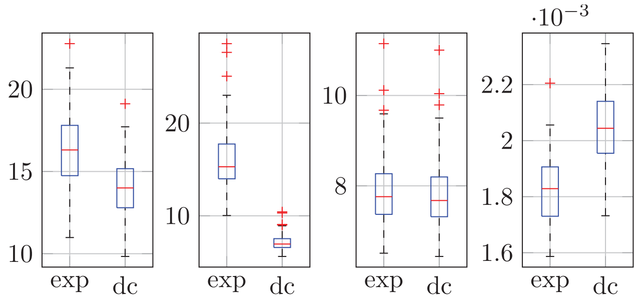

From left to right: Comparison of i) worst-case cost of exploitation phase, ii) cost of exploration, iii) true cost of exploitation phase, iv) measure of uncertainty.

- The uncertainty achieved by **dc** is larger than that achieved by **exp**.
- The performance of the worst-case cost of exploitation phase is lower.
- **dc** is reducing the uncertainty in a **structured** way, targeting directions which 'matter most for control'.

Future reading

[1] Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *to appear in Foundations of Computational Mathematics (arXiv:1710.01688)*, 2017.

[2] Mina Ferizbegovic, Jack Umenberger, Håkan Hjalmarsson, and Thomas B Schön. Learning robust lq-controllers using application oriented exploration. *IEEE Control Systems Letters*, 4(1):19–24, 2019.

Online Learning with Linear Constraints

Vidit Saxena, vidits@kth.se
KTH Royal Institute of Technology Ericsson Research

Background

- Online learning agents operate in *a priori* unknown environment to optimize a well-defined objective.
- Online learning problems are commonly modeled as sequential decision problems, where past actions and rewards guide the choice of the current action.
- We study *multi-armed bandits* (MAB), which model online learning problems with i.i.d. rewards. MABs have been applied in online advertising, dynamic pricing, and wireless rate selection.

Problem Description

- We study online problems with Bernoulli-distributed rewards (for action i, an i.i.d. reward r_i is obtained with probability μ_i.)
- The goal is to maximize the cumulative reward while keeping the reward probability above η.

$$\text{maximize } \mathbb{E}\Big[\sum_{t=1}^{T} r_{i(t)}\mu_{i(t)}\Big]$$

$$\text{subject to } (1/T)\mathbb{E}\Big[\sum_{t=1}^{T} \mu_{i(t)}\Big] > \eta.$$

Proposed Approach

Thompson Sampling

- For MABs *without* constraints, Thompson Sampling (TS) is a widely-used heuristic.
- TS assigns priors $\tilde{\mu}_i(0)$ for the unknown parameters, and calculates the posteriors $\tilde{\mu}_{i(t)}(t)$ after each time step.
- For TS, the loss in the optimam cumulative reward (the *regret*) is bounded below $O(\log(T))$.

Constrained Thompson Sampling

- We propose Linearly Constraint TS (LinConTS) for the studied problem.
- At each time step, LinConTS solves a linear program using the latest parameter estimates.
- We show that the regret as well as constraint violations scale as $O(\log T)$.

Algorithm 1 LinConTS

```
Input: Reward Values $r_{\{1,\dots,N\}}$, Constraint $\eta$
Initialize: $\alpha_{\{1,\dots,N\},0} = 1, \beta_{\{1\dots N\},0} = 1$.
for Time index $t = 1$ to $T$ do
    if $t < N$ then
        $i(t) = t$
    else
        Sample $\tilde{\mu}_{i,t} \sim \text{Beta}(\alpha_{i,t-1}, \beta_{i,t-1})$ for each arm $i = 1, \dots, N$.
        Solve, if feasible, the linear program:
            $LP(\tilde{\mu}_t) : \text{maximize} \sum_i x_{i,t}\tilde{\mu}_{i,t} r_i$
            subject to $\begin{cases} \sum_i x_{i,t}\tilde{\mu}_{i,t} \geq \eta \\ \sum_i x_{i,t} = 1 \\ x_{i,t} \geq 0 \quad \forall i \in \{1,\dots,N\} \end{cases}$ ,    (6)
        if a (feasible) optimal solution existed then
            Sample $i(t) \sim [x_{1,t}, \dots, x_{N,t}]$
        else
            Sample $i(t)$ uniformly from $\{1,\dots,N\}$.
        end if
    end if
    Observe: Reward event $c_{i(t)} \in \{0,1\}$.
    Update:
        $\alpha_{i(t),t} = \alpha_{i(t),t-1} + c_{i(t)}$
        $\beta_{i(t),t} = \beta_{i(t),t-1} + (1 - c_{i(t)})$.
end for
```

Performance Metrics

- Expected Cumulative Regret

$$\mathcal{R}(T) = \mathbb{E}\Big[Tr^* - \sum_{t=1}^{T} r_{i(t)} x_{i(t)}\Big]_+$$

