



# **WASP WINTER CONFERENCE 2022**

## **Poster Catalogue Autonomous Systems (AS)**

# WASP WINTER CONFERENCE 2022

## POSTER CATALOGUE 3/4

### AUTONOMOUS SYSTEMS [AS]

Author	Pages
Ahmad, Faseeh.....	102 A+B
Alshnakat, Anoud.....	103 A+B
Baravdish, Gabriel.....	104 A+B
Batkovic, Ivo.....	105 A+B
Brissman, Emil.....	106 A+B
Bruns, Leonard.....	107 A+B
Ceylan, Ciwan.....	108 A+B
Charitidou, Maria.....	109 A+B
Eryonucu, Cihan.....	110 A+B
Faris, Muhammed.....	111 A+B
Ferizbegovic, Mina.....	112 A+B
Forough, Javad.....	113 A+B
Fredriksson, Teodor.....	114 A+B
Gyllenhammar, Magnus.....	115 A+B
Hellander, Anja.....	116 A+B
Heskebeck, Frida.....	117 A+B
Hynén Ulfsjö, Carl.....	118 A+B
Iovino, Matteo.....	119 A+B
Jakobsson, Erik.....	120 A+B
Jensen, Maarten.....	121 A+B
Johansson, Simon.....	122 A+B
Johnander, Joakim.....	123 A+B
Jonnarth, Arvi.....	124 A+B
Kaalen, Stefan.....	125 A+B
Kampik, Timotheus.....	126 A+B
Khosravi, Hedieh.....	127 A+B
Krook, Jonas.....	128 A+B
Kullberg, Anton.....	129 A+B
Lapandić, Dženan.....	130 A+B
Larsson, Martin.....	131 A+B
Marta, Daniel.....	132 A+B
Mayr, Mattias.....	133 A+B
Moliner, Olivier.....	134 A+B
Mollevik, Iris.....	135 A+B
Narri, Vananda.....	136 A+B

# WASP WINTER CONFERENCE 2022

## POSTER CATALOGUE 3/4

### AUTONOMOUS SYSTEMS (AS)

Author	Pages
Nelson, Christian.....	137 A+B
Nielsen, Kristin.....	138 A+B
Nikbakht Bideh, Pegah.....	139 A+B
Nordlöf, Jonas.....	140 A+B
Nyberg, Truls.....	141 A+B
Oxenstierna, Johan.....	142 A+B
Parsa, Javad.....	143 A+B
Peng, Haori.....	144 A+B
Präntare, Fredrik.....	145 A+B
Ranawaka, Piyumal.....	146 A+B
Rasheed, Fahran.....	147 A+B
Rodriguez-Deniz, Hector.....	148 A+B
Rosdahl, Christian.....	149 A+B
Ruuskanen, Johan.....	150 A+B
Saleh Sedghpour, Mohammed Reza.....	151 A+B
Salt Ducaju, Julian M.....	152 A+B
Schuppe, Georg.....	153 A+B
Shoja, Shamisa.....	154 A+B
Song, Qunying.....	155 A+B
Svahn, Caroline.....	156 A+B
Varnai, Peter.....	157 A+B
Vladu, Emil.....	158 A+B
Wingqvist, Birgitta.....	159 A+B
Xie, Yiping.....	160 A+B
Xu, Xuechun.....	161 A+B
Zhu, Xiaomeng.....	162 A+B
Åström, Hampus.....	163 A+B

Ahmad, Faseeh  
Lund University

## Towards Generalized Robotic Skills and Knowledge Integration

Using skill-based systems to solve robotic tasks in industries has been gaining some popularity. These systems offer easily programmable robotic skills (modular software blocks) that are shareable among robots and rely on integration of planning, learning, sensing and execution. However, one big challenge is the design of a robotic skill (modular software block) that works in general settings. The aim of this research is to develop a framework to solve a variety of industrial tasks using generalized skills with knowledge integration. The framework involves dealing tasks with a lot of parameters and skills, a general way learning new tasks, and knowledge transfer among different hierarchical spaces.

# Towards Generalized Robotic Skills and Knowledge Integration

Faseeh Ahmad, Lund University

Department of Computer Science, Robotics and Semantic Systems (RSS)  
Advisors: Volker Krueger, Elin Anna Topp, Jacek Malek

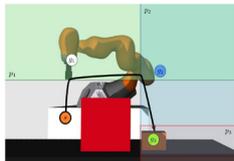


## Motivation & Research Goals

Using skill-based systems to solve robotic tasks in industries has been gaining some popularity. These systems offer easily programmable robotic skills (modular software blocks) that are shareable among robots and rely on integration of planning, learning, sensing and execution. However, one big challenge is the design of a robotic skill (modular software block) that works in general settings. The aim of this research is to develop a framework to solve a variety of industrial tasks using generalized skills with knowledge integration. The framework involves dealing tasks with a lot of parameters and skills, a general way learning new tasks, and knowledge transfer among different hierarchical spaces.

## Background

- Learning Parameters in Behavior Trees (BT) for Movement Skills[1]
  - ❖ BlackDrops[2]
  - ❖ CMAES
  - ❖ Behavior Trees
  - ❖ Obstacle Avoidance and Peg Insertion
- Learning Tacit Knowledge of Robot Tasks through Planning, Knowledge Integration and Multi-objective Optimization
  - ❖ SkiROS[3]
  - ❖ Bayesian Optimization (BO)
  - ❖ Hypermapper[4]
  - ❖ Peg Insertion and Push Task



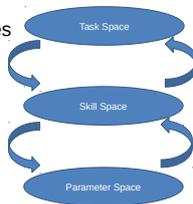
Obstacle Avoidance



Push Task

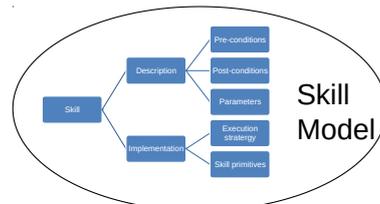
## Focus Points

- Solving tasks with more parameters and skills
- Generalized task specific learning
- Derivative skills
- Knowledge transfer between hierarchical spaces



## References

- Mayr, M., Chatzilygeroudis, K., Ahmad, F., Nardi, L. and Krueger, V., 2021. Learning of Parameters in Behavior Trees for Movement Skills. arXiv preprint arXiv:2109.13050.
- Chatzilygeroudis, K. and Mouret, J.B., 2018, May. Using parameterized black-box priors to scale up model-based policy search for robotics. In 2018 IEEE International Conference on Robotics and Automation (ICRA) (pp. 5121-5128). IEEE.
- Rovida, F., Grossmann, B. and Krüger, V., 2017, September. Extended behavior trees for quick definition of flexible robotic tasks. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 6793-6800). IEEE.
- Nardi, L., Koepflinger, D. and Olukotun, K., 2019, October. Practical design space exploration. In 2019 IEEE 27th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS) (pp. 347-358). IEEE.
- [https://en.wikipedia.org/wiki/Covariance\\_matrix](https://en.wikipedia.org/wiki/Covariance_matrix)

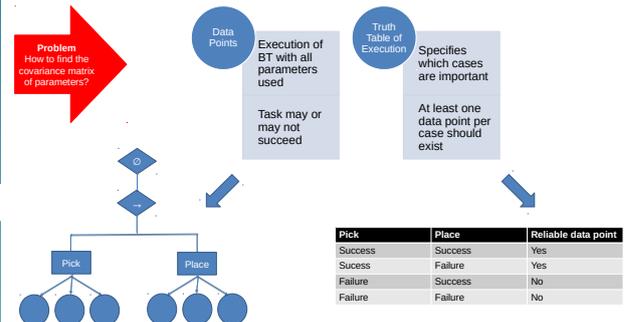
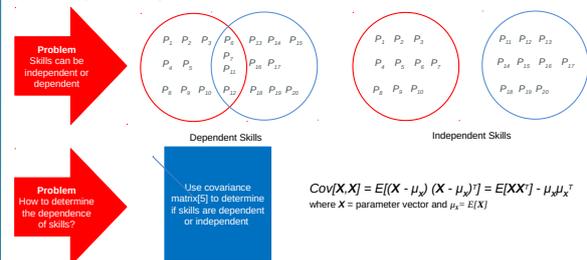


## Splitting Tasks

Learning a task with multiple skills and arbitrarily large parameters by maintaining a balance between quality and budget (time)

Case Study: Pick an object and place it at target location  
 Pick Skill: Parameters:  $P_1, P_2, \dots, P_{10}$   
 Place Skill: Parameters:  $P_{11}, P_{12}, \dots, P_{20}$

Solution: Divide the task into sub-tasks and learn the skill independently



## P4 Formalization & Verification

Programming Protocol-Independent Packet Processors (P4) is a Domain Specific Language used to program the data plane of networking targets as smart NICs and multi-port switches. The data plane contains two main programmable blocks called parser and control. We analyzed those and built a structural operational semantics (small step) model and executable semantics in order to prove interesting properties related to the processed packets as well as the P4 programs overall. The final outcomes of the project: a P4 interpreter verified down to the binary level, and automation to prove Hoare triple contracts of P4.



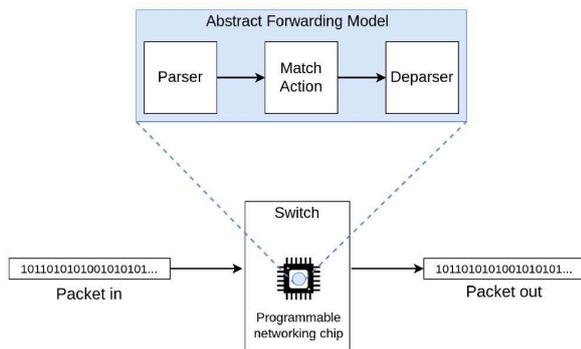
Anoud Alshnakat, Didrik Lundberg  
Roberto Guanciale and Mads Dam

KTH (Theoretical Computer Science Department)

### Introduction

Programming Protocol-Independent Packet Processors (P4) is a Domain Specific Language used to program the data plane of networking targets as smart NICs and multi-port switches. The data plane contains two main programmable blocks called parser and control. We analyzed those and built a structural operational semantics (small step) model and executable semantics in order to prove interesting properties related to the processed packets as well as the P4 programs overall. The final outcomes of the project: a P4 interpreter verified down to the binary level, and automation to prove Hoare triple contracts of P4.

### Generic P4 Architecture



Any packet enters the networking switch supported with a programmable data plane chip (programmed using P4) should be passed in the pipeline to three main stages:

**Parser:** A finite state machine that maps the bits in the input packet into a type representation. extracts the header fields bits from the packet. It can handle both standard header format as well as custom user-defined header format.

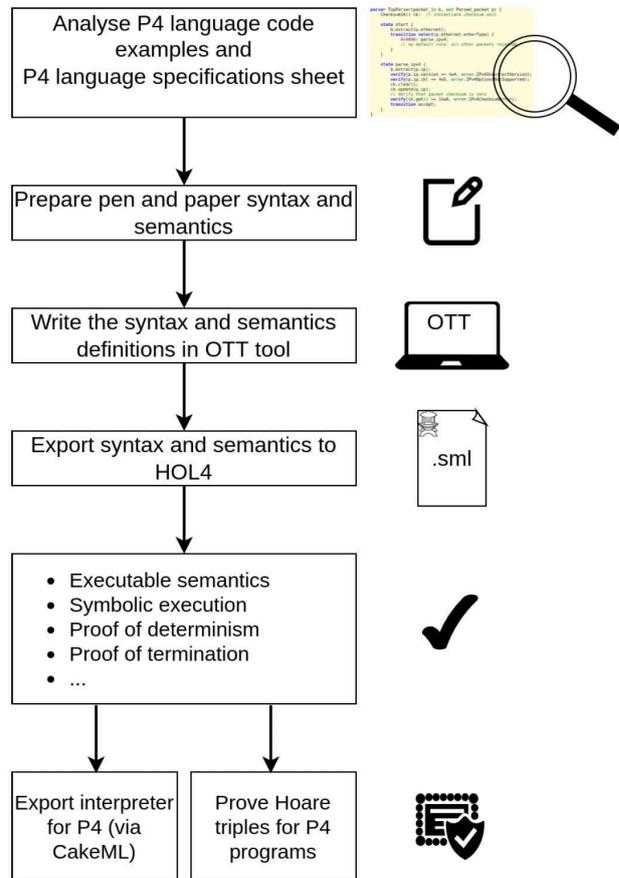
**Match Action tables:** One or more tables that contain keys and matching kinds, which determine the action to be processed on the packet. This stage requires interacting with the control plain.

**Deparser:** A control function that assembles the headers back into a well-formed output packet.

### References

1. "P4-16~ Language Specification," P4.org, 2017. <https://p4.org/p4-spec/docs/P4-16-v1.0.0-spec.html>
2. "ott-lang/ott: The Ott tool for writing definitions of programming languages and calculi," GitHub, Jul. 21, 2021. <https://github.com/ott-lang/ott>

### P4 Formalization



Baravdish, Gabriel  
Linköping University

## GPU Accelerated Sparse Representation of Light Fields

Light field imaging has been an omnipresent research topic during the last decade. With a growing interest, new techniques to capture, sample, and display light fields have been developed. The large amount of data that is produced during the capturing of light fields is a key challenge for acquisition and storage of light fields. We present a GPU-based compression technique based on multidimensional sparse representative. The main goal of the work presented here is to perform light field encoding, which includes an n-mode product, on the GPU and in real time.



# GPU Accelerated Sparse Representation of Light Fields

VISAPP 2019  
Prague 25-27 feb

Gabriel Baravdish, Ehsan Miandji, and Jonas Unger  
Department of Science and Technology, Linköping University

## Goal

Light field imaging has been an omnipresent research topic during the last decade. With a growing interest, new techniques to capture, sample, and display light fields have been developed. The large amount of data that is produced during the capturing of light fields is a key challenge for acquisition and storage of light fields. We present a GPU-based compression technique based on multidimensional sparse representation [3]. The main goal of the work presented here is to perform light field encoding, which includes a tensor-product, on the GPU and in real time. An overview of the method is presented in Figure 1 below.

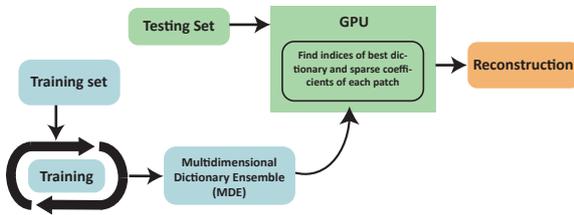


Figure 1: Flowchart of the method. Firstly we train the MDE. Secondly, as an input to the encoding process, we compute the sparse coefficients and membership indices segmentally in parallel. Lastly we reconstruct the patches.

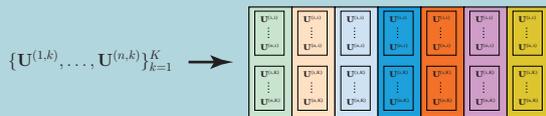
## Training

Initially from a training set of light field data, and as a one-time process, we train a multidimensional dictionary ensemble (MDE) which is obtained by a learning based approach [1].

### Training set



### Multidimensional Dictionary Ensemble - MDE



An MDE is obtained by minimizing:

$$\min_{\mathbf{U}^{j,k}, \mathcal{S}^{(j,k)}, \mathbf{M}_{i,k}} \sum_{i=1}^{N_i} \sum_{k=1}^K \mathbf{M}_{i,k} \left\| \mathcal{L}^{(i)} - \mathcal{S}^{(i,k)} \times_1 \mathbf{U}^{(1,k)} \dots \times_n \mathbf{U}^{(n,k)} \right\|_F^2$$

subject to

$$(\mathbf{U}^{(j,k)})^T \mathbf{U}^{(j,k)} = \mathbf{I}, \quad \forall k = 1, \dots, K, \quad \forall j = 1, \dots, n,$$

$$\left\| \mathcal{S}^{(i,k)} \right\|_0 \leq \tau_i$$

$$\sum_{k=1}^K \mathbf{M}_{i,k} = 1, \quad \forall i = 1, \dots, N_i$$

## GPU Encoding

With the dictionary ensemble, we aim to find the index of the basis where a data point has its most sparse representation with the least error under one of the dictionaries, called the membership index. We use a greedy algorithm [2,3] in order to compute the membership indices segmentally for all data points in parallel on the GPU.

$$\{ \mathcal{S}^{(i)} \}_{i=1}^{N_i}, \mathbf{m}, \mathbf{z} = \text{Test}(\{ \mathcal{T}^{(i)} \}_{i=1}^{N_i}, \{ \mathbf{U}^{(1,k)}, \dots, \mathbf{U}^{(n,k)} \}_{k=1}^K, \epsilon, \tau_i)$$

We project data points of a testing set denoted as  $\mathcal{T}^{(i)}$  onto the dictionaries by parallelizing the tensor product on the GPU, see Equation 2.

$$\text{Tensor-product} \quad \mathcal{X}^{(i)} = \mathcal{T}^{(i)} \times_1 \mathbf{U}^{(1,k)} \dots \times_n \mathbf{U}^{(n,k)} \quad (2)$$

Finally, we find and store the sparse coefficients and their locations in order to reconstruct the data, see Figure 2.

## N-mode product on GPU

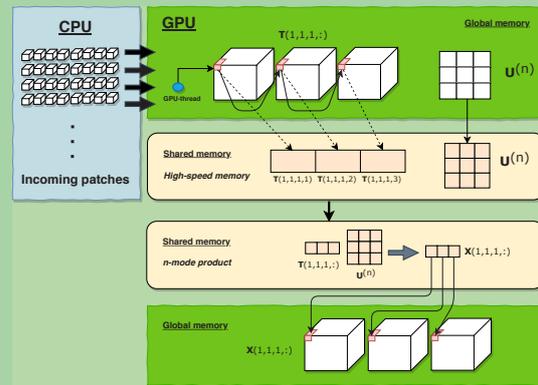


Figure 2: The computation of the tensor-product on the GPU.

We first load  $N$  number of patches, represented as tensors, on the GPU. We then let each thread traverse along the  $n$ -mode fiber and store the fiber on the fast on-chip memory called shared memory. The projection of a patch on a dictionary, see Equation 2, is computed by a vector-matrix product. This process can be simplified as a massive parallelized dot-product operation. We maximize thread collaboration and minimize bank conflicts by assigning a fiber to each thread. One of the biggest challenges with GPGPU computations is the overhead from memory transactions. We overcome this limitation by processing the patches segmentally in large batches.

## Results

We maintain the quality of the data from [1] and gain more than a tenfold speedup for the encoding process. The results for the GPU-implementation were achieved with a GTX Tian Xp. We compare our results against a CPU-implementation, see Fig. 3, where we used four Xeon E7-4870 with a total of 40 cores.

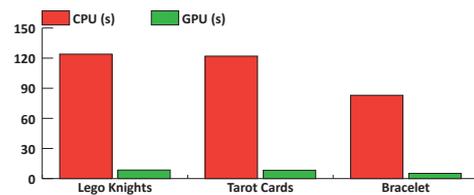


Figure 3: The measured time for the Lego Knights data set with  $K = 64$  dictionaries, sparsity  $\tau = 300$  and threshold  $\epsilon = 5 \times 10^{-5}$ . Further we have  $K = 64$ ,  $\tau = 390$  and  $\epsilon = 7 \times 10^{-5}$  for the Tarot Cards. Lastly, we have  $K = 64$ ,  $\tau = 412$  and  $\epsilon = 5 \times 10^{-5}$  for the Bracelet data set.

## References

- [1] Miandji, Ehsan and Hajisharif, Saghri and Unger, Jonas (2018). A Unified Framework for Compression and Compressed Sensing of Light Fields and Light Field Videos. ACM Transactions on Graphics. ACM, Accepted.
- [2] Miandji, E., Krenander, J., and Unger, J. (2015). Compressive image reconstruction in reduced union of subspaces. Comput. Graph. Forum, 34(2):33-44.
- [3] Miandji, E., Krenander, J., and Unger, J. (2013). Learning based compression of surface light fields for realtime rendering of global illumination scenes. In SIGGRAPH Asia 2013 Technical Briefs, SA '13, pages 24:1-24:4. New York, NY, USA, ACM.