- Expected Cumulative Violation

$$\mathcal{V}(T) = \mathbb{E}\Big[T\eta - \sum_{t=1}^{T} x_{i(t)}\Big]_+$$

Numerical Results

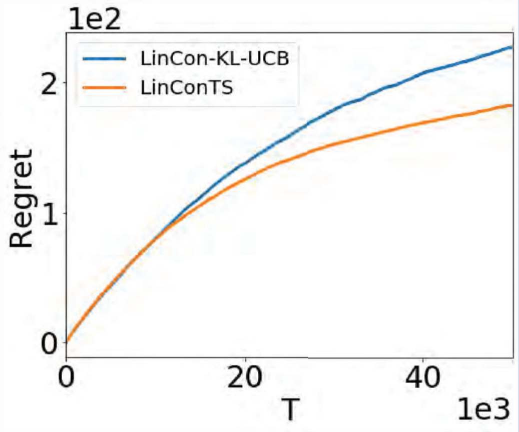

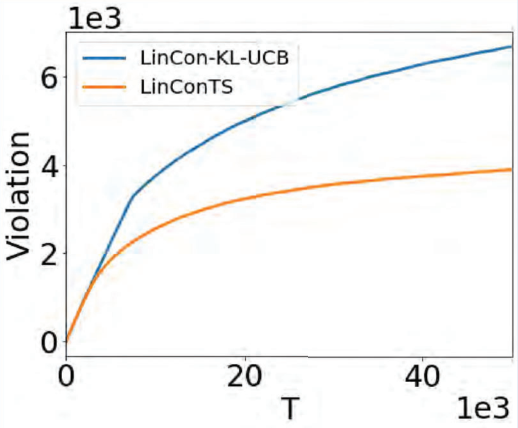

Multi-Camera Extrinsic Calibration from Human Pose

Olivier Moliner, Sony, Lund University
Centre for Mathematical Sciences
Main advisor: Kalle Åström

Motivation & Research goals

The generalization of multi-camera setups in the public space enables novel applications in video analysis, surveillance and human-computer interaction. Accurate multi-view image analysis requires camera systems to be thoroughly calibrated, which usually involves manually collecting calibration data, a time-consuming process that must be repeated if the cameras are moved. The aim of this research is to study methods for extrinsic camera calibration based on 2D human pose detection, enabling automatic calibration of camera systems for simple and scalable installations, robust to camera pose changes.

Method

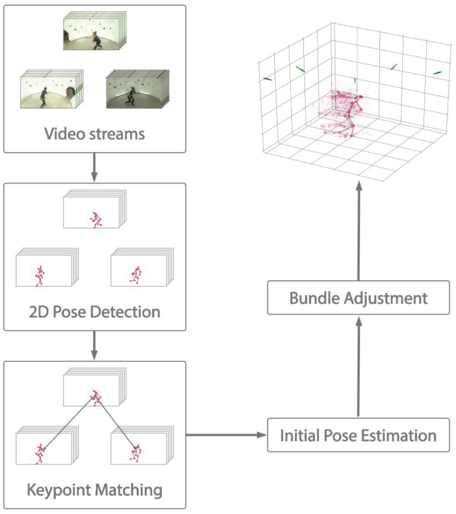

1. **Input:** video streams from intrinsically calibrated cameras
2. **Detect 2D human poses** in each frame, using e.g. [1]
3. **Match the joints** across all views and frames
4. **Initialize the estimated camera poses and 3D feature point** positions using conventional structure-from-motion techniques
5. **Refine** the reconstruction with **bundle adjustment**

Main challenge: 2D human pose estimation has large pixel error (up to tens of pixels)

References

1. Z. Cao *et al.*, "OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields," *IEEE Trans. Pattern Anal. Mach. Intell.*, Jul. 2019.
2. J. Puwein *et al.*, "Joint Camera Pose Estimation and 3D Human Pose Estimation in a Multi-camera Setup," *ACCV 2014*.
3. G. Pavlakos *et al.*, "Expressive Body Capture: 3D Hands, Face, and Body from a Single Image," CVPR 2019.
4. Y. Chen *et al.*, "Adversarial PoseNet: A Structure-Aware Convolutional Network for Human Pose Estimation," ICCV *2017*
5. J. Butepage *et al.*, "Deep representation learning for human motion prediction and classification, CVPR 2017

Human Body Priors for Bundle Adjustment

Approach: to compensate errors in pose estimation, integrate human body constraints and motion models in the bundle adjustment optimization [2].

A simple smooth motion model

We assume that the objects in the scene follow a **near-constant velocity** motion model with Gaussian acceleration noise

$$z_{j_{t+1}} = F z_{j_t} + w_t$$

$$F = \begin{bmatrix} I & I\Delta t \\ 0 & I \end{bmatrix} \qquad w_t \sim \mathcal{N}(0, Q)$$

To enforce the 3D points to follow this motion model, we add a motion term to the standard bundle adjustment cost function:

$$E = (1-\alpha) \underbrace{\sum_{i,j,t} \nu_{ijt} \frac{1}{\sigma^2} \left\| (u_{ij_t} - \pi_i(U_{jt})) \right\|^2}_{\text{Reprojection error}} + \alpha \underbrace{\sum_{j,t} \left\| L_m(z_{j_{t+1}} - F z_{j_t}) \right\|^2}_{\text{Motion error}}$$

Results on synthetic data

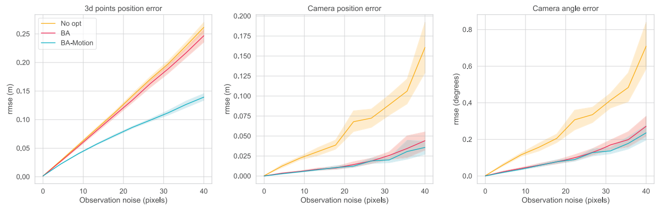

Ongoing & Future Work

- ► **Body pose prior**
  Use a generative model to estimate a probability density function for human poses [3,4], and penalize less probable poses during bundle adjustment.