Batkovic, Ivo  
Chalmers / Zenseact AB

## Using MPC to Enable Safe Autonomous Driving

The rapid development of autonomous driving technologies in the past decades has been driven by the objectives of enabling safer and more efficient transportation. However, in order to enable such automated systems to be deployed on a global scale, problems regarding safety must be addressed. In particular, a self-driving vehicle must be able to safely interact with a surrounding environment consisting of other road users, whose intentions cannot be perfectly known. In this poster, we briefly mention how Model Predictive Control (MPC) can be used to ensure safe autonomous driving in uncertain environments.

# Using MPC to Enable Safe Autonomous Driving

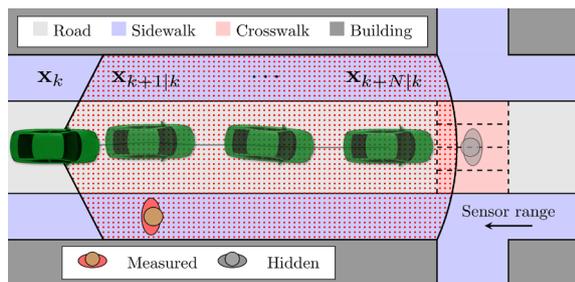
Ivo Batkovic, Ind. PhD, Zenseact AB and Chalmers University of Technology  
Dept. of Electrical Engineering, Mechatronics group  
Supervisors: Prof. Paolo Falcone (CTH) and Dr. Mohammad Ali (Zenseact AB)



## Motivation & Research Goals

In the past decade both the research community and industry have spent a vast amount of time and resources to further develop autonomous driving technologies with the objective of increasing safety and efficiency of passengers and goods transportation. However, in order to fully deploy highly automated driving functionalities, vehicles need not only to reliably sense their surrounding environment, but also *safely* interact with it. In order to overcome such problems one has to address the question of how to design a vehicle controller that is safe by design, but also what requirements need to be set on the sensor-suite and prediction algorithms in order to enable safe autonomous driving. This poster presents an MPC-based approach to ensure safe autonomous driving in uncertain environments by slightly modifying the standard MPC controller design.

## Methods



The objective is to control the autonomous vehicle, given by a nonlinear model  $\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k)$ , such that the a-priori *known* constraints  $h(\mathbf{x}, \mathbf{u}) \leq 0$  and a-priori *unknown* constraints  $g(\mathbf{x}, \mathbf{u}) \leq 0$  are satisfied. In this setting, function  $h$  can model actuator limitations or the allowed distance to the lane boundaries, and is known beforehand. The unknown constraint  $g$  on the other hand models the uncertainty and collision-avoidance w.r.t moving obstacles.

In order to ensure that the vehicle always plans a *safe* trajectory that satisfies the a-priori known and unknown constraints, we formulate the vehicle controller as the following *Model Predictive Control Problem*

$$V(\mathbf{x}) := \min_{\mathbf{x}, \mathbf{u}} \sum_{n=k}^{k+N-1} q(\mathbf{x}_{n|k} - \mathbf{r}_{n|k}^{\mathbf{x}}, \mathbf{u}_{n|k} - \mathbf{r}_{n|k}^{\mathbf{u}}) + p(\mathbf{x}_{k+N|k} - \mathbf{r}_{k+N|k}^{\mathbf{x}})$$

subject to

$$\begin{aligned} \mathbf{x}_{k|k} &= \mathbf{x}_k, \\ \mathbf{x}_{n+1|k} &= f(\mathbf{x}_{n|k}, \mathbf{u}_{n|k}), & \forall n \in [k, k+M-1] \\ h_n(\mathbf{x}_{n|k}, \mathbf{u}_{n|k}) &\leq 0, & \forall n \in [k, k+M-1] \\ g_{n|k}(\mathbf{x}_{n|k}, \mathbf{u}_{n|k}) &\leq 0, & \forall n \in [k, k+M-1] \\ \mathbf{x}_{n|k} &\in \mathcal{X}_r^s, & \forall n \in [k, k+M-1] \\ \mathbf{x}_{k+M|k} &\in \mathcal{X}_{\text{safe}}, \end{aligned}$$

where  $(\mathbf{r}^{\mathbf{x}}, \mathbf{r}^{\mathbf{u}})$  is a predefined reference,  $\mathcal{X}_r^s$  is a standard stabilizing set, and  $\mathcal{X}_{\text{safe}}$  is a safe set. By placing mild assumptions on the structure of the unknown constraint  $g_{n|k}$  and on the existence of a *safe* terminal set  $\mathcal{X}_{\text{safe}}$ , we show that recursive feasibility (safety) can be proven using a controller based on MPC [1,2].

## References

- [1] Safe Trajectory Tracking in Uncertain Environments  
Ivo Batkovic, Mohammad Ali, Paolo Falcone, and Mario Zanon  
Provisionally accepted to IEEE Transactions on Automatic Control
- [2] Experimental Validation of Safe MPC for Autonomous Driving in Uncertain Environments  
Ivo Batkovic, Paolo Falcone, and Mario Zanon  
To be submitted to IEEE Transactions on Control Systems Technology

## Selected Results

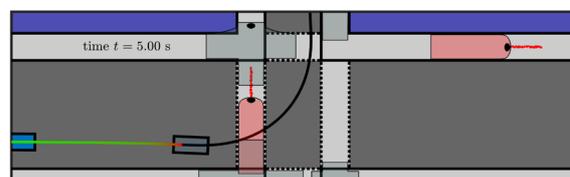
In order to apply the theory from [1], one must first design a safe set, which essentially is a robust invariant set where the a-priori unknown constraints must be inactive. In other words, we assume the following.

There exists a robust invariant safe set  $\mathcal{X}_{\text{safe}}$  such that for all  $\mathbf{x}_{n|k} \in \mathcal{X}_{\text{safe}}$ ,  $h(\mathbf{x}, \mathbf{u}) \leq 0$  and  $g(\mathbf{x}, \mathbf{u}) \leq 0$ , and there exists a *safe* control action that entails that  $f(\mathbf{x}_{n|k}, \mathbf{u}_{\text{safe}}) \in \mathcal{X}_{\text{safe}}$ .

We deploy the safe MPC framework in a real Volvo XC90 vehicle, and assume that a suitable safe set for urban autonomous driving situations is given by a vehicle that has come to a complete stop. The motivation behind this choice is that practical settings, where safety is emphasized, typically consider a system to be safe at steady-state, i.e., a vehicle that is parked in a safe configuration (e.g., a parking lot) is not responsible for collisions with other road users.



The framework was verified in a four-way intersection at a test track, where the vehicle had to track a predefined reference (black line below) while avoiding collisions with moving pedestrians. To do so successfully, the future pedestrian motion had to be predicted (seen by the red and gray boxes) so that the a-priori unknown constraints  $g$  could be formulated. The framework demonstrated real-time capabilities, while providing a comfortable driving behavior and avoiding collisions with the moving pedestrians.



Brissman, Emil  
Linköping University



## Representation Learning and Visual Localization

Exploring Representation Learning and Visual Localization within different research areas such as Video Object Segmentation, Video Instance Segmentation and Instance Segmentation.

# Representation Learning and Visual Localization

Emil Brissman<sup>1,2</sup>, supervisors Michael Felsberg<sup>2</sup> and Per-Erik Forssén<sup>2</sup>  
<sup>1</sup>Saab Dynamics, <sup>2</sup>Computer Vision Laboratory (CVL) Linköping University

## Background & Motivation

- Vehicle, autonomous or semi-autonomous.
- Need vision during operation.
- Onboard resources, limited computational power.
- Try to avoid data overload.
- Mimic human visual attention.
- Choosing salient areas from the total visual scene.

## Research Goal & Questions

- Develop visual methods (CNN) for eg. vehicle localization with single image input.
- What kind of image knowledge can be used?
- In order to estimate eg. risk of collision with other objects or ego-pose.
- Handle objects, varying in shape, with different representation for object detection

Detection

Tracking

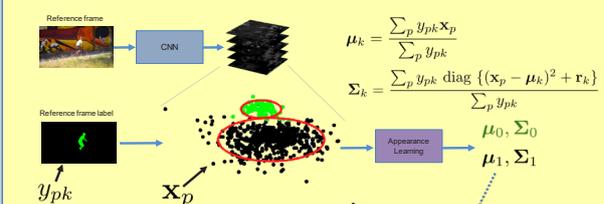
Segmentation

## Methods & Preliminary Results

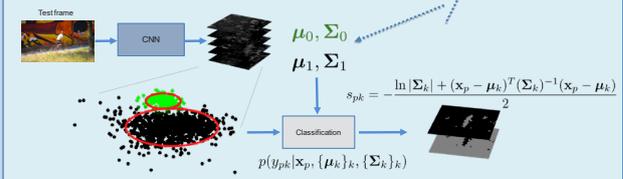
- **Video Object Segmentation** is the task of tracking and segmenting one or multiple objects.
  - **Input:** image sequence and segmentation mask in frame #1,  $t=1$
  - **Output:** segmentation masks for frames  $t > 1$
  - **Assumptions:** ground truth segmentation given in first frame (semi-supervised)
  - **Challenge:** build an object model from the first frame, appearance changes, distractor objects (semantically similar).
- Co-author on AGAME [1], presented at CVPR2019
  - The idea is to probabilistically model the feature distribution of the object and the background.
  - Feature vectors are explicitly classified via the probabilistic object model which is updated after each frame.

[1] J. Johnander, M. Danelljan, E. Brissman, F. S. Khan, and M. Felsberg. A generative appearance model for end-to-end video object segmentation. In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2019.

## Generative Appearance Model: Learning



## Generative Appearance Model: Classification



- In **Video Instance Segmentation** detection is added, no segmentation masks in frame #1.
  - **Input:** image sequence
  - **Output:** segmentation masks for each instance in each frames
  - **Challenge:** appearance changes, semantically similar objects, in-and-out of frame, occlusion
  - Paper [2] published at DAGM GCPR 2021, where it won an honorable mention (co-author)
- **Visual Instance Segmentation** is the task of detecting all objects in an image.
  - **Input:** a single image
  - **Output:** box, label and mask for each object, objects of the same class shall still be treated as separate entities
  - **Assumption:** Each object shape and object location approximately described by one or more anchor boxes
  - **Challenge:** Some objects are well described by a rectangle, however, this is hardly always the case
  - Paper [3] published at SAIS workshop 2021, exploring SDFs for Instance Segmentation
- **Keypoint detection** is the task of finding salient regions in two or more images
  - **Challenge:** model should learn uniqueness
  - **Application:** eg. pose estimation

[2] J. Johnander, E. Brissman, M. Danelljan, and M. Felsberg. "Video Instance Segmentation with Recurrent Graph Neural Networks." In: *Pattern Recognition: 43rd DAGM German Conference, DAGM GCPR 2021*  
 [3] E. Brissman, J. Johnander and M. Felsberg, "Predicting Signed Distance Functions for Visual Instance Segmentation," *2021 Swedish Artificial Intelligence Society Workshop (SAIS)*, 2021

## Towards Real-world Editable 3D Maps Using Deep Learning

Various mapping frameworks used in robotics allow to build dense 3D world representations. Most of these representations use monolithic data structures such as octrees, point clouds, or meshes. While such representations are well suited for navigation tasks, interaction with such representations is difficult. An important step towards more interactive maps is to infer the 3D shape and pose of objects in the scene from partial observations. With SDFest we propose an analysis-by-synthesis pipeline for joint pose and shape estimation using signed distance fields. The pipeline combines an initialization network, a generative shape model and a differentiable renderer to enable joint estimation of 7-DoF pose and shape from RGB-D images.

# Towards Real-world Editable 3D Maps Using Deep Learning

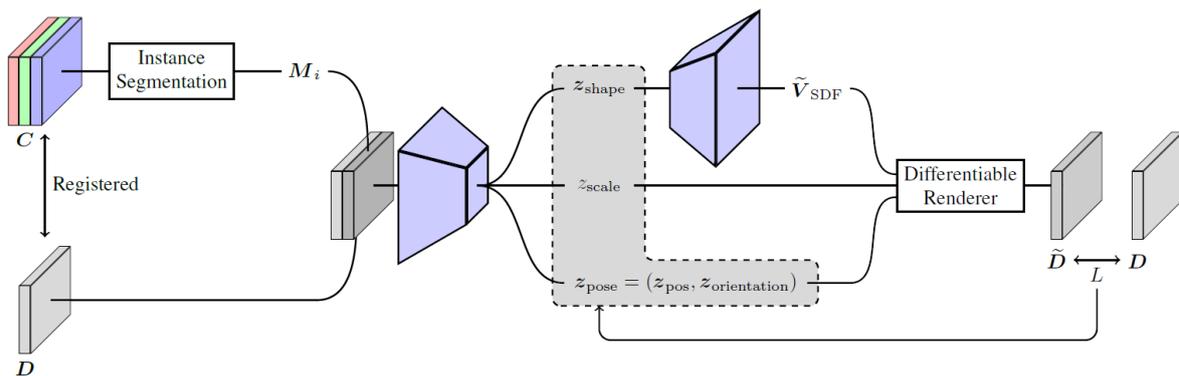
Leonard Bruns, KTH Royal Institute of Technology

Robotics, Perception and Learning



## Motivation & Research Goals

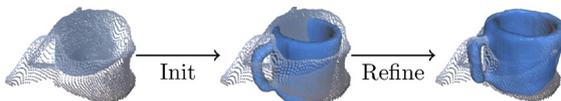
Various mapping frameworks used in robotics allow to build dense 3D world representations. Most of these representations use monolithic data structures such as octrees, point clouds, or meshes. While such representations are well suited for navigation tasks, interaction with such representations is difficult. An important step towards more interactive maps is to infer the 3D shape and pose of objects in the scene from partial observations. With **SDFest** we propose an analysis-by-synthesis pipeline for joint pose and shape estimation using signed distance fields. The pipeline combines an initialization network, a generative shape model and a differentiable renderer to enable joint estimation of 7-DoF pose and shape from RGB-D images.



## Introduction

We propose an analysis-by-synthesis pipeline for categorical shape and pose estimation. The pipeline consists of three main components: an initialization network, a generative shape model, and a differentiable renderer. The method currently works on a per-category level and only requires a collection of aligned meshes to be trained. No real-world data annotation is required.

## Method Overview



**Generative Shape Model** We train a variational autoencoder (VAE) to compress the shape. Shape is encoded as discretized signed distance fields (SDFs).

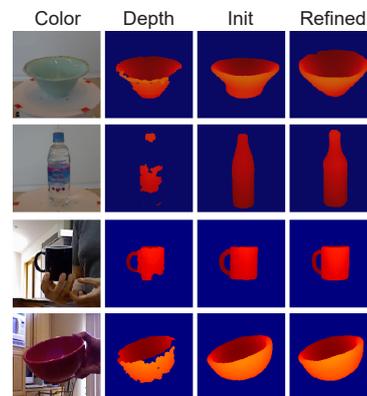
**Initialization Network** We train a neural network which regresses position, orientation, scale and the latent shape of the object from a partial observation. The network is trained in a supervised manner on synthetic data generated from the VAE.

**Differentiable Renderer** We use a differentiable renderer inspired by SDFDiff [1] to render a posed, discretized SDF, while obtaining gradients for the pose, scale, and SDF.

**Iterative Optimization** From the initial estimate, we start an iterative optimization procedure by decoding the signed distance field, rendering it in the current pose and formulating a loss comparing the current estimate with the measured depth.

## Results

We evaluate our approach on various synthetic and real-world datasets. Below we show qualitative results on the RGB-D Object dataset and the Redwood dataset, which includes handheld object in arbitrary orientations.



## References

[1] Y. Jiang, et al. SDFDiff: Differentiable rendering of signed distance fields for 3d shape optimization. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 1251-1261, 2020.

Ceylan, Ciwan  
KTH / SEB

## Feature Extraction from Transaction Graph

Banks are required to analyse large transaction datasets as a part of the fight against financial crime. Today, this analysis is either performed manually by domain experts or using expensive feature engineering. As part of my PhD, I investigate how vector representations can be learned in an unsupervised way from transaction data. I here present a published method for learning node features from transaction amounts, and concurrent work on extracting vector features from the graph structure.

# Feature Extraction from Transaction Graphs

Ciwan Ceylan, KTH Royal Institute of Technology

Division of Robotics, Perception and Learning, SEB

Main advisor: Danica Kragic

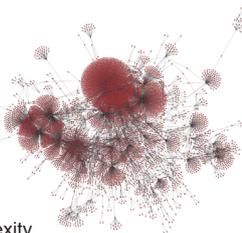


## Motivation & Research goals

Each year a vast amount of dirty money is laundered through the financial systems and financial institutions are under pressure to address this. The problem of discovering such schemes nevertheless remains very challenging from a machine learning perspective due to lack of labelled data, large data volumes and data secrecy. My research relates to advanced analytics for detecting financial crime occurring in transaction networks. Specifically, I'm interested in how recent advances in machine learning, network science and graph signal processing can be applied to financial transaction data in order to discover anomalous customer behaviour, which could indicate illegal activity. A crucial first step is to develop methods which can extract representative vector features from transaction graphs.

## Transaction Networks

Row	block_timestamp	from_address	to_address	amount
1	2018-06-01 07:06:17 UTC	0x8e5d083f	0x30727e88	150000.0
2	2018-06-01 07:12:35 UTC	0x30727e88	0x876eabf4	149999.9
3	2018-06-01 07:16:15 UTC	0x876eabf4	0x742d35cc	135130.0
4	2018-06-06 06:53:46 UTC	0x8e5d083f	0x052af704	259457.0
5	2018-06-08 06:33:24 UTC	0x53a46c10	0xf0881f9	188087.0
6	2018-06-12 16:42:34 UTC	0x052af704	0x35ab0d0f	150000.0



**Network structure** -- topological complexity

**Timestamps** -- dynamical complexity

**Transferred amount** -- data specific complexity

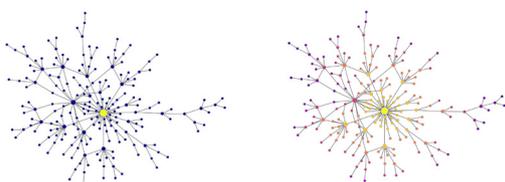
Given transactions observed in a time-window, a transaction network is a directed, weighted and possibly attributed graph, with an additional flow  $F$  on the edges:

$$G_T = (V, E, W, A, F)$$

For application of anomaly detection algorithms on graph nodes, a common approach is to represent the nodes as  $D$ -dimensional embedding vectors [1].

$$f(G_T) = X_T \quad X_T \in \mathbb{R}^{|V| \times D}$$

The mapping  $f$  should be found in an **unsupervised** way.



## Network Structure Embeddings

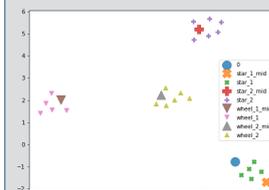
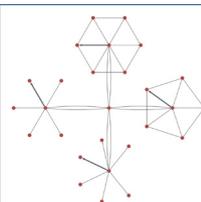
### Two types of node embeddings

In the literature, one finds two types of node embeddings.

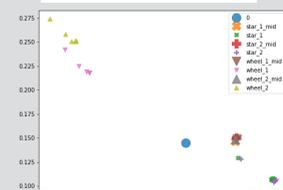
Adjacency embeddings are similar for nodes close in the network [2, 3].

Role embeddings are similar for nodes with locally similar topology [4, 5].

Most SotA role embeddings are only defined for undirected graphs.



NetMF [3]



GraphWave [4]

## Gated Gradient Model

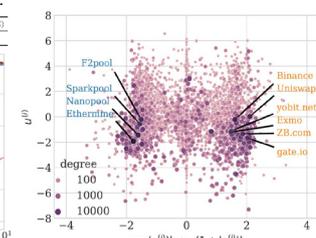
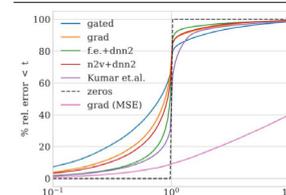
The network flow  $F$  can be captured as node feature vectors. This can be done using either feature engineering (f.e.) or a gradient model (grad) which learns scalar potentials for each node in the network.

The **gated gradient model** [6] extends the gradient model so that vector potentials can be learned instead. This is achieved by introducing a gate function:

$$f^{(ij)} = \sigma(\mathbf{u}^{(i)}, \mathbf{u}^{(j)}) \cdot (\mathbf{z}^{(j)} - \mathbf{z}^{(i)})$$

This model is evaluated by its flow prediction performance on a subgraph of the Ethereum transaction graph. The results generalize to bank internal transaction data.

GATED	GRAD	F.E.+DNN2	NV+DNN2	KUMAR ET AL.	SPRUS	GRAD (MSE)
0.34	10.39	13.10	11.25	14.19	14.61	390.63



## References

1. Akoglu, Leman, Hanghang Tong, and Danai Koutra. "Graph based anomaly detection and description: a survey." *Data mining and knowledge discovery* 29,3 (2015): 626-688.
2. Grover, Aditya, and Jure Leskovec. "node2vec: Scalable feature learning for networks." *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016.
3. Qiu, Jiezhong, et al. "Network embedding as matrix factorization: Unifying deepwalk, line, ple, and node2vec." *Proceedings of the eleventh ACM international conference on web search and data mining*, 2018.
4. Donnat, Claire, et al. "Learning structural node embeddings via diffusion wavelets." *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018.
5. Qiu, Jiezhong, et al. "Gcc: Graph contrastive coding for graph neural network pre-training." *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020.
6. Ceylan, Ciwan, Salla Franzén, and Florian T. Pokorny. "Learning Node Representations Using Stationary Flow Prediction on Large Payment and Cash Transaction Networks." *International Conference on Machine Learning*. PMLR, 2021.

## Decentralized Control of Dynamical Systems under Signal Temporal Logic Specifications

Autonomous systems need often to perform a variety of complex tasks at dynamic environments within certain time intervals. Examples of such tasks could be "reaching a known area within 5 sec" or "move with other agents in formation between 5 and 10 sec until the leader agent reaches a safety area". Each task may often evolve several agents that need to cooperatively design their future actions towards ensuring its satisfaction.

Nevertheless, when agents are working in large environments, communication among them might be difficult, costly or hard to establish. To that end, we propose a decentralized control framework that allows the satisfaction of a global formula with no need of communication. In our work we consider a set of complex tasks expressed as signal temporal logic formulas (STL), the satisfaction of which may depend on several or all agents in the team. As a first step, we decompose the global formula into local formulas whose satisfaction depends on given sub teams of agents using a convex optimization approach. Then, a receding horizon scheme (RHS) is proposed ensuring satisfaction of the local formulas and hence, satisfaction of the global formula. The proposed method is applied in a formation control example.



# Decentralized Control of Dynamical Systems under Signal Temporal Logic Specifications

Maria Charitidou, Dimos V. Dimarogonas, KTH  
Division of Decision and Control Systems, EECS

## Abstract

Autonomous systems need often to perform a variety of complex tasks at dynamic environments within certain time intervals. Examples of such tasks could be "reaching a known area within 5 sec" or "move with other agents in formation between 5 and 10 sec until the leader agent reaches a safety area". Each task may often evolve several agents that need to cooperatively design their future actions towards ensuring its satisfaction. Nevertheless, when agents are working in large environments, communication among them might be difficult, costly or hard to establish. To that end, we propose a decentralized control framework that allows the satisfaction of a global formula with no need of communication. In our work we consider a set of complex tasks expressed as signal temporal logic formulas (STL), the satisfaction of which may depend on several or all agents in the team. As a first step, we decompose the global formula into local formulas whose satisfaction depends on given sub-teams of agents using a convex optimization approach. Then, a receding horizon scheme (RHS) is proposed ensuring satisfaction of the local formulas and hence, satisfaction of the global formula. The proposed method is applied in a formation control example.

## 1. STL Decomposition

Signal Temporal logic is a specification language defined over continuous time signals. Let  $\mu \in \{\perp, \top\}$  be a predicate defined after the evaluation of a continuously differentiable predicate function  $h: \mathbb{R}^n \rightarrow \mathbb{R}$  as follows:

$$\mu = \begin{cases} \top, & h(x) \geq 0 \\ \perp, & h(x) < 0 \end{cases}$$

In our work, we consider a restricted STL fragment defined as follows:

$$\begin{aligned} \varphi &= G_{[a,b]} \mu \mid F_{[a,b]} \mu, \\ \phi &= \bigwedge_{i=1}^p \varphi_i, \end{aligned}$$

where  $\mu$  is a predicate and  $[a, b] \in \mathbb{R}_{\geq 0}$ . We consider a global formula  $\phi$  and a set of disjoint subteams of agents  $\mathcal{V}_l, l = 1, \dots, v$  with  $\bigcup_{l=1}^v \mathcal{V}_l = \mathcal{V}$ . Then, for every  $\varphi_i, i = 1, \dots, p$ , let  $V_i \subseteq \{1, \dots, v\}$  denote the indices of the subteams including at least one agent contributing to the satisfaction of  $\varphi_i$  and  $\bar{z}_i^l \in \bar{Z}_i^l$  are the states of  $\mathcal{V}_l$  satisfying  $h_i(x) = h(\bar{z}_i^1, \dots, \bar{z}_i^{v_i})$ , where  $l_1, \dots, l_{v_i} \in V_i$  and  $\bar{Z}_i^l$  a compact, convex, nonempty set. Then, the local formulas  $\bar{\varphi}_i^l, l \in V_i$  are defined as follows [1]:

$$\bar{\varphi}_i^l = \mathcal{T}_{[a_i^l, b_i^l]}(h_i^l(\bar{z}_i^l; \theta_i^l) \geq 0),$$

where

$$\begin{aligned} \mathcal{T} &= \begin{cases} F, & i \in I_F \\ G, & i \in I_G \end{cases} \\ [a_i^l, b_i^l] &= \begin{cases} [a_i, b_i], & i \in I_G \\ [t_i, t_i], & i \in I_F \end{cases} \end{aligned}$$

$$h_i^l(\bar{z}_i^l; \theta_i^l) = r_i^l - \|\bar{z}_i^l - c_i^l\|_{\infty},$$

$t_i \in [a_i, b_i]$  and  $\theta_i^l = (r_i^l, c_i^l) \in \mathbb{R}_{\geq 0} \times \bar{Z}_i^l$  are parameters found as the solution to the following optimization problem:

$$\begin{aligned} \max_{\theta_i^l} \sum_{i \in V_i} r_i^l \\ \bar{z}_i^l \in \{\xi \in \bar{Z}_i^l: \xi(\eta) = r_i^l + c_i^l(\eta) \text{ or } \xi(\eta) = -r_i^l + c_i^l(\eta)\}, \\ \bar{z}_i^l = [\bar{z}_i^1]_{i \in V_i} \in \text{int}\{x \in X: h_i(x) \geq 0\}, \\ \theta_i^l = (r_i^l, c_i^l) \in \mathbb{R}_{\geq 0} \times \bar{Z}_i^l, \end{aligned}$$

where  $\xi(\eta)$  is the  $\eta$ -th element of  $\xi$  and  $X$  is a known, compact, convex, nonempty set. Then, the local formula corresponding to  $\mathcal{V}_l$  is defined as:

$$\varphi^l = \bigwedge_{i \in \mathcal{V}_l} \bar{\varphi}_i^l.$$

## 2. Results

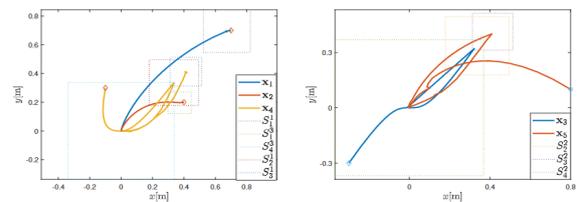
Given each local formula  $\varphi^l$ , we encode the STL constraints using the barrier function:

$$b(z_l, t) = -\ln\left(\sum_{i \in J_l} \exp(-b_i(z_l^i, t))\right),$$

where  $b_i(z_l^i, t) = -\gamma_i(t) + h_i^l(z_l^i; \theta_i^l)$  and  $\gamma_i: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$  are performance functions to be designed according to [2] ensuring satisfaction of all  $\bar{\varphi}_i^l, i \in J_l$  with a minimum robustness  $r$ . Then, satisfaction of  $\varphi^l$  is ensured, when  $b(z_l, t)$  remains nonnegative for every  $t \geq 0$ . Given the dynamical system:

$$\dot{z}_l = A z_l + B u_l,$$

we may define the trajectories of agents in  $\mathcal{V}_l$  using the receding horizon control scheme proposed in [3].



The proposed approach is implemented in a formation control problem of 5 agents. We consider 3 sub-teams  $\mathcal{V}_1 = \{1,4\}$ ,  $\mathcal{V}_2 = \{3,5\}$  and  $\mathcal{V}_3 = \{2\}$  and the STL formula  $\phi$  defined as:

$$\begin{aligned} \phi &= G_{[0,2]}(\|x_1 - x_2 - p_x\|^2 \leq 0.1) \wedge G_{[2.5,4]}(\|x_3 - x_4\|^2 \leq 0.2) \\ &\wedge F_{[3,7]}(\|x_5 - x_4\|_{P_1}^2 \leq 0.2) \wedge F_{[8,10]}\|x_5 - x_2\|_{P_2}^2 \leq 0.25). \end{aligned}$$

The STL formulas are decomposed to local formulas  $\varphi^l$ . The trajectories of the agents, shown in the Figure satisfy the local formulas. This ensures the satisfaction of  $\phi$  with minimum robustness 0.005.

## References

1. M. Charitidou, D. V. Dimarogonas, "Signal Temporal Logic Task Decomposition via Convex Optimization", IEEE Control Systems Letters, vol. 6, pp. 1238-1243, 2022.
2. L. Lindemann and D. V. Dimarogonas, "Decentralized Control Barrier Functions for Coupled Multi-Agent Systems under Signal Temporal Logic Tasks", European Control Conference, 2019, pp. 89-94.
3. M. Charitidou, D. V. Dimarogonas, "Barrier Function-based Model Predictive Control under Signal Temporal Logic Specifications", European Control Conference, 2021.

## Sybil-Based Attacks on Google Maps or How to Forge the Image of City Life

Participatory sensing (PS) applications collect all sorts of data by many users to maintain up-to-date data on everyday life, contributing to our well-being. Beyond occasional faults, it is often assumed that users are benign, thus strong security is not deployed. Controlling multiple users, an attacker can submit a large volume of forged data to dominate the PS-collected data. The result can be outright manipulation of the sensing process. We showcase the importance of this issue by selecting one of the most popular applications, Google Maps. The attacker model in our system is modest yet effective and efficient that are Sybil-based, leveraging non-existing, fake users. We instantiate automated attacks we term script- and emulator-based. The former submits crafted traffic in volume to manipulate the application data. The latter trades-off attack efficiency for increased versatility to attack other Google Maps features. We complete this investigation with human-based false data injection. This is the motivation of this work: to raise awareness on such a vulnerability and risk and improve the trustworthiness of such a popular application. We responsibly disclosed our findings to Google that acknowledged the issue and granted a reward.

Eryonucu, Cihan  
KTH



## Optimal Coordination of Mixed-Traffic Vehicles

In the context of mixed-traffic, the presence of Human-Driven Vehicles (HDVs) can pose several challenges to vehicles coordination due to uncertain, non-cooperative behaviors. In this work, we present an optimal control-based strategy for handling the HDVs by exploiting and coordinating Connected and Automated Vehicles (CAVs) forming the so-called mixed-platoons. A timeslot-based approach is used to schedule the vehicles that are going to occupy any intersection or roundabouts, with respect to safety requirements and physical limitations. In addition, we conduct a study and analyze the impact of human drivers' uncertainties in vehicles coordination.

# Optimal Coordination of Mixed-Traffic Vehicles

Muhammad Faris, Chalmers

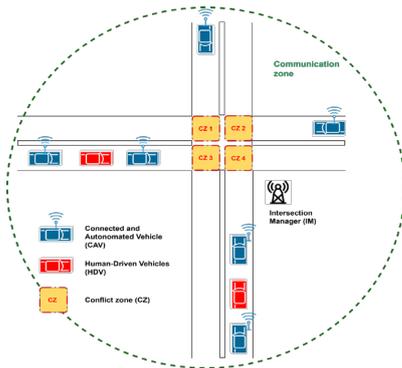
Division of Systems and Control  
Main supervisor: Paolo Falcone



## Abstract

In the context of mixed-traffic, the presence of Human-Driven Vehicles (HDVs) can pose several challenges to vehicles coordination due to **uncertain, non-cooperative behaviors**. In this work, we present an **optimal control-based strategy** for handling the HDVs by exploiting and coordinating Connected and Automated Vehicles (CAVs) forming the so-called **mixed-platoons**. A timeslot-based approach is used to schedule the vehicles that are going to occupy any intersection or roundabouts, with respect to safety requirements and physical limitations. In addition, we conduct a study and analyze the impact of human drivers' uncertainties in vehicles coordination.

## Mixed-Platoons



We consider the following

- **High-penetration** market rate of CAVs hence automated unsignalized intersections can be realized
- **Small** numbers of HDVs are squeezed in between **high** numbers of CAVs. CAVs are utilized as traffic **actuators** and **sensors** to regulate the HDVs.
- **Timeslot-based** intersection occupancy and safe distance approaches are applied to **avoid collisions** among vehicles with conflicting trajectories

## Control Methods

The vehicles coordination problem under consideration can be formally described in a constrained optimization problem

$$\min_{U,T} J(X)$$

$$\text{s.t. } x^+ = f(x, u),$$

$$X \in \mathcal{X}, U \in \mathcal{U}, T \in \mathcal{T}$$

... that can be recast as a receding horizon optimal control problem. Particularly, we propose a bi-level optimization based on **Model Predictive Control (MPC)** [1].

## References

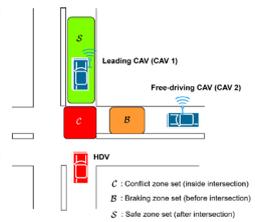
[1] J. B. Rawlings, D. Q. Mayne, and M. M. Moritz, *Model predictive control: Theory, Computation, and Design 2nd Edition*, vol. 197. 2019.  
[2] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*, vol. 1, no. 4. 2017.

## Reachability Analysis

We study the impact of uncertainties from HDVs on the vehicles coordination problem using **reachability tools** [2]

$$K_t(S, W) = \text{Pre}(K_{t-1}(S, W), W) \cap X$$

In particular, we evaluate the initial states feasible set of a free-driving CAV against a preceding mixed-platoon in terms of its ability to maintain safety, e.g., braking, under additive uncertainties of a human driver.



## Results

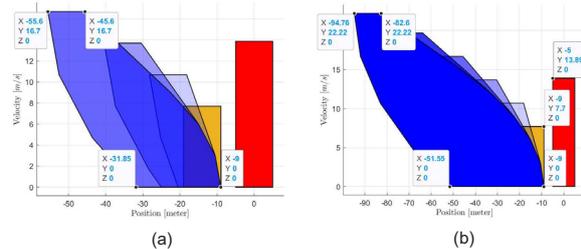
We perform numerical simulations of mixed-traffic vehicles coordination using two different HDV prediction models: (1) car-following and (2) constant measurements in a small-scale setting. The result is given in the following table

Model	Clearance time
(1) Car-following	2.9 sec
(2) Constant measurements	2.4 sec

Conclusions:

- The algorithm successfully prevents any collision or platoon cut-in.
- Using a model like in (1) might take a much longer time to clear an intersection due to misleading assumptions.

Furthermore, we carry out reachability analysis and present feasible sets of two different cases: (a) one with added uncertainty and (b) the other is not. The results are shown in the figures below



Conclusions:

Uncertainties shrink the feasible sets of initial states. The higher the uncertainty, the smaller the area.

Ferizbegovic, Mina  
KTH

## The fundamental lemma based on second order moments

In this paper, we propose variations of the fundamental lemma that utilize second-order moments such as correlation functions in the time domain and power spectra in the frequency domain. We believe that using a formulation with estimated correlation coefficients is suitable for data compression, and possibly can reduce noise.

Also, the formulations in the frequency domain can enable modeling of a system in a frequency region of interest.

# Willems' fundamental lemma based on second-order moments

Mina Ferizbegovic, Håkan Hjalmarsson, Per Mattsson, Thomas B. Schön



UPPSALA  
UNIVERSITET



## Summary and contributions

We propose:

- a variation of Willems' fundamental lemma based on the correlation functions, which can be useful for large noisy datasets,
- a variation of Willems' fundamental lemma using spectra, which can be useful for modeling in a frequency region of interest.

## Willems' fundamental lemma

Willems' fundamental lemma [1]: it is possible to describe all trajectories of a linear deterministic system from a single input-output data trajectory under some assumptions.

We assume:

- collected data  $(u^d, y^d)$  of length  $T$
- the system is controllable
- input  $u^d$  is persistently exciting of order  $L + n_x$ :

$$\begin{bmatrix} u_0^d & u_1^d & \dots & u_{T-L-n_x}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L+n_x-1}^d & u_{L+n_x}^d & \dots & u_{T-1}^d \end{bmatrix} \text{ full row rank}$$

For any trajectory  $(u, y)$  of length  $L$  of the system there exists  $g \in \mathbb{R}^M$

$$\begin{bmatrix} u_0 \\ \vdots \\ u_{L-1} \\ y_0 \\ \vdots \\ y_{L-1} \end{bmatrix} = \begin{bmatrix} u_0^d & u_1^d & \dots & u_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L-1}^d & u_L^d & \dots & u_{M-1}^d \\ y_0^d & y_1^d & \dots & y_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ y_{L-1}^d & y_L^d & \dots & y_{M-1}^d \end{bmatrix} g \quad M = T - L$$

## Data driven simulation

Given:  $u_{ini} = \begin{bmatrix} u_{-L_0} \\ \vdots \\ u_{-1} \end{bmatrix}$ ,  $y_{ini} = \begin{bmatrix} y_{-L_0} \\ \vdots \\ y_{-1} \end{bmatrix}$ ,  $u = \begin{bmatrix} u_0 \\ \vdots \\ u_{L-1} \end{bmatrix}$ , data  $(u^d, y^d)$

estimate  $x_0$                       predict output

Find:  $y = \text{col}(y_0 \dots y_{L-1})$

Classic: Find  $A, B, C, D$ . Using  $u_{ini}, y_{ini} \xrightarrow{A, B, C, D} x_0, u \xrightarrow{A, B, C, D, x_0} y$

Assumptions: the system is minimal,  $u^d$  is persistently exciting,  $\text{lag} \leq L_0$

$$\begin{bmatrix} U_p \\ Y_f \end{bmatrix} = \begin{bmatrix} u_0^d & u_1^d & \dots & u_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L_0-1}^d & u_{L_0}^d & \dots & u_{M+L_0-2}^d \\ u_{L_0}^d & u_{L_0+1}^d & \dots & u_{M+L_0-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L-1}^d & u_L^d & \dots & u_{M-1}^d \end{bmatrix} \begin{matrix} Y_p, Y_f \\ \text{similarly as} \\ U_p, U_f \end{matrix}$$

Solution[2]: Compute a solution of  $\begin{bmatrix} U_p \\ Y_f \end{bmatrix} g = \begin{bmatrix} u_{ini} \\ y_{ini} \\ u \end{bmatrix} \implies y = Y_f g$

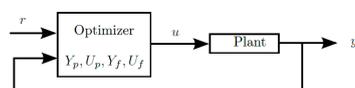
## Data driven predictive control

- A data-driven simulation used instead of a traditional model in the MPC algorithm.
- We obtain the following optimization problem:

$$\min_{g, u, y} \sum_{k=0}^{L-1} (y_k - r_{t+k})^T Q (y_k - r_{t+k}) + u_k^T R u_k + \lambda \times \text{regularization}$$

s.t.  $\begin{bmatrix} U_p \\ Y_p \\ U_f \\ Y_f \end{bmatrix} g = \begin{bmatrix} u_{ini} \\ y_{ini} \\ u \\ y \end{bmatrix}$ ,  $u_k \in \mathcal{U}$ ,  $y_k \in \mathcal{Y}$

offline noise  $\rightarrow$  s.t.  $U_p, Y_p$   
offline data  $\rightarrow$   $U_f, Y_f$   
online data  $\rightarrow$   $u, y$



- We optimize  $g$  online (receding horizon).
- If  $T$  is large,  $\dim(g) \approx T \implies$  high dimensional optimization problem. **Data compression!** [3]

## Willems' Lemma Using Correlation Functions

- autocorrelation  $R_{uu}(\tau)$  of the input  $u_t^d$
- cross-correlation  $R_{yu}(\tau)$  between output  $y_t^d$  and input  $u_t^d$  signals
- the system is controllable and stable
- $u^d$  is persistently exciting of order  $L + n_x$

$$\begin{bmatrix} R_{uu}(0) & R_{uu}(1) & \dots & R_{uu}(L-1) \\ R_{uu}(1) & R_{uu}(0) & \dots & R_{uu}(L-2) \\ \vdots & \vdots & \ddots & \vdots \\ R_{uu}(L-1) & R_{uu}(L-2) & \dots & R_{uu}(0) \end{bmatrix} > 0$$

For any trajectory  $(u, y)$  of length  $L$  of the system there exists  $g \in \mathbb{R}^M$

$$\begin{bmatrix} u_0 \\ \vdots \\ u_{L-1} \\ y_0 \\ \vdots \\ y_{L-1} \end{bmatrix} = \begin{bmatrix} R_{uu}(n_x) & R_{uu}(n_x-1) & \dots & R_{uu}(n_x+1-M) \\ R_{uu}(n_x+1) & R_{uu}(n_x) & \dots & R_{uu}(n_x+2-M) \\ \vdots & \vdots & \ddots & \vdots \\ R_{uu}(n_x+L-1) & R_{uu}(n_x+L-2) & \dots & R_{uu}(n_x+L-M) \\ R_{yu}(n_x) & R_{yu}(n_x-1) & \dots & R_{yu}(n_x+1-M) \\ R_{yu}(n_x+1) & R_{yu}(n_x) & \dots & R_{yu}(n_x+2-M) \\ \vdots & \vdots & \ddots & \vdots \\ R_{yu}(n_x+L-1) & R_{yu}(n_x+L-2) & \dots & R_{yu}(n_x+L-M) \end{bmatrix} g$$

## Comparison to the Willems' Lemma

- We have assumed that the input is a stationary process and that the system is stable.
- Data compression (reducing dimension of  $g$ )
- Ongoing work in a finite-time domain using the following approximation of correlation coefficients:  $R_{xy}(k) \approx \frac{1}{T-k} \sum_{i=0}^{T-k-1} x_{i+k} y_i^T$
- Using sample correlations reduce the effect of noise in the data compared to using raw data.