- ► **Body motion model**
  Learn a model for human motion (e.g. [5]) and use its predictions as regularizer for bundle adjustment.

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Developing Tools and Analyze Methods for Secure Software Update

Pegah Nikbakht Bideh, Martin Hell

Department of Electrical and Information Technology, Lund University, Sweden

Introduction

Nowadays software is built on many OSS external components. Due to:

- Increased number of sophisticated attacks in software
- Increasing attention to identifying security vulnerabilities in the last decade
- Increasing number of software intensive systems
- Shift towards cloud computing
- Increasing popularity of IoT devices

It is important to handle the process of identifying, evaluating and prioterizing new vulnerabilities, in order to be able to patch or update the software ontime.

For that, we proposed a new maturity model which can be used to identify and evaluate vulnerabilities in software development companies using OSS components. Also, we proposed a new recommender system for prioritizing identified vulnerabilities based on user preferences.

We further tried to do an actual OTA update in an IoT environment using CoAP and MQTT protocols to see how security can affect the energy consumption of IoT devices.

Background

Maturity models: are used to assess processes/structures, and objects/technology within an organization. A maturity model can be seen as a tool that helps organizations improve the way they work. Maturity models will help organizations identify the issues in need for improving and prioritizing the efforts.

Recommender systems': goal is to present a set of recommendations of items to a set of users. The recommenders are widely used today (Netflix, Spotify, and Amazon). Recommenders have differernt kinds of design:

- Knowledge-based systems
- Content-based systems
- Collaborative filtering

IoT protocols: CoAP and MQTT are the most common application layer protocols used in IoT environments. CoAP uses UDP at transport layer, while MQTT uses TCP as transport protocol. To add security at transport layer, the natural choice is to use DTLS for CoAP and TLS for MQTT. There are four security modes such as Nosec, PreSharedKey, RawPubicKey and Certificates which can be used for both CoAP and MQTT protocols.

HAVOSS: A Maturity Model for Handling Vulnerabilities in Third Party OSS Components

The existed and related maturity models are very broad and cover many aspects related both to software development, maintenance, and organizational aspects, but they are not detailed enough to cover all aspects of handling vulnerabilities in third party code

HAVOSS: A maturity model focusing on managing vulnerabilities in third party libraries and code, and the subsequent software update activities that are required to limit a product's exposure to attacks

Our model is not a replacement for the other models, it is a complement to other maturity models. The model inclues capability areas below:

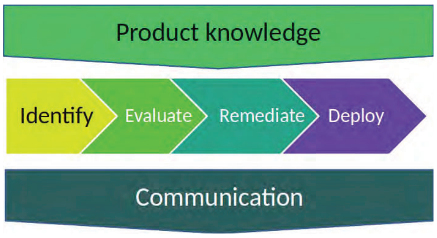

Product knowledge

Identify → Evaluate → Remediate → Deploy

Communication

Our model can represent the maturity of an organization with 5 different maturity level from 0 to 4. The first level is level 0 means that no effort is spend at all. Level 1 means that the process is carried out in some way but it is often unclear how it is done. Level 2 means there are defined approaches and routines. Level 3 represents a state where there is a standardized process in place. Finally, level 4 means experiences are collected from using the standardized procedures.

A Recommender System for User-Specific Vulnerability Scoring

Common Vulnerabilitiy Scoring System (CVSS) is commonly used to score the severity of a vulnerability. But it does not take user preferences into account. The score is based on different properties which are weighted together by a formula. In our recommender system, we use some extra features than CVSS score (Impact metrics, Access vector, and etc) such as:

- Published date
- Metasploit exploits
- Google hits

The general overview of our recoomender system is shown below:

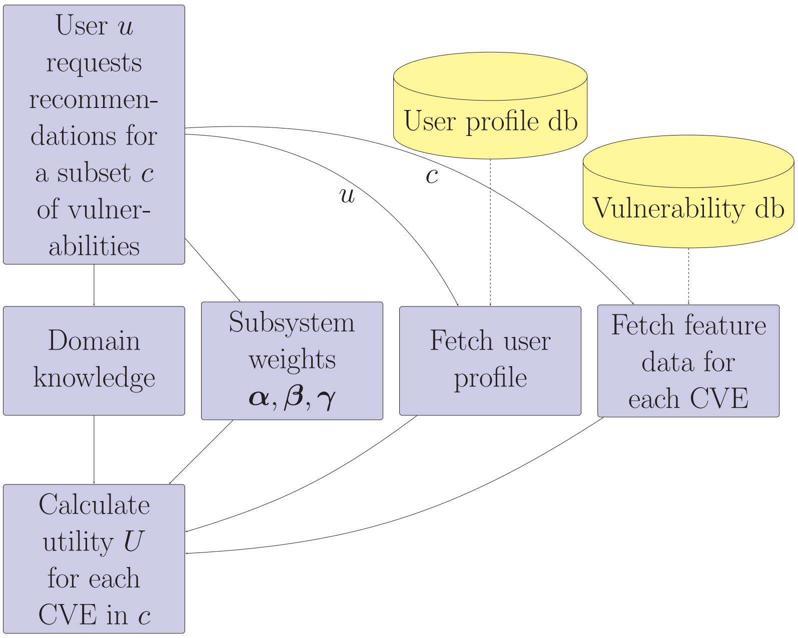

User u requests recommendations for a subset c of vulnerabilities

User profile db

Vulnerability db

u

c

Domain knowledge

Subsystem weights $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}$

Fetch user profile

Fetch feature data for each CVE

Calculate utility U for each CVE in c

Ongoing work: Analyzing Security Overhead of CoAP and MQTT Protocols

The ubiquitious nature of IoT devices often require them to run on batteries, making energy efficiency a primary concern. Adding security to the communication will add additional overhead. Thus, it is important to understand to which extent security affects the energy consumption of the devices in real use cases below:

- Sending various sizes of encrypted data
- Firmware update using OTA (Over-The-Air) update

We considred differernt factors in our energy consumption measurements such as different AES modes of operations, key sizes and differernt packet loss rates.

Results indicated that adding security to CoAP has much more overhead than MQTT for large payload sizes.

Future Works

- We aim to use new technologies such as SDN (Software Defined Networking) for key distribution which can also be used in IoT environments.
- We aim to develop a method for software updates in autonomous environments, with considering several aspects such as how critical the update is? how it needs to be securely deployed? what are the limitations of autonomous systems in comaprision to non-autonomous systems during update procedure?

WASP | WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Prescribed Performance Control Guided Policy Improvement for Satisfying Signal Temporal Logic Tasks

P. Varnai, D. V. Dimarogonas

varnai@kth.se

Motivation

Our goal is to develop a **practical learning framework** that allows dynamical systems to learn to satisfy **complex spatiotemporal tasks optimally** as well as **adjust to changes** in them in real-time.

Problem description

Consider the system:

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{g}(\boldsymbol{x})\boldsymbol{u} + \boldsymbol{w}$$

where only $\boldsymbol{g}(\boldsymbol{x})$ is known. Given:

- a task description ϕ and
- a cost $C(\boldsymbol{x}_{[0,T]}, \boldsymbol{u}_{[0,T]})$,

find control actions $\boldsymbol{u}$ which minimize the cost $C(\cdot)$ while satisfying ϕ with robustness $\rho^\phi > \rho_{\text{goal}}$.

Signal Temporal logic

- Extends Boolean logic with temporal operators
- Logical predicates μ_i stem from functions of system signals:

$$\mu_i := \begin{cases} \text{true if } h^{\mu_i}(\boldsymbol{x}) \geq 0 \\ \text{false if } h^{\mu_i}(\boldsymbol{x}) < 0 \end{cases}$$

- More complex expressions are formed recursively:

$$\phi = \top \mid \mu_i \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 U_{[a,b]} \phi_2 \mid F_{[a,b]}\phi \mid G_{[a,b]}\phi$$

- Robustness metric ρ^ϕ quantifies degree of satisfaction

Example: satisfying $F_{[0,8]}G_{[0,\infty]}\mu$ with the controller:

$$\boldsymbol{u}(\boldsymbol{x},t) := \kappa(\boldsymbol{x},t)\frac{K}{\|\boldsymbol{v}(\boldsymbol{x})\|_2^2 + \Delta}\boldsymbol{v}(\boldsymbol{x}), \quad \boldsymbol{v}(\boldsymbol{x})^{\mathrm{T}} := \frac{\partial \rho^\mu(\boldsymbol{x})}{\partial \boldsymbol{x}}\boldsymbol{g}(\boldsymbol{x})$$

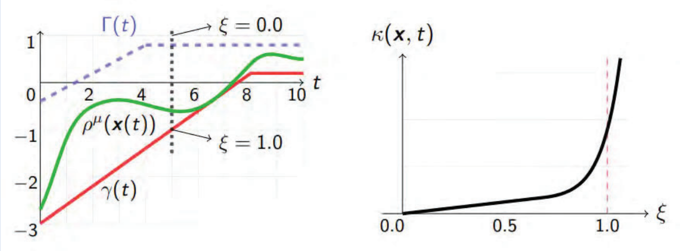

Approach: Guided PI^2

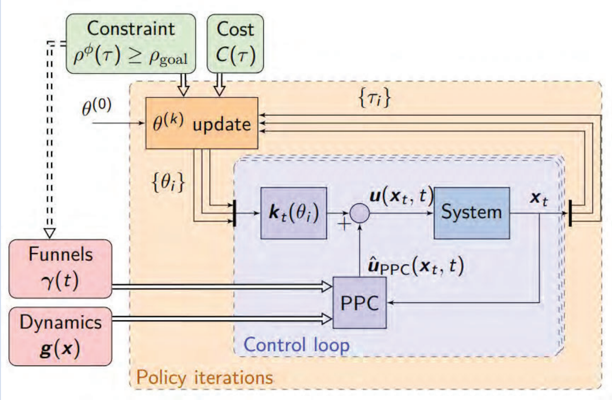

- The so-called PI^2 learning algorithm aims to find the parameterized **feedforward** terms $\boldsymbol{k}_t(\theta)$
- Prescribed performance control (PPC) aids satisfying the STL task in a **feedback** manner

Visualization:

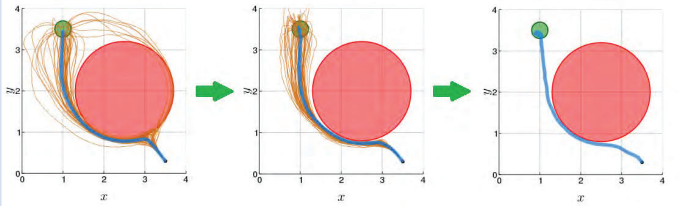

Sample results

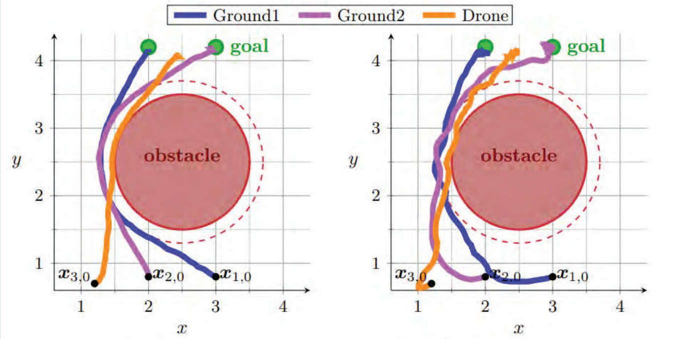

(a) G-PI^2: $\rho^\phi = 0.006$, $C = 7.61$ (b) PI^2: $\rho^\phi = -0.061$, $C = 10.07$

Ongoing work

- New guidance controllers
- New robustness metrics
- Funnel adaptation
- Multi-agent variant

Compressed Gradient Methods for Hessian-Aided Error Compensation

Sarit Khirirat, Sindri Magnússon and Mikael Johansson

School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology

Parallel and distributed optimization is the computational workhorse of large-scale ML and signal processing. As problem sizes continue to grow, communication is becoming the major performance bottleneck. Gradient compression reduces communication, but impairs solution accuracy. It has been empirically shown that gradient compression errors can be compensated for, but existing error-compensation schemes have very limited theoretical support. To quantify such improvements, we develop intuition on quadratic problems and propose Hessian-aided error compensation, which outperforms existing schemes in numerical evaluations.

Parallel and distributed ML

$$\min_{x\in\mathbb{R}^d} \; f(x) = \frac{1}{n}\sum_{i=1}^{n} f_i(x)$$

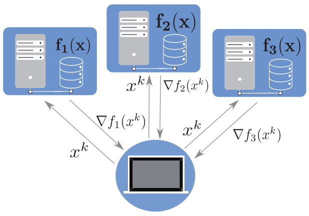

$$x^{k+1} = x^k - \frac{\gamma}{n}\sum_i \nabla f_i(x^k)$$

- Empirical risk minimization problems with finite-sum structure
- Each f_i describes risk (loss) of x on data stored by worker i
- Huge communication to transmit high-dimensional gradients
- Performance bottleneck shifts from computation to communication

Motivation

- Compressors save communicated bits, but impairs accuracy
- Error compensation improves both solution time and solution accuracy
- Novel theoretical understanding of the benefit of error compensation

Error Compensation on Quadratic Problems

Develop insight into error accumulation by studying quadratic problems

$$\underset{x\in\mathbb{R}^d}{\text{minimize}} \quad \frac{1}{2}x^T Hx + b^T x.$$

The iterates $\{x^k\}_{k\in\mathbb{N}}$ generated by compressed gradient descent

$$x^{k+1} = x^k - \gamma Q\left(\nabla f(x^k)\right)$$

satisfy

$$x^k - x^\star = \underbrace{A_\gamma^k(x^0 - x^\star)}_{\text{linear rate}} + \gamma \cdot \underbrace{\sum_{j=0}^{k-1} A_\gamma^{k-1-j}(\nabla f(x^j) - Q(\nabla f(x^j)))}_{\textbf{Accumulation of previous compression errors}},$$

while the iterates of the error-compensated gradient descent

$$x^{k+1} = x^k - \gamma Q(\nabla f(x^k) + A_\gamma e^k)$$
$$e^{k+1} = \nabla f(x^k) + A_\gamma e^k - Q(\nabla f(x^k) + A_\gamma e^k)$$

with $A_\gamma = I - \gamma H$ satisfy

$$x^k - x^\star = A_\gamma^k(x^0 - x^\star) + \underbrace{\gamma e^k}_{\textbf{No accumulation}}.$$

SGD with Hessian-Aided Error Compensation

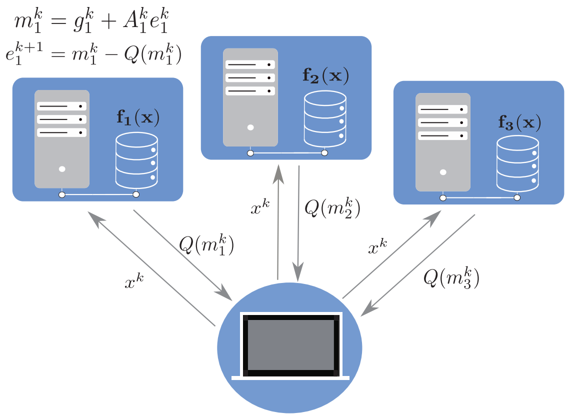

$$x^{k+1} = x^k - \frac{\gamma}{n}\sum_i Q(m_i^k)$$

Note: SGD with direct compression (C-SGD) is EC-CSGD with $e_i^k = 0$.