## Calculating correlation coefficients

Using that  $\begin{bmatrix} x_0 & x_1 & \dots & x_{M-1} \\ u_0^d & u_1^d & \dots & u_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L-1}^d & u_L^d & \dots & u_{T-1}^d \end{bmatrix}$  is full row rank, we can obtain

$$\begin{bmatrix} R_{uu}(k) \\ \vdots \\ R_{uu}(k+L-1) \\ R_{yu}(k) \\ \vdots \\ R_{yu}(k+L-1) \end{bmatrix} = \begin{bmatrix} u_0^d & u_1^d & \dots & u_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ u_{L-1}^d & u_L^d & \dots & u_{M-1}^d \\ y_0^d & y_1^d & \dots & y_{M-1}^d \\ \vdots & \vdots & \ddots & \vdots \\ y_{L-1}^d & y_L^d & \dots & y_{M-1}^d \end{bmatrix} g$$

## Willems' Lemma Using Spectra

- spectrum  $\Phi_{uu}(\omega)$  of the input  $u_t^d$ , at different  $M$  frequencies
- and the cross-spectrum  $\Phi_{yu}(\omega)$  between output  $y_t^d$  and input  $u_t^d$  signals, at different  $M$  frequencies
- the system is controllable and stable
- full row rank

$$[W_{L+n_x}(\omega_1) \otimes \Phi_{uu}(\omega_1) \quad \dots \quad W_{L+n_x}(\omega_M) \otimes \Phi_{uu}(\omega_M)],$$

where  $W_L(\omega) = [1 \ e^{j\omega} \dots \ e^{j\omega(L-1)}]^T$

For any trajectory  $(u, y)$  of length  $L$  of the system there exists  $g \in \mathbb{R}^M$

$$\begin{bmatrix} y_0 \\ \vdots \\ y_{L-1} \\ u_0 \\ \vdots \\ u_{L-1} \end{bmatrix} = \begin{bmatrix} W_L(\omega_1) \otimes \Phi_{yu}(\omega_1) & \dots & W_L(\omega_M) \otimes \Phi_{yu}(\omega_M) \\ W_L(\omega_1) \otimes \Phi_{uu}(\omega_1) & \dots & W_L(\omega_M) \otimes \Phi_{uu}(\omega_M) \end{bmatrix} g$$

## Future reading

- [1] J. C. Willems, P. Rapisarda, I. Markovsky, and B. De Moor. A note on persistency of excitation. *Systems & Control Letters*, 2005.
- [2] I. Markovsky and P. Rapisarda. Data-driven simulation and control. *International Journal of Control*, 2008.
- [3] J. Coulson, J. Lygeros, and F. Dörfler. Data-enabled predictive control: In the shallows of the DeePC. In *18th European Control Conference (ECC)*, 2019.
- [4] M. Ferizbegovic, H. Hjalmarsson, P. Mattsson, and T. B. Schön. Willems' fundamental lemma based on second-order moments. In *60th IEEE Conference on Decision and Control (CDC)*, 2021.

## Anomaly Detection and Countermeasures for Edge Clouds

The accelerated growth of the Internet of Things (IoT) and emerging 5G infrastructure has opened up opportunities to develop intelligent applications that transform data into business and societal value in plenty of application domains such as public services, intelligent transportation, augmented reality, industrial automation, and smart healthcare. The centralized cloud computing model has shown to have inherent problems when it comes to meet certain requirements of bandwidth-hungry or response-time-critical applications at the edge of the network. Thus, centralized clouds cannot provide services with high performance and reliability for such applications. Edge clouds are distributed computing infrastructures comprising edge nodes, fog nodes, and distant clouds, where the massive amount of data moves back-and-forth between the edge and distant cloud datacenters that concern data privacy and security issues. The ultimate goal of this project is to design, develop and deploy decentralized autonomous anomaly detection and countermeasures for ensuring performance and security in edge clouds using emerging machine learning models against unexpected service performance, security flaws and cyber attacks.

Forough, Javad  
Umeå University



# Anomaly Detection and Countermeasures for Edge Clouds

Javad Forough, Umeå University

Department of Computing Science  
Supervisor: Erik Elmroth  
Co-supervisor: Monowar Bhuyan



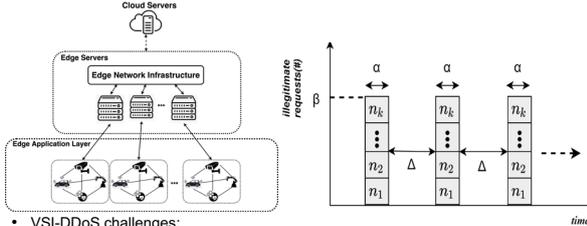
## Motivation

The accelerated growth of the Internet of Things (IoT) and emerging 5G infrastructure has opened up opportunities to develop intelligent applications that transform data into business and societal value in plenty of application domains such as public services, intelligent transportation, augmented reality, industrial automation, and smart healthcare. The centralized cloud computing model has shown to have inherent problems when it comes to meet certain requirements of bandwidth-hungry or response-time-critical applications at the edge of the network. Thus, centralized clouds cannot provide services with high performance and reliability for such applications. Edge clouds are distributed computing infrastructures comprising edge nodes, fog nodes, and distant clouds, where the massive amount of data moves back-and-forth between the edge and distant cloud datacenters that concern data privacy and security issues. The ultimate goal of this project is to design, develop and deploy decentralized autonomous anomaly detection and countermeasures for ensuring performance and security in edge clouds using emerging machine learning models against unexpected service performance, security flaws and cyber attacks.

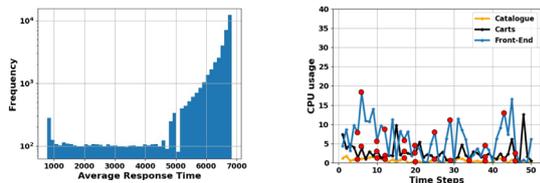


## Works Done

- Forough, J., Bhuyan, M., & Elmroth, E. (2021, August). Detection of VSI-DDoS Attacks on the Edge: A Sequential Modeling Approach. In *The 16th International Conference on Availability, Reliability and Security* (pp. 1-10). (With 20.34% acceptance rate)
- Forough, J., Bhuyan, M., & Elmroth, E. (2021, October). DELA: A Deep Ensemble Learning Approach for Cross-layer VSI-DDoS Detection on the Edge. Submitted to *The 37th ACM/SIGAPP Symposium On Applied Computing (SAC2022)*



- VSI-DDoS challenges:
  - Degradation of Quality of Service (QoS) for the legitimate users
  - Hard to detect, as the pattern of the monitored metrics remains very similar to the one in non-attack periods
  - Impact is even worse when attack happens on different layers of the system

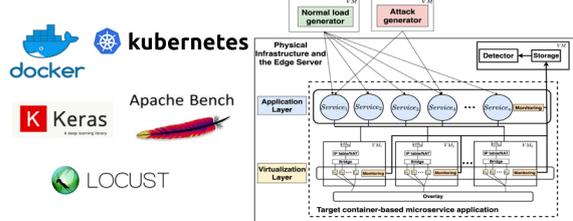


## References

- Shan, H., Wang, Q., & Yan, Q. (2017). Very short intermittent DDoS attacks in an unsaturated system. In *International Conference on Security and Privacy in Communication Systems* (pp. 45-66). Springer.
- Park, J., Nyang, D. and Mohaisen, A., (2018). Timing is almost everything: Realistic evaluation of the very short intermittent ddos attacks. In *2018 16th Annual Conference on Privacy, Security and Trust (PST)* (pp. 1-10). IEEE.
- Sherstinsky, A. (2020). Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network. *Physica D: Nonlinear Phenomena*, 404, 132306.
- Luong, M.T., Pham, H. and Manning, C.D., 2015. Effective approaches to attention-based neural machine translation. arXiv preprint arXiv:1508.04025.

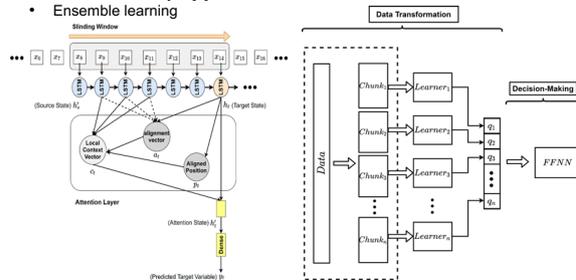
## Testbed Setup

- Container-based microservice application
- Docker and swarm for service deployment and orchestration
- Locust for normal load generation
- Apache Bench for VSI-DDoS implementation
- Keras for model implementation



## Proposed Methods

- Deep sequential modelling
- Long Short-Term Memory (LSTM) [3]
- Local attention layer [4]
- Ensemble learning



## Current Work

- The Quality of Service (QoS) of microservice applications experienced by users on the edge clouds is always exposed to security and performance anomalies such as slow-http, DDoS, spike overload, resource failure, etc. To address such problems of anomalies, we need to first identify the root cause of the problem as a "Security" or "Performance" anomaly, to be able to mitigate the issues in the next step. Hence, in this work, we plan to develop a comprehensive Anomaly Detection System (ADS) which will be able to detect both security and performance anomalies.

## Machine Learning Algorithms for Automatic Labelling

Obtaining labels for semi-supervised learning can be an extravagant and tedious task because of manual labeling. Because of this, industries are looking for automated solutions for data labeling. Semi-supervised machine learning algorithms (SSL) are used to automatically label datasets where few labels are available. However, it is time-consuming for practitioners in industry and academia to choose the optimal labeling algorithm for a particular problem. Therefore it is relevant to provide research that provides practitioners knowledge to choose the optimal algorithms for their specific use cases.

# Machine Learning Algorithms for Automatic Labelling



Teodor Fredriksson, Chalmers  
Computer Science and Engineering

## Motivation

- Supervised classification tasks requires labeled data.
- In industrial settings, datasets are rarely labeled.
- Data Labeling might be costly in terms of time and money.
- Automatic Labeling approaches exists but are not widely used in industry.
- Lack of research to help new practitioner's choose optimal AL approach based on their situation.

## Semi-Supervised Learning

- **Semi-supervised classification algorithms learns from both labeled and unlabeled data.**
- Assumes small amount of labeled data

## Future Research

In future research we wish to evaluate more state-of-the-art machine learning and deep learning algorithms for data labeling and evaluate them utilizing different datasets and settings based on industry.

## Previous Work

- **Systematic Literature Review and Mapping Study [1], [2].**
  - In what research fields can we apply active and semi-supervised learning
  - What kind of machine learning algorithms are used?
  - What is the popularity of datatypes among the different methods?
  - What are the datasets used to evaluate these algorithms?
  - What algorithm(s) should be used for each application?
- **Case Study with Industry [3].**
  - What are the key-challenges that practitioners face in the process of labeling data?
  - What are the mitigation strategies that practitioners use to overcome these challenges?
- **Empirical Evaluation of Graph-based Semi-Supervised Learning Algorithms .**
  - Evaluates 13 different SSL algorithms on 24 different datasets divided into three datatypes, (numerical, text, images).
  - What is the ranking of the algorithms in terms of highest accuracy w.r.t aggregated data, manual effort and datatype?
- **Assessing the Sustainability of Semi-Supervised Learning Datasets using Item Response Theory [4].**
  - What datasets are suitable to compare different graph-based SSL algorithms compared.
  - How can different graph-based SSL algorithms be compared?

## References

- [1] Fredriksson, T., Bosch, J. and Olsson, H.H., 2020. Machine Learning Models for Automatic Labeling: A Systematic Literature Review. In *ICSOFT* (pp. 552-561).
- [2] Fredriksson, T.A., Mattos, D.I., Bosch, J. and Olsson, H.H., 2020. Machine Learning Algorithms for Data Labeling: An Empirical Evaluation
- [3] Fredriksson, T., Mattos, D.I., Bosch, J. and Olsson, H.H., 2020, November. Data labeling: an empirical investigation into industrial challenges and mitigation strategies. In *International Conference on Product-Focused Software Process Improvement* (pp. 202-216). Springer, Cham.
- [4] Fredriksson, T., Mattos, D.I., Bosch, J. and Olsson, H.H., 2021, September. Assessing the Suitability of Semi-Supervised Learning Datasets using Item Response Theory. In *2021 47th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)* (pp. 326-333). IEEE

Gyllenhammar, Magnus  
Zenseact

## Considerations for safety assurance of ADSs

Safety assurance of Automated Driving Systems (ADS) is arguably one of the largest outstanding challenges before large-scale deployment of such systems on public roads. In my research I focus on the aspects of providing, not only effective, but also efficient safety assurance of ADSs. Central to safety assurance is the compilation of a compelling safety (assurance) case that presents evidence-supported arguments for the system's safety fulfilment. I have investigated different ways to approach safety assurance and break down this task by using the Operational Design Domain (ODD), but also by looking at different assurance methods to support the safety case construction and maintenance.

# Considerations for safety assurance of ADSs

Magnus Gyllenhammar, Zenseact, KTH

Mechatronics

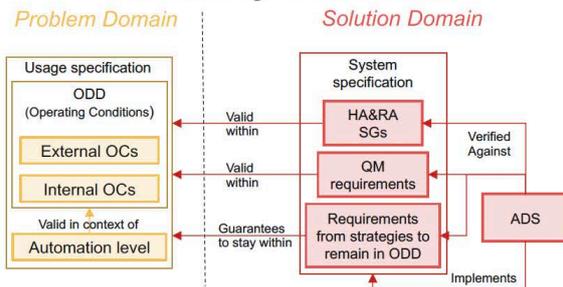


zenseact

## Abstract

Safety assurance of Automated Driving Systems (ADS) is arguably one of the largest outstanding challenges before large-scale deployment of such systems on public roads. In my research I focus on the aspects of providing, not only effective, but also efficient safety assurance of ADSs. Central to safety assurance is the compilation of a compelling safety (assurance) case that presents evidence-supported arguments for the system's safety fulfilment. I have investigated different ways to approach safety assurance and break down this task by using the Operational Design Domain (ODD) [1], but also by looking at different assurance methods to support the safety case construction and maintenance [2].

## Operational Design Domain (ODD) for safety assurance



The Operational Design Domain (ODD) of an ADS has been proposed to limit or constrain the development, design, verification and validation activities by limiting the operations of the ADS to this ODD. This thus splits the problem of safety assurance into two parts:

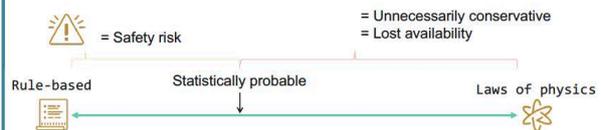
- Designing, developing and providing evidence for a safe ADS within the limits of the ODD, and
  - Ensuring that the ADS does not operate outside the ODD.
- For the second aspect the following strategies have been proposed [1]:
- I. Internal, inherent in ADS feature definition
  - II. External, checking mission when accepting strategic task
  - III. External, statistically defined spatial and temporal triggering conditions
  - IV. External, run-time measurable triggering conditions related to operating conditions

## References

1. Gyllenhammar, Magnus, et al. "Towards an operational design domain that supports the safety argumentation of an automated driving system." *10th European Congress on Embedded Real Time Systems (ERTS 2020)*. 2020.
2. Gyllenhammar, Magnus, Carl Bergenhem, and Fredrik Warg. "ADS Safety Assurance—Future Directions." *CARS 2021 6th International Workshop on Critical Automotive Applications: Robustness & Safety*. 2021.
3. Gyllenhammar, Magnus, et al. "Minimal Risk Condition for Safety Assurance of Automated Driving Systems." *CARS 2021 6th International Workshop on Critical Automotive Applications: Robustness & Safety*. 2021.

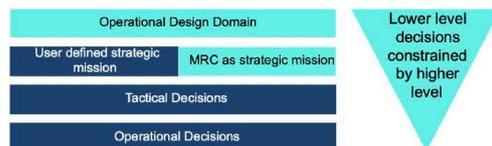
## Data and models for assurance

Throughout the development processes as well as the assurance there is need for using models. These models can either be based on conservative assumptions, such as using the laws of physics for determining possible actions of other traffic participants. On the other end of the spectrum, one could consider using predefined rules, with some presumptions about where and how they are applicable. However, both these extremes will result in a suboptimal system considering both safety and performance. A better approach is letting data support those models.



## Minimal Risk Condition for safety assurance

To fulfil the ODD, one should ensure that the strategies (I) – (IV) is determined for all types of ODD exits. However, the system itself could also experience deterioration or degradation to its capabilities. In such cases it might not be possible to safely continue fulfilling the mission. If this is the case the system should go into MRC. To contextualise that decision, one could use the following decision hierarchy [3]:



The key element of this is the split of the decisions on the strategic level. The mission of the ADS is either to fulfil the user defined mission or to go into MRC. This is in relation to the ODD at the top or in the case of degradations this ODD is exchanged with the Restricted Operational Domain (ROD).

Hellander, Anja  
Linköping University

## Unified task and motion planning

Many robotic applications involve both high-level (discrete) task planning and low-level (continuous) motion planning. Solving the two planning problems separately one after the other often leads to suboptimal solutions, or no feasible solution at all. This doctoral project aims at tightly integrating methods for task planning with methods for optimal-control-based motion planning in order to solve the task and motion planning problems simultaneously. This poster gives a brief background to the problem, presents the overarching research questions of the doctoral project and presents the work that is currently ongoing.

Hellander, Anja  
Linköping University

# Unified Task and Motion Planning

Anja Hellander, Linköping University  
Department of Electrical Engineering

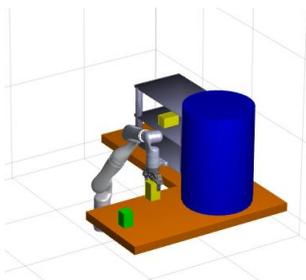


## Abstract

Many robotic applications involve both high-level (discrete) task planning and low-level (continuous) motion planning. Solving the two planning problems separately one after the other often leads to suboptimal solutions, or no feasible solution at all. This doctoral project aims at tightly integrating methods for task planning with methods for optimal-control-based motion planning in order to solve the task and motion planning problems simultaneously. This poster gives a brief background to the problem, presents the overarching research questions of the doctoral project and presents the work that is currently ongoing.