Convergence Analysis

Assumptions:

- Each function $f_i(\cdot)$ has L-Lipschitz continuous gradient
- g_i^k satisfies $\mathbf{E}[g_i^k] = \nabla f_i(x^k)$ and $\mathbf{E}\|g_i^k - \nabla f_i(x^k)\|^2 \le \sigma^2$
- $A_i^k = I - \gamma H_i^k$ with $\mathbf{E}[H_i^k] = \nabla^2 f_i(x^k)$ and $\mathbf{E}\|H_i^k - \nabla^2 f_i(x^k)\|^2 \le \sigma_H^2$

Non-convex problems: If $\gamma < 1/(3L)$, then C-SGD satisfies

$$\min_{l\in[0,k]} \mathbf{E}\|\nabla f(x^l)\|^2 \le \tfrac{A}{k+1}\tfrac{1}{\gamma}(f(x^0) - f(x^\star)) + B\epsilon^2 + \gamma \cdot C\sigma^2,$$

while EC-CSGD satisfies

$$\min_{l\in[0,k]} \mathbf{E}\|\nabla f(x^l)\|^2 \le \tfrac{A}{k+1}\tfrac{1}{\gamma}(f(x^0) - f(x^\star)) + \boldsymbol{\gamma^2} \cdot B\epsilon^2 + \gamma \cdot \tilde{C}\sigma^2,$$

for positive constants $A = 2/\beta, B = 3L/\beta, C = (1 + 3L\gamma)/\beta$, and $\tilde{C} = \alpha/\beta$ where $\beta = 1 - 3L\gamma$ and $\alpha = L^2 + 2(1 + 3L\gamma)(\sigma_H^2 + L^2)$.

Note: EC-CSGD error can be made arbitrarily small by decreasing γ.

Experimental Results (Least-Squares)

We highlight strong performance of Hessian-aided error compensation.

- `EC-Vr.1` is EC-CSGD with $A_i^k = I$
- `EC-Hessian` is EC-CSGD with $A_i^k = I - \gamma H_i^k$
- `EC-diag-Hessian` is EC-CSGD with $A_i^k = I - \gamma\text{diag}(H_i^k)$

Here, H_i^k is the Hessian associated with g_i^k at iteration k.

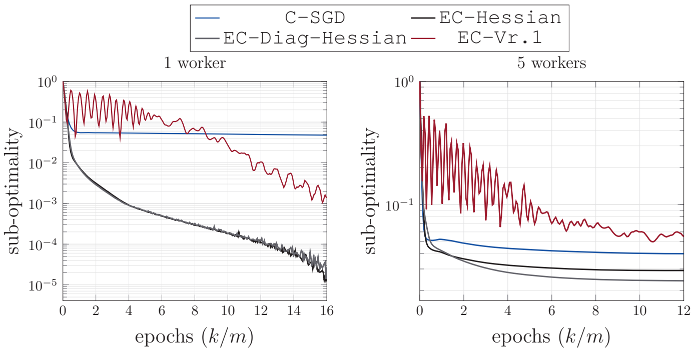

Fig. 1. Performance of algorithms when the binary compressor is used.

INTELLIGIBLE EXPLANATIONS IN INTELLIGENT SYSTEMS

Sule Anjomshoae and Kary Främling

Department of Computing Science, Umeå University

Introduction

EXPLAINABLE ARTIFICIAL INTELLIGENCE (XAI): Artificial intelligence whose actions and decisions are understandable by humans.

- **Motivations**
  -General Data Protection Regulation (GDPR) (right to explanations).
  -Increasing trust and transparency for intelligence systems.
- Previous works are mostly dealing with **data-driven XAI**.
- We provide a systematic literature review on **goal-driven XAI** [1].
- Then, present the Contextual Importance and Utility method for generating and communicating intelligible explanations [2].

Background

EXPLAINABILITY IN GOAL-DRIVEN DOMAIN:

- Explaining the actions and decisions of autonomous agents and robots.
- Inform user to understand the capabilities and limits of the agents.
- Explanations lead to better human-agent collaboration.
- Other terms; readability, legibility, explicability, and predictability.

EXPLAINABILITY IN DATA-DRIVEN DOMAIN:

- Understanding predictions made by machine learning algorithms.
- Mostly for experts to evaluate the accuracy of the predictions.
- Other terms; comprehensibility, justification, and intelligibility.

Systematic Literature Review

A systematic literature review is provided to gain insights into how current works solve the problem of generating and communicating intelligible explanations. Highlights from the systematic literature review are listed below;

Generating Explanations:

- The most widely used explanation type is **introspective informative explanations** that are based on the reasoning process which leads to a decision.
- Several studies suggested generating explanation facilities based on practically relevant **social science concepts**.

Communicating Explanation:

- Most studies selected **single modality** to communicate explanations (e.g. text-based, visuals, speech).
- A few works addressed the issues of **personalization and context-awareness**.

Evaluating Explanations:

- Most of the studies lack evaluation or tackle simple scenarios.
- The number of empirical studies is more than the user-based evaluations probably due to limitations in time and subject availability to conduct a user study.