## Background

Task and motion planning are naturally interdependent in many robotic applications where robots have to perform both high-level task planning in order to achieve some objective as well as low-level motion planning in order to determine how to perform actual movements. Hierarchical approaches where the task and motion planning are performed separately often give suboptimal or even infeasible solutions. In order to achieve reliable task and motion planning it is therefore necessary to formulate and solve a single integrated planning problem where the task and motion planning problems are solved jointly.



Example of a task and motion planning problem. The manipulator robot must perform task planning in order to determine which objects to pick and place, and motion planning in order to determine how to perform the operations.

## Research questions

- How to extend an existing action language (e.g. PDDL) to include specifications of optimal control problems for dynamic systems?
- How to tightly integrate methods behind classical AI planners with methods behind motion planners using graph search and numerical optimal control?
- How to develop efficient heuristics for problems that tightly integrate task and motion planning?

## Ongoing work

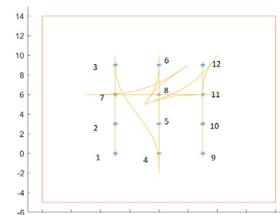
The current focus is developing a framework for unified task and motion planning that rather than only finding a feasible solution (if one exists) to a task and motion planning can also perform optimization of this solution (at least to some degree).

The ongoing work is therefore focused on:

- How to handle that the motion planning problem has continuous variables whereas the task planning problem is discrete. How should the discrete values be generated? In advance, during the search? Randomly or deterministically?
- The search will require calls to some function in order to determine if feasible motion plans exist or not, which will be expensive. How can the number of calls be reduced?
- How can optimization be integrated into already existing frameworks for task and motion planning?

## Ongoing work: Drill planning

- **Setting:** A drill rig must drill a number of holes at given positions. Once a hole has been drilled, the rig cannot pass over it. The drill holes are positioned densely relative the drill rig size.
- **Problem:** Decide in which order the holes are to be drilled, and plan feasible paths between them for the rig to follow.
- **Our approach:** Discretize the configuration space of the drill rig. Graph search (backward) in a state space with state consisting of current position (hole), current (discretized) heading and previously drilled holes. Call to lattice-based motion planner to determine if a feasible path between two states exists.



Example of a resulting path for the center of the vehicle's rear axle.

## An Adaptive Approach for Task-Driven BCI Calibration

Brain-Computer Interfaces (BCI) use brain signals as inputs and machine learning algorithms to decipher the meaning of these. A BCI system needs to be calibrated before usage, i.e., the machine learning algorithm needs to be trained. The overall goal is to solve a task as fast as possible. The calibration can be terminated with an adaptive approach when the BCI system is good enough to solve the task. Here we present a structure for such a system and show some initial results.



# An Adaptive Approach for Task-Driven BCI Calibration



FRIDA HESKEBECK AND CAROLINA BERGELING – DEPARTMENT OF AUTOMATIC CONTROL

## Background

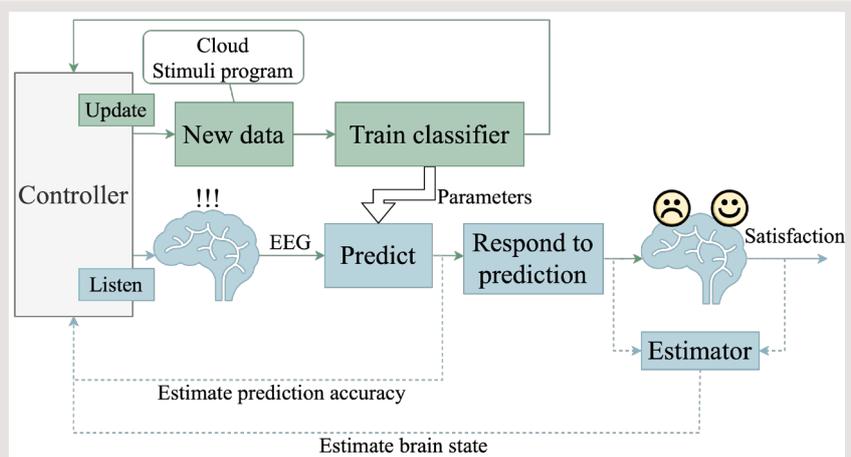
**Brain-Computer Interfaces (BCI)** use brain signals as inputs and machine learning algorithms to decipher the meaning of these. A BCI system needs to be calibrated before usage, i.e., the machine learning algorithm needs to be trained. The calibration phase is often tedious for the user. Hence we focus on improving the calibration of BCI systems.

## Recent work

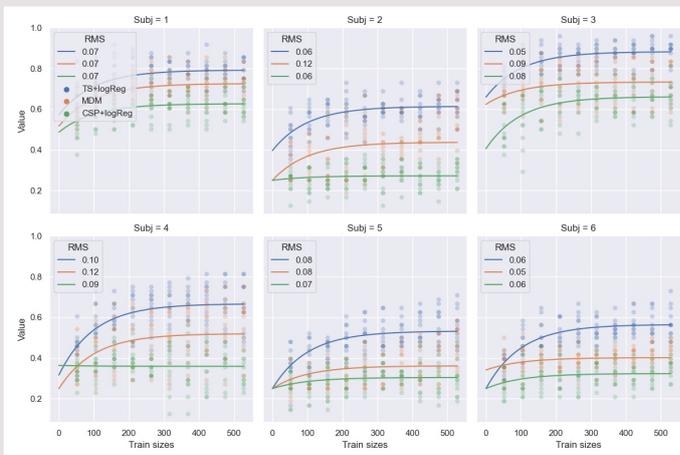
- ❖ Aim: Fast task solving
  - Hearing aid - Attenuate sound source
- ❖ Controller decides action
  - *Update* – retrain machine learning algorithm with new data
  - *Listen* – predict user intent.

## Conclusion

Our suggestion for an adaptive Brain-Computer Interface **automatically decides** whether to 1) **listen** to the brain and respond accordingly, using the current machine-learning algorithm to predict the meaning of the brain signals, or 2) **update** the machine learning algorithm to make better predictions in the future.



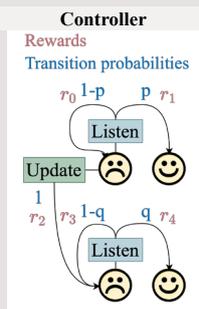
- ❖ Model learning rates (shown in the graphs).
  - Estimate  $p$  and  $q$ .
  - Model: First order system, fixed  $T$ , individual start and end points.
  - Across sessions and individuals.



- ❖ Markov Decision Process (MDP) to find best action

- If the inequality below is true, *update* is best action.
- Parameters: **Transition probabilities, rewards, and discount factor.**

$$\frac{(1-p)r_0 + pr_1}{1-\gamma(1-p)} < r_2 + \gamma \frac{(1-q)r_3 + qr_4}{1-\gamma(1-q)}$$



## Future

- Resources for *update* action:
  - Time to do experiment.
  - Computer power to retrain machine learning model.
  - Money for accessing the cloud with previous data.
- Controller.
  - Extend Markov Decision Process.
  - Use some "error" based on brain state.
- Estimate prediction accuracy.
  - Without validation data?
- Estimate user's brain state.
  - Use error potentials – a type of brain response.
  - Estimation with sporadic input from user.

## Contact

Frida Heskebeck | [frida.heskebeck@control.lth.se](mailto:frida.heskebeck@control.lth.se) | heskebeck.com

Hynén Ulfsjöö, Carl  
Linköping University

## Motion-planning and decision-making under uncertainty for heavy vehicles.

To safely maneuver a heavy vehicle in complex traffic situations, the uncertainty in the prediction of the surrounding vehicles must be considered during planning. In this poster a two-stage approach to this problem is presented that tightly couples a POMDP with scenario-based stochastic MPC, to be able to exploit the properties of both methods.

This is applied to a highway driving situation where the ego vehicle wants to overtake a vehicle in dense traffic, where the prediction of the environment is uncertain and there is uncertainty in how cooperative each driver is. The resulting two-stage motion planner is able to safely plan in this situation and the inclusion of the MPC-step is shown to drastically improve the solution from just using the POMDP.

# Motion-planning and decision-making under uncertainty for heavy vehicles.

Carl Hynén Ulfsjöö, Linköping University  
Dept. of Automatic Control, ISY  
Supervisor: Prof. Daniel Axehill (LiU)



## Introduction

To safely and efficiently maneuver a heavy vehicle in a complex traffic situation, the driver needs to perceive, interpret and predict the motion of multiple surrounding vehicles. Then based on that prediction it must choose an appropriate action that considers the large level of uncertainty in the prediction, without becoming overly conservative.

To realize this in a motion planner it should:

- take the uncertainty in prediction into account
- exploit interactions between the ego and surrounding vehicles
- make joint discrete and continuous decisions.

## Method

The developed motion planner is based on a two-stage approach. First a general partially observable Markov decision process (POMDP) is solved, then the solution is used in the second stochastic model predictive control (MPC) step, which improves the solution. This results in a motion planner where the POMDP:

- makes discrete decisions
  - handles general uncertainty in perception and prediction
  - outputs a coarsely discretized control signal
- and the MPC:
- improves the solution locally
  - handles unimodal Gaussian uncertainty
  - outputs a finely discretized control signal.

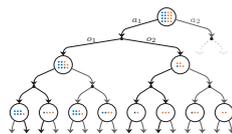
## Partially observable Markov decision process

The POMDP models a decision process where the noisy system dynamics are known but the underlying state cannot directly be measured. It tries to find the optimal policy ( $\pi$ ) that maps a probability distribution over the state-space to an action. The optimization problem that it solves can be written as:

$$\begin{aligned} & \underset{\pi(\cdot)}{\text{maximize}} && \mathbb{E} \left[ \sum_{k=0}^N \gamma^k R(x_k, u_k) \right] \\ & \text{subject to} && x_{k+1} = f(x_k, u_k) + w_t(x_k, u_k) && \text{(prediction)} \\ & && y_k = h(x_k, u_k) + v_k(x_k, u_k) && \text{(observation)} \\ & && b_k \sim p(b_k | b_{k-1}, y_k, u_k) && \text{(belief propagation)} \\ & && u_k \sim \pi(b_k), x_0 \sim b_0 \\ & && x \in \mathcal{X}, u \in \mathcal{U}, y \in \mathcal{O}. \end{aligned}$$

The POMDP is solved using the on-line POMDP solver DESPOT that uses sampling to approximate the uncertainty, which converts the problem to a tree-search problem.

The POMDP solver can handle very general uncertainty and directly consider partial observability. However the resulting tree scales poorly with regards to  $|\mathcal{U}|$ , which in practice means that the control signal must be coarsely discretized.



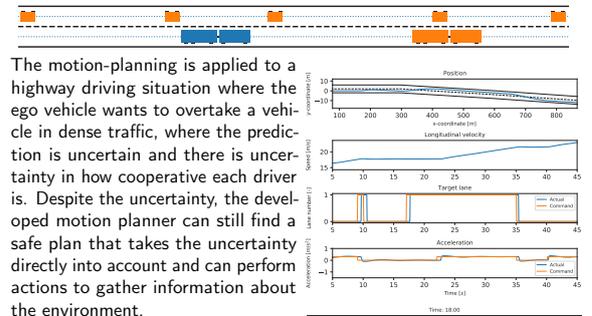
## Stochastic model predictive control

The stochastic MPC step is introduced to compensate for the coarse discretization in the solution to the POMDP. A scenario-based stochastic MPC formulation is used, because of the multimodal nature of the prediction of the environment (a surrounding vehicle might or might not yield). This uses discrete scenarios to represent the different modes, and for each mode typical stochastic MPC techniques are used to represent the local uncertainty.

As the MPC step is based on the solution to the POMDP, the solution can be used in several ways to tighten the coupling between them.

- The scenarios in the MPC can be derived from the sampled scenarios in the POMDP, and scenario reduction techniques can be used to only include the most relevant scenarios.
- The solution to the POMDP can directly be used to warm start the optimization solver.
- The POMDP solution can be used to define nonanticipatory constraints in the MPC, which determine when different modes are indistinguishable.

## Preliminary results



The motion-planning is applied to a highway driving situation where the ego vehicle wants to overtake a vehicle in dense traffic, where the prediction is uncertain and there is uncertainty in how cooperative each driver is. Despite the uncertainty, the developed motion planner can still find a safe plan that takes the uncertainty directly into account and can perform actions to gather information about the environment.

An example of this can be seen in the first figure to the right where the planner first commands a lane change at time 10s as it believes that vehicle in the passing lane is cooperative enough. However, as the vehicle does not react cooperatively it postpones the lane change until after the vehicle has passed. In the second figure to the right the result of using the two-stage approach is shown. The blue vehicle is only using the POMDP solution while the green is using the improved solution and is therefore able to perform the lane change much faster.

## Conclusions and future work

Combining a POMDP with stochastic MPC makes it possible to exploit the best properties of both methods. The method shows promising results in experiments on a typical highway driving situation. As future work the coupling between the two methods needs to be further investigated, additionally the implementation must be improved to make it real-time capable.

## Combining verbal-HRI with Behavior Trees to disambiguate human demonstrations

Fast changing tasks in unpredictable, collaborative environments are typical for medium-small companies, where robotised applications are increasing. Thus, robot programs should be generated in short time with small effort, and the robot able to react dynamically to the environment. To address this, a method exists that combines context awareness and planning to learn Behavior Trees (BTs), a reactive policy representation that is becoming more popular in robotics. The method allows to learn BTs from human demonstration. In those tasks in which the robot is required to fetch items for subsequent manipulation tasks, ambiguities might originate from the presence of identical objects in the scene. To disambiguate the scene, we propose a method that exploits visual data and uses verbal-HRI to request the human intervention, asking questions to understand the target item for the task. We combine this method to the existing BT learning framework to endow the robot with the capability of solving the task in ambiguous scenarios.

## Combining verbal Human Robot Interaction to solve ambiguities in Behavior Tree execution

Matteo Iovino, Irmak Doğan, Christian Smith, Iolanda Leite, KTH  
Robotics Perception and Learning

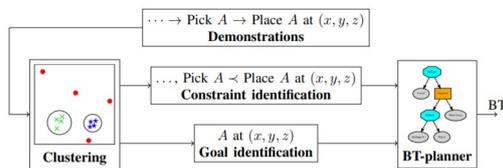


### Motivation & Research goals

Fast changing tasks in unpredictable, collaborative environments are typical for medium-small companies, where robotised applications are increasing. Thus, robot programs should be generated in short time with small effort, and the robot able to react dynamically to the environment. To address this, a method exists that combines context awareness and planning to learn Behavior Trees (BTs) from demonstration [1]. However, situations may arise where the robot is tasked to fetch an item that is present in multiple copies. The robot faces an ambiguous scenario that has to be disambiguated for the task to continue. We propose to combine the existing LfD method for BTs with verbal-HRI that uses visual data to query the scene for the target object and asks questions to the human to disambiguate it [2].

### Learn BTs from demonstration

At a high level, our proposed algorithm learns BTs from demonstrations in four steps. Human demonstrations are clustered to infer the context of each action and similarities between them, and then to infer task constraints and goal conditions, which are finally fed to a planner that builds the BT.



#### Demonstrations

The teaching method is kinesthetic and there are three actions available: a Pick action will close the robot grippers around the target object and a Place or Drop action will open the grippers, releasing the object. For all actions, the pose of the end-effector is recorded as the target pose for that action.

#### Behavior Tree Synthesis

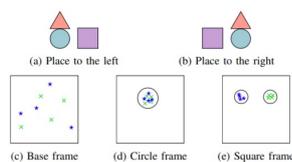
The BT is synthesised using the planner proposed in [3], leveraging the idea of backchaining. We run the plan offline because it is preferable to have the full tree available before running it on a real robot.

#### Goal and Constraints identification

The algorithm infers task constraints by observing the order in which actions appear in the demonstrations and adding each pair of ordered actions to the list of constraint and translated into preconditions that must be fulfilled before executing an action. Conflicting constraints are removed.

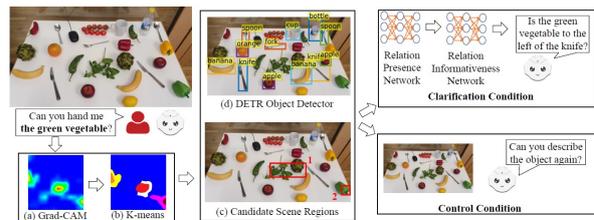
#### Clustering of demonstrated actions

Different actions might also be executed in different reference frames. Thus, equivalent actions across demonstrations have to be identified and their reference frame inferred. If an action possibly belongs to multiple clusters, we can infer the context in which the action is performed.



### Resolve ambiguities with verbal-HRI

When a human queries the robot to identify an object, situations may arise where the target object is present in multiple copies in the scene. Using RGB data from the camera, the robot uses Grad-CAM [3] to compute the activation regions corresponding to the query and then K-means to cluster them. Then, the DETR Object Detector [4] is used to detect the objects in the scene together with their bounding boxes. The bounding boxes are compared against the clusters to output candidate scene regions. Deep Neural Network techniques are used to parse the natural language sentence to find the target object and to formulate clarification conditions, using other identified items in the scene and referring spatial expression (*left of item\_x*). A conversation is then started to finally disambiguate the object.



### Combined method

We propose to combine the two methods to disambiguate the task during the execution of a BT learned from demonstration. We assume that the scene is not ambiguous during the demonstration and hence the robot is able to successfully grasp the target object for the task. If the task is ambiguous at execution the BT will fail and the disambiguation pipeline is triggered.

### References

- Gustavsson et al. (2021). Combining Context Awareness and Planning to Learn Behavior Trees from Demonstration. *arXiv e-prints*, arXiv-2109.
- Dogan et al. (2021). Asking follow-up clarifications to resolve ambiguities in human-robot conversation. *Preprint*.
- Selvaraju et al. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision* (pp. 618-626).
- Kamath et al. (2021). Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 1780-1790).
- Colledanchise et al. (2019, May). Towards blended reactive planning and acting using behavior trees. In *2019 International Conference on Robotics and Automation (ICRA)* (pp. 8839-8845). IEEE.

Jakobsson, Erik  
Linköping University / Epiroc Rock Drills AB

## Condition Monitoring for Hydraulic Rockdrills

In this work we investigate data driven methods for classifying patterns in pressure measurements from hydraulic rock drills. By using signatures from induced faults, we aim to handle different configurations and individuals without the need for obtaining training data from all possible configurations. The key is to generate features that capture the difference from a fault in relation to a non-faulty reference. These relative features should also be insensitive to differences from different configurations.

Jakobsson, Erik  
Linköping University / Epiroc Rock Drills AB



# Condition Monitoring for Hydraulic Rockdrills

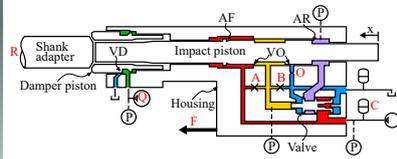
Erik Jakobsson  
Erik Frisk, Robert Pettersson, Mattias Kryssander  
erik.jakobsson@epiroc.com



## DESCRIPTION

In this project we aim to monitor the internal condition of hydraulic rock drills. This is done using measurements during operation, combined with machine learning schemes to classify different faulty behavior. An important aspect is the high variability between different applications/configurations. The use of Non-Fault reference data is a key technique.

## BACKGROUND & MOTIVATION



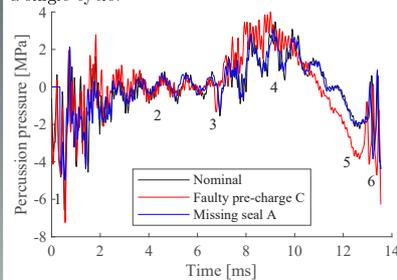
The basic functionality of a rock drill is simple. An impact piston is connected via hydraulic channels to a valve. The interaction between the two components results in a self oscillating mecha-

nism, where impacts between the impact piston and shank adapter generate stress waves used to drill holes in hard rock. Side effects of the oscillations are severe vibrations and pressure pulsations, making the rock drill a difficult application to monitor. The hope is to use a very low number of sensors, possible even positioned away from the machine.