CIU Method

The utility of the Contextual Importance and Utility (CIU) [3] method for providing intelligible explanations are briefly summarized below;

Explanation presentation:

- CIU values can be represented as text, visuals, or images.
- The variability in representing explanations could improve the interaction quality.

Complete and Contrastive explanations:

- CIU allows generating both complete and contrastive explanations.
- Complete explanations are the causes of an individual prediction or a decision.
- Contrastive instances are why a certain outcome is more probable than another.

Personalized explanations:

- CIU values can be represented with different levels of details based on the intended user.
- CIU allows providing explanations that are tailored to users' specification (e.g. explanations for patient vs physician).

Explanations of Black-Box Model Predictions by CIU

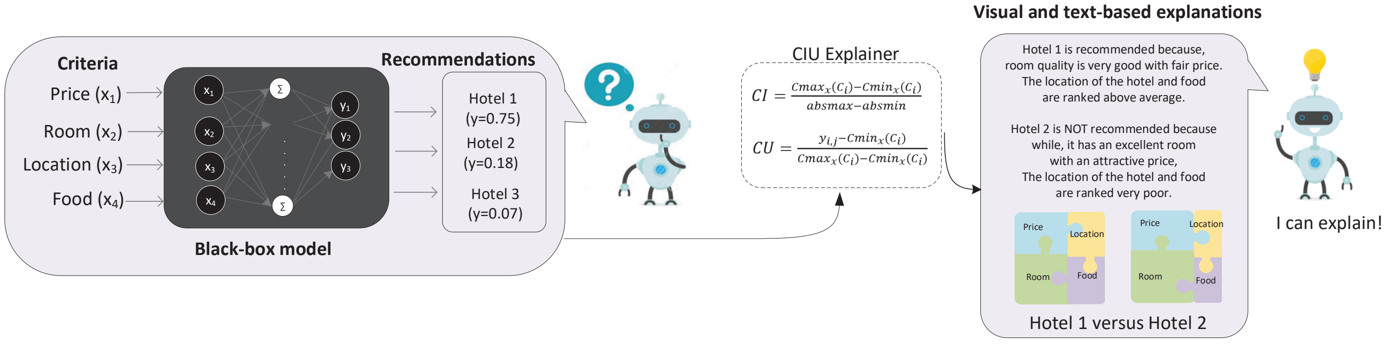

$$CI = \frac{Cmax_x(C_i) - Cmin_x(C_i)}{absmax - absmin}$$

$$CU = \frac{y_{i,j} - Cmin_x(C_i)}{Cmax_x(C_i) - Cmin_x(C_i)}$$

Hotel 1 is recommended because,
room quality is very good with fair price.
The location of the hotel and food
are ranked above average.

Hotel 2 is NOT recommended because
while, it has an excellent room
with an attractive price,
The location of the hotel and food
are ranked very poor.

Future Research Directions

This extended abstract summarizes our efforts towards providing intelligible explanations for intelligent systems. Future work will focus on;

- integrating the CIU explanation method in practical applications.
- producing personalized explanations by considering user's characteristic.
- investigating the usability of the explanations in real-world settings.

References

[1] Anjomshoae, S., Najjar, A., Calvaresi, D., Främling, K., Explainable Agents and Robots: Results from a Systematic Literature Review. in Proceedings of the 18th International Conference on Autonomous Agents and Multi Agent Systems. 2019.

[2] Anjomshoae, S., Främling, K., Najjar, A., Explanations of Black-Box Model Predictions by Contextual Importance and Utility. In EXTRAAMAS 2019.

[3] Främling, K. and D. Graillot. Extracting Explanations from Neural Networks. In Proceedings of the ICANN. 1995. Citeseer.

UMEÅ UNIVERSITY

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

Department of Computing Science, SE-901 87 Umeå, www.umu.se

Explicit-duration Hidden Markov Model on DNA Base-calling

Xuechun Xu

chunx@kth.se

Department of Information Science and Engineering

DNA Base-calling

Nanopore base-caller

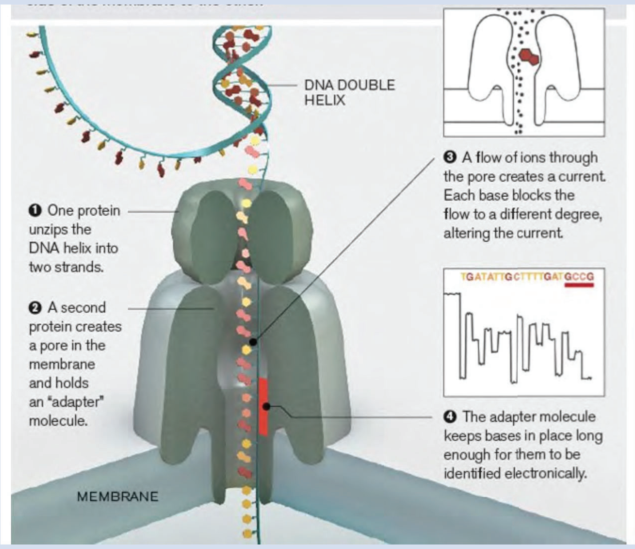
DNA DOUBLE HELIX
❶ One protein unzips the DNA helix into two strands.
❷ A second protein creates a pore in the membrane and holds an "adapter" molecule.
❸ A flow of ions through the pore creates a current. Each base blocks the flow to a different degree, altering the current.
TGATATTGCTTTTGATGCCG
❹ The adapter molecule keeps bases in place long enough for them to be identified electronically.
MEMBRANE