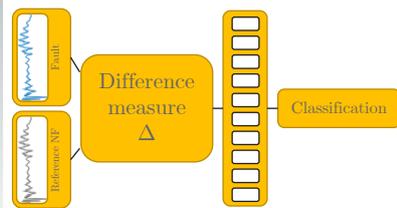
Knowing the current condition of the rock drill is an important step towards autonomous mining, where the information will be used for maintenance planning, logistics and prevention of secondary damage.

## CONTRIBUTIONS

**Pressure Signature Fault Detection** A single pressure sensor is used to classify the current condition of a rock drill in a lab setting. An example of such pressure data is shown below, for a single cycle.



The faulty conditions are seen to slightly differ from the nominal No-fault case. However, configuration changes such as different hose lengths can have similar effects and need to be accounted for. We do this by looking at the difference from the No-Fault case, from the same configuration.



We wish to find a function  $\Delta$  such that the same fault  $A$  in different configurations  $i, j$  give the same output,

$$\Delta(NF_i, A_i) = \Delta(NF_j, A_j) \quad \forall i \neq j \quad (1)$$

but different faults  $A, B$  give different outputs

$$\Delta(NF_i, A_i) \neq \Delta(NF_i, B_i) \quad \forall A \neq B \quad (2)$$

**Dynamic Time Warping Feature Vector** So far, the best  $\Delta$  found consists of a number of difference measures based on Dynamic Time Warping (DTW), as a way to handle differences in stroke duration and pressure wave propagation.

$$\Delta = \begin{bmatrix} p(\delta_{amp}(A^1, NF^1), \delta_{amp}(A^R, NF^R)) \\ \delta_{fs}(A^1, NF^1) \\ \dots \\ \delta_{osc}(A^1, NF^1) \\ p(\delta_{lag}(A^1, NF^1), \delta_{lag}(B^R, NF^R)) \end{bmatrix}$$

A feature vector is generated for each sample by measuring pairwise difference  $p(\delta_i, \delta_{ref})$  and frequency differences to a set of reference samples using various measures  $\delta$ . SVM classification using such feature vectors give the following accuracy:

True Class \ Predicted Class	no A	sm	worn co	Large do	low damp	Low PR	nom	cut valve	low pack	no B	thick steel
no A	50.4%	36.4%	0.4%	0.8%						12.0%	
sm	4.8%	91.2%								3.6%	0.4%
worn co	2.8%	0.4%	20.8%	15.2%	24.0%	4.4%	18.8%	13.6%			
Large do			1.2%	50.0%	25.6%	0.4%	0.4%	22.0%			0.8%
low damp	0.8%	0.8%	10.4%	26.8%	26.0%	0.4%	1.2%	33.2%	0.4%		
Low PR	3.6%	8.8%				1.2%	95.6%	22.0%	14.0%		
nom	0.8%	1.2%	2.4%	2.4%	3.2%	33.6%	99.2%	7.6%		0.4%	
cut valve	10.0%	0.4%	3.6%	0.8%	4.0%	2.4%	1.2%	71.2%		0.4%	
low pack	2.0%	1.2%		2.0%					84.8%		
no B	0.4%	0.4%	0.6%	11.2%	12.0%	0.4%	0.8%	18.4%		50.8%	
thick steel											100.0%

## RESEARCH GOAL & QUESTION

A number of research questions define the area:

1. How can a rock drill be modeled and monitored in order to predict future failure? What data should be collected to maximize the information gathered without creating a too complex product?



2. Can a low number of non-dedicated sensors be used to monitor multiple components, for example different parts of a drilling system?

3. How can No-Fault data be used to give a reference to handle differences between configurations?

An important aspect of the research is to understand how condition monitoring methods can be applied for products with a relatively low volume, high customization, and in a very harsh environment.

## PUBLICATIONS

- [1] Jakobsson et al. "Data driven modeling and estimation of accumulated damage in mining vehicles using on-board sensors" published in *Proceedings of Annual Conference of the Prognostics and Health Management Society, St. Petersburg, Florida, USA, 2017*.
- [2] Jakobsson et al. "Fatigue Damage Monitoring and Prognostics for Mining Vehicles using Data Driven Models" published in the *International Journal of Prognostics and Health Management (IJPHM), 2019*.
- [3] Jakobsson et al. "Automated Usage Characterization of Mining Vehicles For Life Time Prediction" published in *Proceedings of IFAC World Congress, Berlin, 2020*.
- [4] Åstrand et al. "A System for Underground Road Condition Monitoring" published in *International Journal of Mining Science and Technology, 2019*.
- [5] Jakobsson et al. "Fault Identification in Hydraulic Rock Drills from Indirect Measurement During Operation" published in *Proceedings of IFAC MMM, Nancy, 2021*.

## ROADMAP & MILESTONES

- Three conference papers accepted [1],[3],[5], two journal papers [2],[4] accepted.
- Licentiate thesis presentation, December 2019.
- Ongoing (Final): Condition monitoring of hydraulic rock drills journal paper.

Jensen, Maarten  
Umeå University

## Contextual Deliberation

Determining the context can help AI systems and agents in interacting with or resembling humans. However proper context dependent reasoning systems do not exist yet. Our aim is to create a framework that can make human like context dependent decisions. The poster gives a simple example that shows how we as humans intuitively use context determination to make decisions. Then it follows up with an initial conceptual framework that is a start for contextual deliberation in machines.



# WHAT IS YOUR CONTEXT?

## Contextual Deliberation

This is **you** at this poster presentation

### Questions about context

#### Where are **you** now?

In front of the poster, at home (zoom) at a conference, in a city, in Sweden, in Europe, on planet Earth

#### What time is it?

Exact? 15:00, afternoon, Wednesday, 12 January, winter, 2022, modern age, age of mankind (not dinosaur age)

#### What is the activity?

Poster presentations, working, doing science, networking, socializing

#### Who are around **you**? (Relations)

PhD students, professors, colleagues, friends, opposition, aliens, unknown

For humans context determination is intuitive, but how should machines do it?

### Why did **you** move to this poster? (Different Ways of Thinking)

- I moved sequentially from poster to poster (**Repetition**)
- There were more people at this poster than others (**Imitation**)
- I saw the title of the poster from a distance and it seemed most related to my research (**Rational choice**)
- If I'm at another person's poster I can lure people to my poster (**Game theory**)
- I met the author of the poster before and promised to see his poster (**Game theory or (Moral) values**)

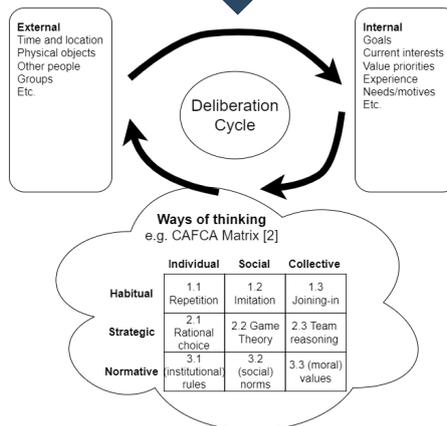
Maybe like this?

### Our project

Determining the context can help **AI systems** and **agents** in interacting with or resembling humans. However proper context dependent reasoning systems do not exist yet. Our aim is to create a framework that can make human like context dependent decisions [1].

### The deliberation cycle

The deliberation cycle (see figure) should explore *context* and reasons using different *ways of thinking* [5] until an action is found. Rather than first determining context and secondly determine which way of thinking to use. It is an iterative process where context influences the ways of thinking **and** the ways of thinking influences the exploration of the decision context.



### Context

This can be both consider from *internal* as well as *external* sources, see [3].

### Different Ways of Thinking [5]

Kahnemann proposed the idea that humans use two different ways of thinking to make decisions, Fast and Slow [4]. We extend this idea with more ways of thinking. The work of [2] proposes a framework that we use as a tool to have a relative complete categorization of different types of reasoning (see 3x3 matrix in figure).

### Learning

Not considered at the moment, but in the future **unsupervised learning** could be an interesting technique for context exploration/determination.

### Questions we will be able to answer

Insight in contextual deliberation can help us with questions like:

- How do people determine the context?
- How do people react to a new law?
- How do people react in a lockdown?
- How do we program a realistic human-interaction system or robot?

- [To be published] Jensen, M. Verhagen H., Vanhée L., & Dignum F. (2021) Towards Efficient Context-Sensitive Deliberation
- Elsenbroich, C., Verhagen, H.: The simplicity of complex agents: a contextual action framework for computational agents. *Mind & Society*15(1), 131-143 (2016)
- Zimmermann, A., Lorenz, A., & Oppermann, R. (2007, August). An operational definition of context. Springer, Berlin, Heidelberg.
- Daniel, K. (2017). Thinking, fast and slow?
- Minsky, M. (2007). The emotion machine, Simon and Schuster.

Johansson, Simon  
Chalmers / AstraZeneca

## Decision Making for Design of Chemical Libraries

The need for data in a standardized format grows stronger within the machine learning modeling for chemistry in the pharmaceutical area. One of the popular formats for data generation are chemical libraries, which can now rapidly be designed by generative models such as RNNs. In my project I propose a method for filtering the output of the focused generative models ( $\sim 10^5$ ) down to the typical size used for library design in a lab ( $\sim 10^2$ ). This process combines the fields of generative modeling, retrosynthesis prediction, chemical property prediction and decision-making to filter the compound selection.

Johansson, Simon  
Chalmers / AstraZeneca

# Decision Making for Design of Chemical Libraries

Simon Johansson

Supervisors: Alexander Schliep, Morteza Chehrehghani, Ola Engkvist  
University of Gothenburg|Chalmers University of Technology|AstraZeneca



Department of Computer science and Engineering|MolecularAI



UNIVERSITY OF GOTHENBURG

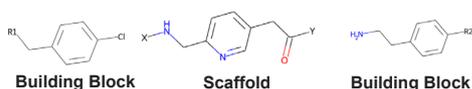
## Abstract

The need for data in a standardized format grows stronger within the machine learning modeling for chemistry in the pharmaceutical area. One of the popular formats for data generation are chemical libraries, which can now rapidly be designed by generative models such as RNNs. In my project I propose a method for filtering the output of the focused generative models ( $\sim 10^5$ ) down to the typical size used for library design in a lab ( $\sim 10^2$ ). This process combines the fields of generative modeling, retrosynthesis prediction, chemical property prediction and decision-making to filter the compound selection.

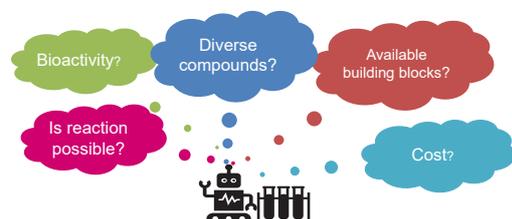
## Introduction

The development of strong data-driven models for chemistry in the pharmaceutical area has led to needs for more standardized data [1]. A **chemical library** is a collection of molecules synthesized under the same conditions with variations on the functional groups to represent a dense area in the chemical space.

Through **generative models** such as LibINVENT [2], thousands of molecules for libraries can be designed in an instant from a given core *scaffold*. By attaching **building blocks (BBs)** to this scaffold, we ensure a core similarity in the library.



This generative throughput is larger than the number of molecules that can be synthesized physically and a data-driven **system for compound selection** is needed to filter the list of suggestions to a manageable library. This can be done using numerous selection criteria.



## References

1. *AI-Assisted Synthesis Prediction*, Johansson, S et al., *Drug Discovery Today: Technologies* 32, 65-72. (2019)
2. *LibINVENT: Reaction-based Generative Scaffold Decoration for In Silico Library Design*, Fialkova, V et al., *JCIM*, (2021) <https://doi.org/10.1021/acs.jcim.1c00469>
3. *AiZynthFinder: a fast, robust and flexible open-source software for retrosynthetic planning*, Genheden, S. et al, *JChemInf*, (2020), <https://doi.org/10.1186/s13321-020-00472-1>
4. *VennABERs predictors*, Vovk, V & Petej, I. (2014) [arXiv:1211.0025](https://arxiv.org/abs/1211.0025)
5. *The coincidence approach to stochastic point processes*. Macchi, O. (1975) *Advances in Applied Probability*, 7(1), 83-122. doi:10.2307/1425855

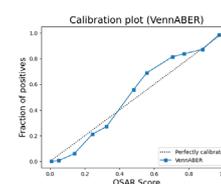
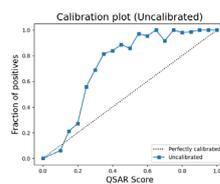
## Methods

LibINVENT is trained using reinforcement learning for 1k epochs. Fragments using the amide coupling and Buchwald-Hartwig reactions to connect to the scaffold were targeted as the focus.

The saved BBs are then evaluated through Monte Carlo Tree Search (MCTS) [3] to explore possible synthesis routes. We compare the routes against a list of available stock.

A random forest model was trained to predict the activity of the generated compounds towards DRD2, and calibrated using a dataset split of 60:20:20 for training:calibration:test with the **VennABERs** predictor [4].

The selection method intended to be used is **Determinantal point processes** [5].



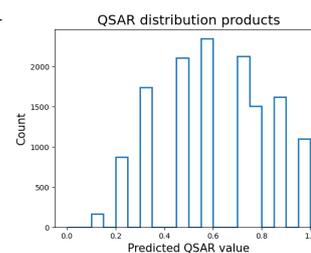
## Results

After collecting all sampled suggestions, a total of 42,448 molecules had been generated. This went through a filtering process:

- BBs used in less than 5 molecules.
- BBs which could not perform the targeted 2 reactions.
- BBs that could not be acquired within one synthesis reaction.

This yielded 100 BBs for amide coupling and 435 for Buchwald-Hartwig.

However, several fragments yield the same products together with the scaffold. The number of unique products yielded were 13600, with a skew towards being active.



The decision-making process for this selection is still in progress.



## Dynamic Visual Learning

This PhD-project studies two aspects of neural networks for computer vision:

Learning to process dynamic visuals, i.e., videos. Autonomous robots act in a dynamic world and need to form a high-level understanding of that world. One of the keys to obtaining such an understanding is the detection and tracking of dynamic objects.

Dynamically learning visual concepts. Neural networks are remarkably adept at image recognition. When trained, they are also very specific. If a new object category is to be recognized, hundreds of examples need to be annotated and the neural network retrained. It might be useful if the neural network could be shown a single (or handful) examples and directly be able to recognize them.

# Dynamic Visual Learning

Joakim Johnander, Linköping University  
Department of Electrical Engineering

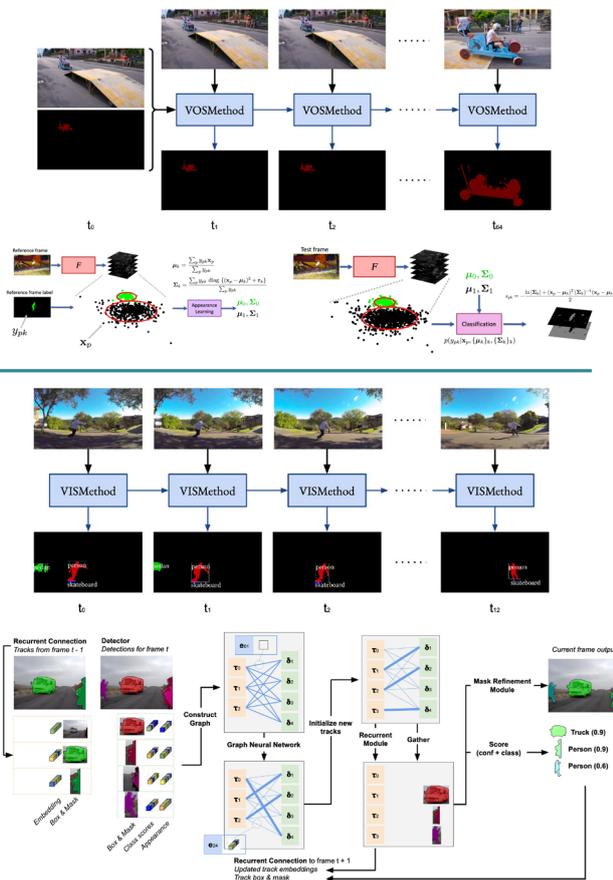


## Aims

- This PhD-project studies two aspects of neural networks for computer vision:
- Learning to process dynamic visuals, i.e., videos. Autonomous robots act in a dynamic world and need to form a high-level understanding of that world. One of the keys to obtaining such an understanding is the detection and tracking of dynamic objects.
  - Dynamically learning visual concepts. Neural networks are remarkably adept at image recognition. When trained, they are also very specific. If a new object category is to be recognized, hundreds of examples need to be annotated and the neural network retrained. It might be useful if the neural network could be shown a single (or handful) examples and directly be able to recognize them.

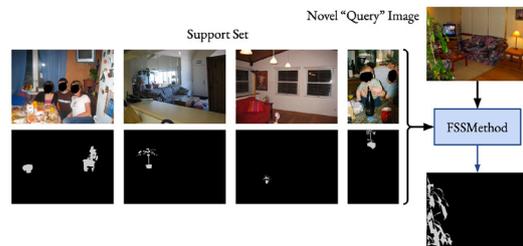
## Tracking Dynamic Objects

This PhD-project studies direct learning of the Video Object Segmentation (VOS) task – tracking and segmentation of generic objects – and Video Instance Segmentation (VIS) – detection, tracking, and segmentation of dynamic objects. One of the key challenges lies in the neural network design. Another challenge lies in the offline training of said neural networks, which is especially challenging in the case of Video Instance Segmentation.



## Few-Shot Segmentation

This PhD-project studies direct learning of the Few-Shot Segmentation (FSS) task – semantic segmentation given one or a handful of training examples. This problem is reminiscent of VOS, but the appearance variation within a single semantic class is far greater than the variations within a single instance.



Prior works study *prototype*-based learning mechanisms, but these struggle to model multi-modal appearance distributions. We instead explore a learning mechanism based on Gaussian Process regression.

$$\begin{aligned}
 \mathbf{x}_{Sk} &= F(I_{Sk}) \in \mathbb{R}^{D \times H \times W} \\
 \mathbf{y}_{Sk} &= G(M_{Sk}) \in \mathbb{R}^{H \times W} \\
 \theta &= \Lambda(\{\mathbf{x}_{Sk}, \mathbf{y}_{Sk}\}_k) = \frac{\sum_{k,h,w} (\mathbf{y}_{Sk})_{h,w} (\mathbf{x}_{Sk})_{h,w}}{\sum_{k,h,w} (\mathbf{y}_{Sk})_{h,w}} \\
 \mathbf{x}_Q &= F(I_Q) \in \mathbb{R}^{D \times H \times W} \\
 (\mathbf{y}_Q)_{h,w} &= f_\theta(\mathbf{x}_Q)_{h,w} = \frac{\langle (\mathbf{x}_Q)_{h,w}, \theta \rangle}{\|(\mathbf{x}_Q)_{h,w}\|_2 \|\theta\|_2} \\
 M_Q &= U(\mathbf{y}_Q, \mathbf{x}_Q, \text{shallow})
 \end{aligned}$$

- Assume  $\begin{pmatrix} \mathbf{y}_Q \\ \mathbf{y}_Q \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_S \\ \mu_Q \end{pmatrix}, \begin{pmatrix} \mathbf{K}_{SS} & \mathbf{K}_{SQ} \\ \mathbf{K}_{SQ} & \mathbf{K}_{QQ} \end{pmatrix}\right)$
- We can predict  $\mathbf{y}_Q | \mathbf{y}_S, \mathbf{x}_S, \mathbf{x}_Q \sim \mathcal{N}(\mu_Q | \mathbf{S}, \Sigma_Q | \mathbf{S})$  where  $\mu_Q | \mathbf{S} = \mathbf{K}_{SQ}^T (\mathbf{K}_{SS} + \sigma_s^2 \mathbf{I})^{-1} \mathbf{y}_S$  and  $\Sigma_Q | \mathbf{S} = \mathbf{K}_{QQ} - \mathbf{K}_{SQ} (\mathbf{K}_{SS} + \sigma_s^2 \mathbf{I})^{-1} \mathbf{K}_{SQ}$ .

## References

- Joakim Johnander, Martin Danelljan, Fahad Shahbaz Khan, and Michael Felsberg. "DCCO: Towards deformable continuous convolution operators for visual tracking." In: *International Conference on Computer Analysis of Images and Patterns*. Springer, 2017, pp. 55–67.
- Goulan Bhat, Joakim Johnander, Martin Danelljan, Fahad Shahbaz Khan, and Michael Felsberg. "Unveiling the power of deep tracking." In: *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 483–498.
- Joakim Johnander, Goulan Bhat, Martin Danelljan, Fahad Shahbaz Khan, and Michael Felsberg. "On the Optimization of Advanced DCF-Trackers." In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, 2018.
- Matej Kristan, Ales Leonardis, Jiri Matas, Michael Felsberg, Roman Pfugfelder, Luka Cehovin Zajc, Tomas Vojir, Goutam Bhat, Alan Lukezic, Abdelrahman Eldehakey, et al. "The sixth visual object tracking v2018 challenge results." In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, 2018.
- Joakim Johnander, Martin Danelljan, Emil Brissman, Fahad Shahbaz Khan, and Michael Felsberg. "A generative appearance model for end-to-end video object segmentation." In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8953–8962.
- Amanda Berg, Joakim Johnander, Flavie Durand de Gevigney, Jorgen Ahberg, and Michael Felsberg. "Semi-automatic annotation of objects in visual-thermal video." In: *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019.
- Joakim Johnander, Emil Brissman, Martin Danelljan, and Michael Felsberg. "Video Instance Segmentation with Recurrent Graph Neural Networks." In: *Pattern Recognition: 43rd DAGM German Conference, DAGM GCPR 2021, Bonn, Germany, September 28 – October 1, 2021, Proceedings*. Ed. by G. Bauckhage et al. (eds). Springer, 2021. doi: 10.1007/978-3-030-92659-5\_13.
- Emil Brissman, Joakim Johnander, and Michael Felsberg. "Predicting Signed Distance Functions for Visual Instance Segmentation." In: *2021 Swedish Artificial Intelligence Society Workshop (SAIS)*. IEEE, 2021, pp. 1–6.
- Joakim Johnander, Johan Edstedt, Michael Felsberg, Fahad Shahbaz Khan, and Martin Danelljan. "Dense Gaussian Processes for Few-Shot Segmentation." In: *arXiv preprint arXiv:210.03674* (2021).

## Learning to Segment Images Without Mask Labels

Deep learning methods have achieved remarkable results in many computer vision tasks, including semantic segmentation, where the task is to classify each pixel in an image to a predefined set of classes, e.g. person, cat or car. Applications include autonomous driving, video surveillance, and medical image analysis. However, training deep segmentation models requires large datasets of costly human-annotated pixel-wise segmentation masks. In this work, we explore a branch called weakly-supervised semantic segmentation, where the only source of supervision are cheap image-level classification labels. We propose two contributions; importance sampling, and feature similarity loss, for approaching this challenging task, and significantly improve contour accuracy over state-of-the-art methods.

# Learning to Segment Images Without Mask Labels

Arvi Jonnarth, Linköping University &amp; Husqvarna

Department of Electrical Engineering, Computer Vision Laboratory  
Supervisors: Michael Felsberg (LiU), Adam Tengblad (Husqvarna)

## Abstract

Deep learning methods have achieved remarkable results in many computer vision tasks, including semantic segmentation, where the task is to classify each pixel in an image to a predefined set of classes, e.g. person, cat or car. Applications include autonomous driving, video surveillance, and medical image analysis. However, training deep segmentation models requires large datasets of costly human-annotated pixel-wise segmentation masks. In this work, we explore a branch called weakly-supervised semantic segmentation, where the only source of supervision are cheap image-level classification labels. We propose two contributions; *importance sampling*, and *feature similarity loss*, for approaching this challenging task, and significantly improve contour accuracy over state-of-the-art methods.

## Methods

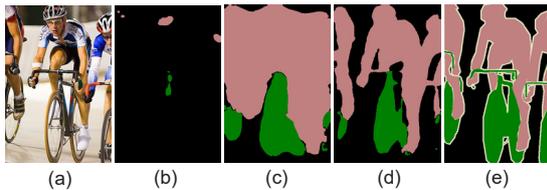


Figure 1. CAM comparison. (a) Input image; pseudo-masks with (b) max pooling, (c) importance sampling, and (d) importance sampling and feature similarity loss; (e) ground truth.

A fully convolutional neural network is trained in three stages:

1. Training of a multi-label classification network to generate class activation maps (CAMs). Max or average pooling is used to go from pixel-wise to image-level predictions.
2. Training of an AffinityNet [2] to predict pixel affinities.
3. A final segmentation network is supervised by pseudo-masks generated by the CAM and AffinityNet networks.

## Contributions

Classification networks are known to (1) mainly focus on discriminative regions, and (2) to produce diffuse CAMs without well-defined prediction contours. We approach both problems with two contributions for improving CAM learning in stage 1.

First, we use **importance sampling** based on  $K$  probability mass functions  $p_k$ , one per class  $k \in \{1, \dots, K\}$ , induced by the CAMs  $a_\theta \in [0, 1]^{W \times H \times K}$  to sample image-level predictions  $\tilde{y}_k$ .

$$p_k(I, J|x) = \Pr(I = i, J = j|x, k) = \frac{1}{Z_k(a)} a_\theta(x)_{ijk},$$

$$\tilde{y}_k = a_\theta(x)_{ijk}, \quad (i, j) \sim p_k(I, J|x).$$

Second, we formulate a **feature similarity loss** term  $\mathcal{L}_{fs}$  which aims to match the prediction contours with edges in the image.

$$\mathcal{L}_{fs}(a, x) = -\frac{1}{(HW)^2} \sum_{ij} w_{ij} g(a_i, a_j) f(\delta(x_i, x_j)),$$

Gaussian spatial weight  $\frac{1}{2\pi\sigma^2} \exp\left(-\frac{\|p_i - p_j\|_2^2}{2\sigma^2}\right)$       gating function  $\frac{1}{2} \|a_i - a_j\|_2$       pixel dissimilarity function  $\tanh\left(\mu + \log\left(\frac{\delta}{1-\delta}\right)\right)$ ,  $\delta = \|x_i - x_j\|_1 / C$

## Selected Results

The model is evaluated on the VOC benchmark dataset, with 20 foreground classes. Qualitative results are shown in Figure 2.

In Table 1 we compare our method with state-of-the-art weakly supervised methods in terms of two complementary metrics:

1. Mean intersection over union (mIoU) based on the area of predicted segmentation masks.
2. F-score based on the contours of segmentation predictions.

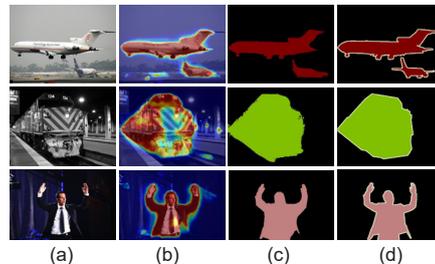


Figure 2. Qualitative results. (a) Input image, (b) foreground class activations of the CAM network, (c) segmentation predictions of the final model from stage 3, and (d) ground truth.

Table 1. Performance comparison on the VOC validation set.

Method	Area mIoU	Contour F-score
SEAM [1]	64.5	35.7
PMM [3]	68.5	42.1
Ours	66.1	48.6

## References

1. Wang et al., *Self-Supervised Equivariant Attention Mechanism for Weakly Supervised Semantic Segmentation*, CVPR, 2020.
2. Ahn et al., *Learning Pixel-Level Semantic Affinity with Image-Level Supervision for Weakly Supervised Semantic Segmentation*, CVPR, 2018.
3. Li et al., *Pseudo-Mask Matters in Weakly-Supervised Semantic Segmentation*, ICCV, 2021.

Kaalen, Stefan  
KTH / Scania

## SMP-tool for quantitative analysis of systems

Systems are growing more and more complex, which makes research and development increasingly relying on model-based development in order to ensure the safety of cyber-physical systems. Stateflow is a tool that supports modeling of systems as finite-state machines and has become the industrial standard practice in among others the automotive industry. However, Stateflow is limited in that it does not explicitly support modeling of stochastic processes, which are essential in model-based safety analysis. In order to overcome this, I have together with my colleagues developed SMP-tool that allow for modeling systems as Stochastic StateFlow (SSF) Models, and analysis of these models by studying the underlying stochastic process on the form of a generalized semi-Markov process.

Kaalen, Stefan  
KTH / Scania

# SMP-tool for quantitative analysis of systems

Stefan Kaalen, KTH  
Mechatronics

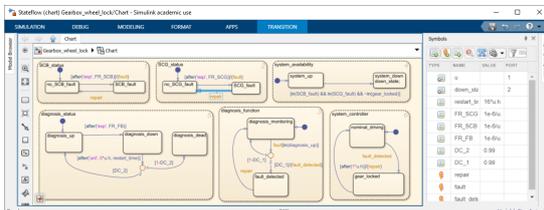


## Abstract

Systems are growing more and more complex, which makes research and development increasingly relying on model-based development in order to ensure the safety of cyber-physical systems. Stateflow is a tool that supports modeling of systems as finite-state machines and has become the industrial standard practice in among others the automotive industry. However, Stateflow is limited in that it does not explicitly support modeling of stochastic processes, which are essential in model-based safety analysis. In order to overcome this, I have together with my colleagues developed SMP-tool that allow for modeling systems as Stochastic StateFlow (SSF) Models, and analysis of these models by studying the underlying stochastic process on the form of a generalized semi-Markov process.

## SSF models

**SSF models** is a stochastic extension of a subset of **Stateflow**. The subset has been chosen according to utility found through numerous case studies joint with a desire to produce a **safe subset** [1]. SSF models can be modeled, although not analyzed directly, in Stateflow. The figure below presents an SSF model in Stateflow of a case study of a subsystem of a gearbox.



The state marked "down\_state" representations the system failure of the wheels of the vehicle locking at their current position caused by an erroneous actuation of the gears in the gearbox. The case study is further explained in [2]

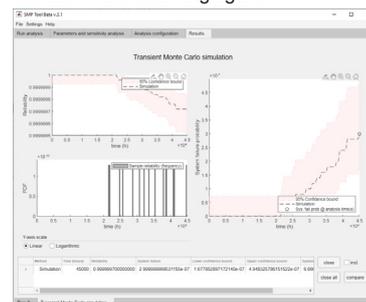
SSF models extends Stateflow both with the option of assigning probability distributions to the waiting time of transitions and with the option to assign discrete probabilistic choices of the destination state of transitions. For full syntax and semantics of SSF models, see [2]

## SMP-tool

SMP tool has the ability to perform multiple types of analysis of SSF models modeled in Stateflow. The analyses include transient analyses of the reliability, parameter sensitivity analysis, and steady state analysis. The tool is free and can be downloaded from [3].

The tool has a simulation engine for SSF models and a symbolic/numerical engine for SSF models where the underlying stochastic process is a Hierarchical Semi-Markov Processes (HSMP) [4].

The result from a transient analysis of the case study presented earlier is presented in the following figure.



The figure presents how the probability of a system failure develops over time given the parameters specified in the model. By a sensitivity analysis, different parameter configurations, and thereby system specifications, can be found which delivers the same reliability.

## References

1. MathWorks Advisory Board. 2020. Control algorithm modeling guidelines using MATLAB, Simulink, and Stateflow. Technical Report. SRI International. <https://www.mathworks.com/solutions/mab-guidelines.html>. Cited 2021-12-22
2. S. Kaalen, A. Hampus, M. Nyberg, and O. Mattsson. 2022. A stochastic extension of Stateflow. In IPCE '22 Proceedings of the 13th ACM/SPEC international conference on performance engineering, 2022. ACM (unpublished).
3. kth.se/itm/smp-tool. Cited 2021-12-22
4. Stefan Kaalen, Mattias Nyberg, and Olle Mattsson. 2021. Transient Analysis of Hierarchical Semi-Markov Process Models with Tool Support in Stateflow. In Quantitative Evaluation of Systems: 18th International Conference, Proceedings, Springer Nature, 2021, p. 105-126. Springer nature, 105—126

## SMP-tool

We have delivered a stochastic extension of Stateflow for the purpose of evaluation the performance of safety critical systems.

Future work includes extending the symbolic/numerical engine to handle all SSF models. Furthermore, to make SMP-tool tractable for use in the industry, future works includes support for modeling complex systems as SSF models.

## Explainable Reasoning and Decision-Making: From Humans to Machines

The line of work that leads to my dissertation studies automated reasoning and its intersection with human decision-making.

Most of the works address one of the following two questions: i) How can principles of human reasoning and decision-making be applied to drawing explainable inferences from knowledge bases with conflicting statements? ii) How can we improve agility and human explainability of complex 'intelligent' software systems?

The research subject of Question i) is formal argumentation, a graph-based method for non-monotonic reasoning, and the primary method is formal analysis. The main research results are different formal methods to ensure consistency when drawing repeated inferences from changing argumentation graphs, and analyses of the ability of different inference functions to support these approaches; in particular, a novel bridge between formal argumentation as a form of non-monotonic reasoning and economically rational decision-making is built. Regarding Question ii), the research subject is (the engineering of) multi-agent systems, which is studied from engineering and human-computer interaction points of view. The main results are new perspectives on and approaches to deploying agents in multi-agent systems in dynamic, Web-based environments and empirical results on how multi-agent systems can be better explained to human users.

Kampik, Timotheus  
Umeå University

## Explainable Reasoning and Decision-Making: From Humans to Machines

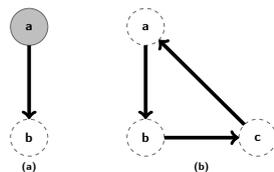
Timotheus Kampik, Umeå University  
Department of Computing Science  
Supervisors: Helena Lindgren and Juan Carlos Nieves



### In a Nutshell

The line of work that leads to my dissertation studies automated reasoning and its intersection with human decision-making. Most of the works address one of the following two questions: *i)* How can principles of human reasoning and decision-making be applied to drawing *explainable* inferences from knowledge bases with conflicting statements? *ii)* How can we improve agility and human explainability of complex "intelligent" software systems? The research subject of Question *i)* is formal argumentation, a graph-based method for non-monotonic reasoning, and the primary method is formal analysis. The main research results are different formal methods to ensure consistency when drawing repeated inferences from changing argumentation graphs, and analyses of the ability of different inference functions to support these approaches; in particular, a novel bridge between formal argumentation as a form of non-monotonic reasoning and economically rational decision-making is built. Regarding Question *ii)*, the research subject is (the engineering of) *multi-agent systems*, which is studied from engineering and human-computer interaction points of view. The main results are new perspectives on and approaches to deploying agents in multi-agent systems in dynamic, Web-based environments and empirical results on how multi-agent systems can be better explained to human users.

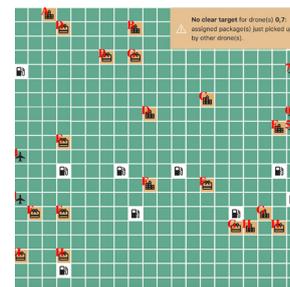
### Example: 'Formal' Part



**Figure 1:** Inconsistent preferences: using many *abstract argumentation* reasoning methods, the left graph implies  $\{a\}$  is preferred over  $\{\}$ , while the right graph implies  $\{\}$  is preferred over  $\{a\}$ .

This semi-formal example illustrates how most inference functions of abstract argumentation, in which conflicts in a set of *arguments* (for example: logical statements, business rules, claims in legal proceedings, *et cetera*) are modeled as a directed graph, violate the *consistent preferences* principle of economic rationality: given any set of choice items  $A$ , a rational agent consistently chooses the same items  $A^* \subseteq A$ , which implies that  $\forall A_c \subseteq A$ , such that  $A_c \neq A^*$ ,  $A^*$  is preferred over  $A_c$ . Given a set of options  $A' \supseteq A$ , the agent must choose  $A'^*$  so that  $A'^* = A^*$  or  $A'^* \not\subseteq A$ . This model is too simple to guide real-life decision-making (as has been shown by a range of behavioral economics research). However, it can be used as a sanity check for decision and reasoning algorithms. *E.g.*, the figure above shows that almost all of the well-established inference functions of abstract argumentation are not compliant with properties of economic rationality.

### Example: 'Engineering' Part



**Figure 2:** A multi-agent simulation of a drone delivery scenario, with an explanation message for a human supervisor in the top right corner.

The figure above shows a drone delivery simulation. Each drone can be thought of as an *autonomous agents* that has a partial view of the world; its knowledge may be incomplete and inconsistent with the knowledge that other drones or the "global" operators have. Hence, unexpected situations may occur, such as two drones attempting to pick up the same package, which in turn may result in a drone having to change directions mid-way. To make the overall behavior of the agents *explainable* to human operators, the state of all agents needs to be aggregated and filtered, and it is not exactly clear how to make the right trade-off to achieve explanation granularity that is useful, yet concise. In a human-interaction case study, we have compared different approaches to filtering explanations.

### Selected Publications

- [I] Ensuring reference independence and cautious monotony in abstract argumentation. Kampik, Nieves & Gabbay. *International Journal of Approximate Reasoning*. 2022
- [II] The quest of parsimonious XAI: A human-agent architecture for explanation formulation. Mualla, Tchappi, Kampik, Najjar, Calvaresi, Abbas-Turki, Galland & Nicolle. *Artificial Intelligence*. 2022
- [III] Governance of autonomous agents on the Web: Challenges and opportunities. Kampik, Mansour, Boissier, Kirrane, Padget, Payne, Singh, Tamma & Zimmerman. *ACM Transactions of Internet Technologies*. 2022
- [IV] Abstract Argumentation and the Rational Man. Kampik & Nieves. *Journal of Logic and Computation*. 2021
- [V] Argumentation-based health information systems: A design methodology. Lindgren, Kampik, Guerrero Rosero, Blusi & Nieves. *IEEE Intelligent Systems*. 2021
- [VI] The burden of persuasion in abstract argumentation. Kampik, Gabbay & Sartor. *International Conference on Logic and Argumentation*. 2021
- [VII] Explanations of non-monotonic inference in admissibility-based abstract argumentation. Kampik & Ćyras. *International Conference on Logic and Argumentation*. 2021
- [VIII] The degrees of monotony-dilemma in abstract argumentation. Kampik & Gabbay. *European Conference on Symbolic and Quantitative Approaches with Uncertainty*. 2021
- [IX] Autonomous agents on the edge of things (demonstration). Kampik, Gomez, Ciorte & Mayer. *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*. 2021
- [X] A Framework for collaborative and interactive agent-oriented developer operations (demonstration). Amaral, Kampik & Cranefield. *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 2020

## MM-wave channel sounding for indoor positioning

In this project we investigate the millimeter-wave (mm-wave) channel capability of being utilized for highly accurate indoor positioning purposes such as smart factory, sensing ,etc. Highly accurate radio based positioning relies on the additional information provided by Multi-Path Components (MPCs) which act as Virtual Anchor (VA) points besides the Line of Sight (LoS) and physical anchor. Hence in the first stage of the project we have focused on characterizing the behavior of the MPCs over time, e.g. the number of tractable ones and their life time, in mm-wave channel, by analyzing the real scenario measured data with high resolution in frequency and spatial domains.

# MM-wave channel sounding for indoor positioning



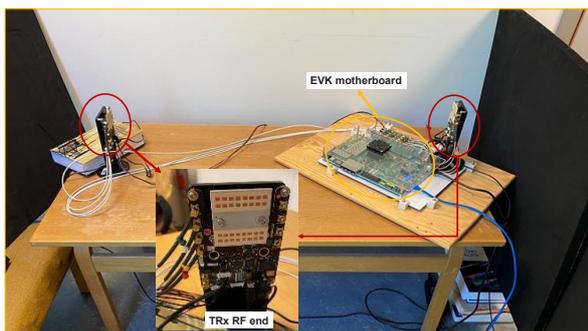
Hedieh Khosravi, Lund university  
Department of Electrical and Information Technology

## Research goals

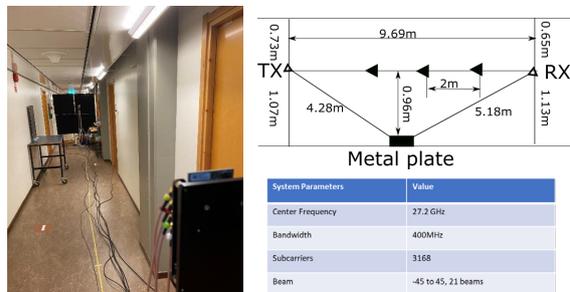
In this project, we investigate the millimeter-wave (mm-wave) channel capability of being utilized for highly accurate indoor positioning purposes such as smart factory, sensing ,etc. Highly accurate radio based positioning relies on the additional information provided by Multi-Path Components (MPCs), which act as Virtual Anchor (VA) points besides the Line of Sight (LoS) and physical anchor. Hence in the first stage of the project we have focused on characterizing the behavior of the MPCs over time, e.g. the number of tractable ones and their life time, in mm-wave channel, by analyzing the real scenario measured data with high resolution in delay and spatial domains.