Basecalling - a challenging problem

- "Basecalling for ONT devices is the process of translating this raw signal into a DNA sequence. It is not a trivial task as the electrical signals come from single molecules, making for noisy and stochastic data."
  – Wicket al., "Performance of neural network basecalling tools for Oxford Nanopore sequencing", Genome Biology (2019)

Problematic

A piece of measured signals

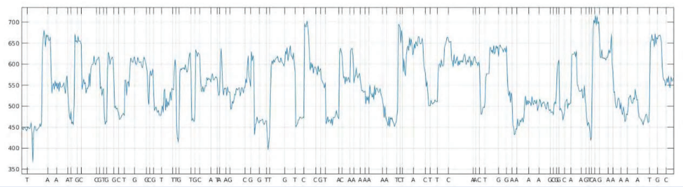

Synchronization

- *DNA go through pore with non-constant speed*
  - The motor protein drags DNA as it wants, i.e. unpredictable
  - The electron fieldwork push the charges DNA against the motor protein to slow down the process
- *Underlying DNA sequence of length L gives N current measurements*
  - Require base–signal level alignment to create training data

Measurements

- The current measurements are affected by K consecutive bases, which leads to model the states/classes as *K-mer*, of which K usually taken 5 – 7
- Each K-mer with continuous current measure. Moreover, the rough inner structure of the biological pore causes large variances
- Unaware of DNA modifications by nature, e.g. methylation, in the data.

State of art is not ideal

- The leading technology is held by Oxford Nanopre, which claims to use DNN. However, the accuracy is not qualified for clinical usage.
- The speed for base-calling is penalizing a lot for high accuracy performance. Updating the hardware will increase the cost.

Method Proposal

Simple Hidden Markov Model

Use K-mer to represent the 4 nucleartide bases, i.e. state space of size 4^K

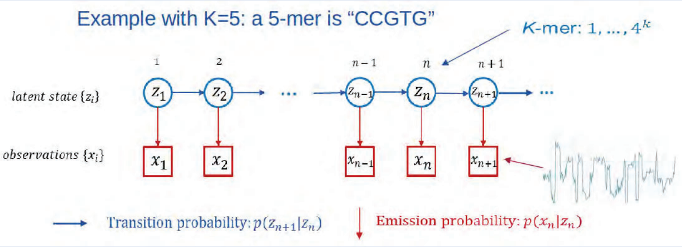

- Transition probability:

$$p(z_n = CGTGx|z_{n-1} = CCGTG)\ where\ x \in A, C, G, T$$

- Emission distribution: $p(x_n|z_n)$

Explicit-duration HMM

- Introduce duration variable $d \in 1, 2, ..., D$
- The state space extends to of size $4^K \times D$

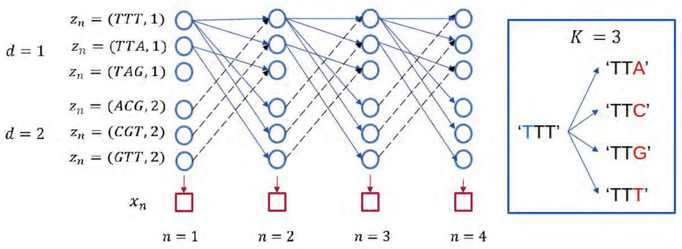

- Transition constraint by duration: $p(z_n = TTA, d-1|z_{n-1} = TTA, d)$
- Transition on next base only allowed at $d = 1$
- *Inference:* We can still use Viterbi or Forward-Backward algorithm to decode the train of hidden states

Training the ED-HMM

- Transition probabilities between K-mers can be learned by simple frequency count
- Duration probability and measurement model can be iteratively learned by Baum-Welch algorithm

Current results

Data: Ecoli and Lambda phage

- Trained with short pieces of Ecoli DNA(with methylation)and tested on Nvidia-dgx station (4 P100 GPU cards)

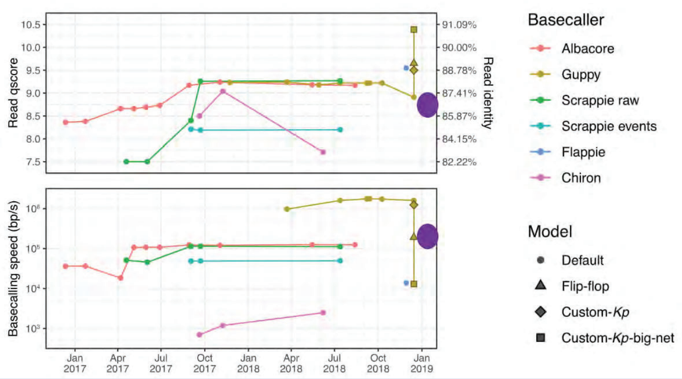
Read qscore
Read identity
Basecalling speed (bp/s)
10^6
10^5
10^4
10^3
Jan 2017
Apr 2017
Jul 2017
Oct 2017
Jan 2018
Apr 2018
Jul 2018
Oct 2018
Jan 2019
Basecaller
Albacore
Guppy
Scrappie raw
Scrappie events
Flappie
Chiron
Model
Default
Flip-flop
Custom-*Kp*
Custom-*Kp*-big-net