## Measurement system

A Sivers-Qamcom 5G transceiver setup is being used to sound the mm-wave channel at 28 GHz. It is a 16+16 channel beamforming transceiver with a complete radio front-end and control functions optimized for high performance 5G NR applications. The setup consists of RF ends and a Xilinx Evaluation Kit (EVK) motherboard which creates the based band signal at 184.32 MHz and configures the beamforming at FPGA. Sivers setup can cover 24-29.5 GHz with the bandwidth of up to 400 MHz and have integrated transmitter and receiver beam book for beam steering and it also supports reconfigurable beam steering, e.g. different beam steering combinations can be activated at the FPGA level or high-level control. The channel sounding measurement can be done at millisecond level with this setup, so it can support indoor scenarios with reasonable mobility and dynamic.

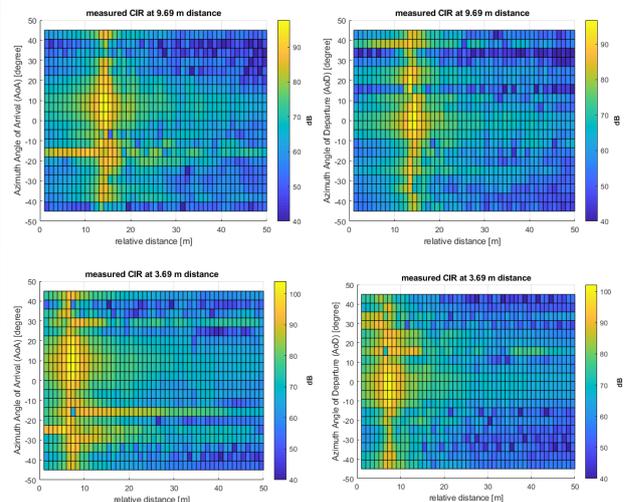


The channel impulse response in the azimuth domain has been shown below for two different Rx positions. Some strong MPCs coming from different directions besides LoS can be seen in both snapshots which is promising for the continuation of the work. To be able to stochastically track the evolution of the MPCs over time, e.g. detect and associate them, the more comprehensive measurement campaigns are planned to be performed.



## Measurement campaign

In our measurements we aim to cover typical indoor scenarios so that the extracted channel impulse responses we rely on, are related to the real environments. The results of one of the measurement campaigns we have done are presented in this section. As shown further, Tx and Rx are placed in a corridor at the department with about 10 m distance at the start point, approaching each other. A metal plate is placed in between them to assure the existence of a strong reflection besides LoS.



## Alternating Stutter Bisimulation

We want to use a fragment of Linear Temporal Logic (LTL) without the next operator to specify safety-critical requirements and synthesize a robust controller that fulfills those specifications on a discrete-time continuous state system that is subject to disturbances. Essentially, the robust controller must ensure that specific subsets of the state space are visited in an order which is allowed by the formal LTLnn specification.

However, the synthesis method cannot be applied directly on a continuous state space because it takes finite-state transition systems as input. One way to bridge this gap is to divide the continuous state space into a finite partition and let each block of the partition be one state in a transition system, which is called an abstract system. We introduce the alternating stutter bisimulation relation to be a basis for constructing the partition.



## Alternating Stutter Bisimulation

Jonas Krook, Robi Malik, Sahar Mohajerani, Martin Fabian  
krookj@chalmers.se, robi@waikato.ac.nz, mohajera@chalmers.se, fabian@chalmers.se

### Defining a partition

For an equivalence relation  $\mathcal{R} \subseteq S \times S$ , the equivalence class of  $s \in S$ , denoted  $[s]_{\mathcal{R}}$ , is the set  $\{s' \in S \mid (s, s') \in \mathcal{R}\}$ . The equivalence classes of  $\mathcal{R}$  form a partition of  $S$ , wherein they are referred to as blocks. We call the union of any number of equivalence classes for a superblock, and the set of all superblocks of  $\mathcal{R}$  for  $SB(\mathcal{R})$ .

A transition system is a tuple  $G = \langle S, \Sigma, \delta, S^{\circ}, AP, L \rangle$  where  $S$  is a set of states;  $\Sigma$  is a set of transition labels;  $\delta \subseteq S \times \Sigma \times S$  is a transition relation;  $S^{\circ} \subseteq S$  is a set of initial states;  $AP$  is a set of atomic propositions;  $L: S \rightarrow 2^{AP}$  is a state labelling function.

A path fragment of  $G$  is a sequence of states  $\pi = s_1 s_2 s_3 \dots \in S^*$  such that  $(s_i, \sigma, s_{i+1}) \in \delta$  for some  $\sigma \in \Sigma$  for all  $i$ .

We say that  $\langle G, s_1 \rangle \models [s_1]_{\mathcal{R}} \mathcal{U} T$ , for  $T \in SB(\mathcal{R})$ , if there exists an  $i > 0$  for each infinite path fragment  $\pi = s_1 \dots s_i s_{i+1}$  such that  $s_j \in [s_1]_{\mathcal{R}}$  for all  $j \leq i$  and  $s_{i+1} \in T$ . Furthermore,  $\langle G, s_1 \rangle \models [s_1]_{\mathcal{R}} \mathcal{W} T$  allows also paths with  $s_i \in [s_1]_{\mathcal{R}}$  for all  $i > 0$ .

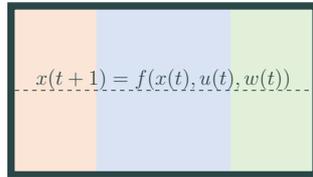
A controller for  $G$  is a function  $C: S^+ \rightarrow 2^{\Sigma}$ . A positional controller is a function  $C: S \rightarrow 2^{\Sigma}$ . The transition system resulting from controlling  $G$  by  $C$  is denoted  $C/G$ .

Let  $\mathcal{R}$  be an equivalence relation over  $S$  and let  $(s, t) \in \mathcal{R}$ .  $\mathcal{R}$  is an alternating stutter bisimulation iff

- (i)  $L(s) = L(t)$
- (ii) if, for some positional controller,  $\langle C_s/G, s \rangle \models [s]_{\mathcal{R}} \mathcal{U} T$  for some  $T \in SB(\mathcal{R})$ , then there exists a positional controller  $\langle C_t/G, t \rangle \models [s]_{\mathcal{R}} \mathcal{U} T$
- (iii) same as (ii) but with  $\mathcal{W}$ .

### Abstraction that preserves $LTL_{\setminus \circ}$ specifications under robust control

We want to use a fragment of Linear Temporal Logic (LTL) without the next operator to specify safety-critical requirements and synthesize a robust controller that fulfills those specifications on a discrete-time continuous state system that is subject to disturbances. Essentially, the robust controller must ensure that specific subsets of the state space are visited in an order which is allowed by the formal  $LTL_{\setminus \circ}$  specifica-

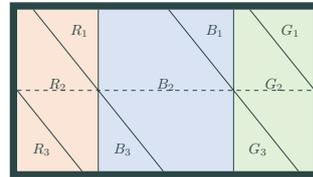


An abstract controller for the abstract system decides on control actions in the form of sets of allowed transitions, and the disturbance or process noise determines which of these allowed transitions are taken. The choices available to the abstract controller in an abstract state are based on the existence of concrete robust positional controllers that can robustly control the concrete system from the states in the corresponding source block to states in a target superblock. The blocks (equivalence classes) of alternating stutter bisimulations are defined in such a way that the system can be controlled to the same target superblocks from any of the source block's states.

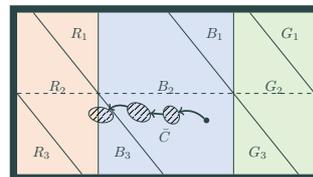
For instance, if a concrete positional controller (e.g.  $\bar{C}$ ) can control the system from one state in block  $B_2$  to a set of states in the superblock  $B_3 \cup R_2$ , then an abstract controller can choose the abstract states  $B_3$  and  $R_2$  as

the next possible states from abstract state  $B_2$ . This works since it must exist concrete controllers for any state in  $B_2$  that can control to  $B_3 \cup R_2$ .

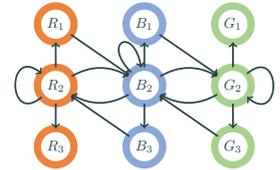
A self-loop is added to an abstract state if there exists a concrete positional controller that lets the concrete system remain in the corresponding block forever.



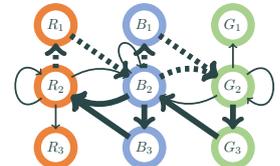
the next possible states from abstract state  $B_2$ . This works since it must exist concrete controllers for any state in  $B_2$  that can control to  $B_3 \cup R_2$ . A self-loop is added to an abstract state if there exists a concrete positional controller that lets the concrete system remain in the corresponding block forever.



into a finite partition and let each block of the partition be one state in a transition system, which is called an abstract system. We use the equivalence classes of an alternating stutter bisimulation as the partition. The transitions in the abstract transition system are then based on how the original, or concrete, dynamical system can be robustly controlled within and between the blocks.



We can now synthesize a controller for the abstract system such that the red and green abstract states are visited infinitely often. Every abstract control action has a corresponding concrete controller forcing the transition, so a concrete controller fulfilling the requirement can be implemented as a sequence of concrete positional controllers.



## On Joint State Estimation and Model Learning using Gaussian Process Approximations

State estimation is of interest in essentially every sector of science and engineering. Typically, techniques for state estimation require the specification of a dynamical model of the system in question. It is often possible to derive a partial description of the system dynamics, but depending on the modeling assumptions, this can potentially lead to bad state estimates, due to an insufficient description of the dynamics. This project explores the combination of such a partial dynamical description with a generic black-box structure to allow online learning of parts of the system dynamics. In this way, the model can be improved over time, as more measurements have been obtained, and in extension improve the resulting state estimate. We provide some initial results in this regard.

Kullberg, Anton  
Linköping University

WASP WALLENBERG AI,  
AUTONOMOUS SYSTEMS  
AND SOFTWARE PROGRAM

## On Joint State Estimation and Model Learning using Gaussian Process Approximations

Anton Kullberg, PhD Student, Linköping University  
Div. of Automatic Control  
Supervisors: Assoc. Prof. Gustaf Hendeby (LiU) and Assoc. Prof. Isaac Skog (LiU)

li.u  
LINKÖPINGS UNIVERSITET

### Motivation & Research Goals

State estimation is of interest in essentially every sector of science and engineering. Typically, techniques for state estimation require the specification of a dynamical model of the system in question. It is often possible to derive a partial description of the system dynamics, but depending on the modeling assumptions, this can potentially lead to bad state estimates, due to an insufficient description of the dynamics. This project explores the combination of such a partial dynamical description with a generic black-box structure to allow online learning of parts of the system dynamics. In this way, the model can be improved over time, as more measurements have been obtained, and in extension improve the resulting state estimate. We provide some initial results in this regard.

### Methods

We consider the general discrete-time description of a dynamical system given by

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{g}_f(\mathbf{w}_k + \mathbf{g}(\mathbf{x}_k, \mathbf{u}_k)) \\ \mathbf{y}_k &= \mathbf{h}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{e}_k. \end{aligned}$$

Here,  $\mathbf{x}_k$ ,  $\mathbf{u}_k$ ,  $\mathbf{y}_k$  are the state, input and measurement at time  $k$ , respectively. Further,  $\mathbf{w}_k$ ,  $\mathbf{e}_k$  are mutually independent white noise processes, particularly,  $\mathbf{w}_k \sim \mathcal{N}(0, \mathbf{Q})$  and  $\mathbf{e}_k \sim \mathcal{N}(0, \mathbf{R})$ . The measurement function,  $\mathbf{h}$ , is assumed known and parts of the system dynamics,  $\mathbf{f}$  and  $\mathbf{g}_f$ , are assumed known. Lastly, the function  $\mathbf{g}$  captures any system dynamics that are not described by  $\mathbf{f}$ .

We model the function  $\mathbf{g}$  as a basis function expansion

$$\mathbf{g}(\mathbf{x}_k) = \sum_j \phi^j(\mathbf{x}_k) \theta^j,$$

with  $\phi^j$  chosen as radial basis functions, i.e.,  $\phi^j$  is a function of  $\|\mathbf{x}_k - \xi^j\|$ , where  $\xi^j$  is the basis function center. The basis functions are placed in the region of the state space in which there is unknown dynamics. Essentially, this constitutes a grid of basis functions, where the extent of the grid determines in what regions unknown dynamics can be learned and the density of the grid determines the fidelity of the learned dynamics. As such, if the state space region of interest is large, the number of parameters  $\theta$  is large. To facilitate learning of the parameters  $\theta$ , an augmented state vector is constructed as  $\mathbf{x}_k^e = [\mathbf{x}_k^\top \ \theta^\top]^\top$ . This enables us to estimate the state trajectories and learn (parts of) the model online in a joint fashion using an EKF [1]. Even though this is theoretically computationally efficient, as the number of parameters grows beyond a few thousand, this is not feasible to do in real-time, limiting the model to either a small state space region or low dynamical fidelity [2]. To resolve this issue, we choose  $\phi^j$  such that

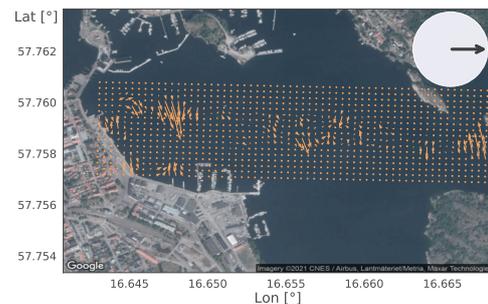
$$\phi^j(\|\mathbf{x}_k - \xi^j\|) \equiv 0, \quad \|\mathbf{x}_k - \xi^j\| > c^j.$$

As such, each basis function  $\phi^j$  only contributes to the function value in a region close to its center  $\xi^j$ , limiting the number of parameters necessary for each function evaluation. With a few modifications to the EKF recursions, this enables real-time online joint state estimation and model learning [2].

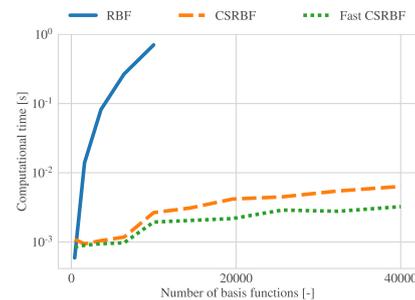
### References

- [1] Learning Driver Behaviors Using A Gaussian Process Augmented State-Space Model  
A. Kullberg, I. Skog and G. Hendeby  
International Conference on Information Fusion (FUSION), July 2020
- [2] Online Joint State Inference and Learning of Partially Unknown State-Space Models  
A. Kullberg, I. Skog and G. Hendeby  
Transactions on Signal Processing 69 2021
- [3] Learning Motion Patterns in AIS Data and Detecting Anomalous Vessel Behavior  
A. Kullberg, I. Skog and G. Hendeby  
International Conference on Information Fusion (FUSION), Nov 2021

### Selected Results



[3] Learned motion dynamics of vessels traveling in a harbor region in Västervik, Sweden. Visualized is the Cartesian acceleration of the vessels traveling out of port. The acceleration was learned from historical data of vessels traveling throughout the port, provided by the Swedish Defense Research Agency (FOI).



[2] Comparison of the computational time for a single propagation step in the EKF. Blue is for the standard formulation using Gaussian radial basis functions (global support). Orange and green both use compact basis functions but differ slightly in how they identify which basis functions to use. After about 10000 basis functions, the global formulation is no longer runnable at all, whereas the compact formulation easily handles a large number of basis functions.

## On Rendezvous in Autonomous Cooperative Landings

We investigate the rendezvous problem for the autonomous cooperative landing of an unmanned aerial vehicle (UAV) on an unmanned surface vehicle (USV). The rendezvous problem is challenging due to several reasons, for example, sudden communication losses or strong disturbances acting on the agents can lead to disastrous consequences. Moreover, even the basic tasks to determine if the rendezvous is possible or not and what strategy to employ when the rendezvous location has to be updated can be complex. Our goal is to create a rendezvous algorithm with an online update of the rendezvous location such that convergence is guaranteed. The preliminary proposed algorithm requires the agents to exchange information only when necessary to maintain the convergence.

# On Rendezvous in Autonomous Cooperative Landings

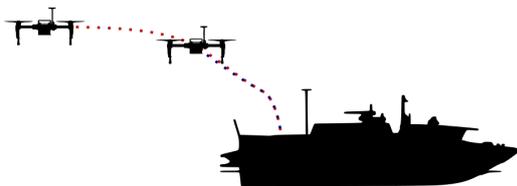
Dženan Lapandić, KTH Royal Institute of Technology  
Division of Decision and Control Systems  
Supervisor: Prof. Bo Wahlberg



## Motivation & Research Goals

We investigate the rendezvous problem for the autonomous cooperative landing of an unmanned aerial vehicle (UAV) on an unmanned surface vehicle (USV). The rendezvous problem is challenging due to several reasons [2], for example, sudden communication losses or strong disturbances acting on the agents can lead to disastrous consequences. Moreover, even the basic tasks to determine if the rendezvous is possible or not and what strategy to employ when the rendezvous location has to be updated can be complex. Our goal is to create a rendezvous algorithm with an online update of the rendezvous location such that the convergence is guaranteed. The preliminary proposed algorithm requires the agents to exchange information only when necessary to maintain the feasibility.

## Methods



We consider two **heterogeneous** agents with **nonlinear dynamics** and additive disturbances. Each agent solves a corresponding distributed optimal control problem formulated as a **Model Predictive Control** problem penalizing the distance to the rendezvous location  $\theta$  while satisfying state and input constraints.

The **control objective** is to steer the relevant states of every agent  $y_i$  to a rendezvous point  $\theta \in \Theta \subseteq \mathbb{R}^p$  in finite time.

- It is assumed that the initial rendezvous location is feasible
- The time planning horizon  $T$  is long enough to reach at least one  $\theta$  in the rendezvous set  $\Theta$ .
- The agents update and share the rendezvous location only when they are not guaranteed to reach it, i.e. to maintain the feasibility

Based on the deviation of the predicted terminal state output from the rendezvous location  $V_o = \|\hat{y}_i(t_k + T; t_k) - \theta(t_k)\|^2$  the agent  $i$  updates  $\theta$  according to the online **update law**

$$\theta(t_{k+1}) = \begin{cases} \theta(t_k) & V_o \leq \varepsilon \\ \theta(t_k) - \eta v_\theta(t_k) & V_o > \varepsilon \end{cases}$$

where  $\eta$  and  $\varepsilon$  are tuning parameters and  $v_\theta(t_k)$  is defined as:

$$v_\theta(t_k) = \frac{\partial V_o}{\partial \theta(t_k)} \left\| \frac{\partial V_o}{\partial \theta(t_k)} \right\|^{-1}$$

Parameter  $\eta$  is a step size that must be chosen as a small value, in order to avoid overshooting, and it quantifies the correction of  $\theta$  in the output space.

## References

[1] Aperiodic Communication for MPC in Autonomous Cooperative Landing. Dženan Lapandić, Linnea Persson, Dimos V. Dimarogonas, Bo Wahlberg. 7th IFAC Conference on NMPC 2021

[2] Model Predictive Control for Autonomous Ship Landing in a Search and Rescue Scenario. Linnea Persson and Bo Wahlberg. In AIAA SciTech 2019 Forum, 1169

## Selected Results

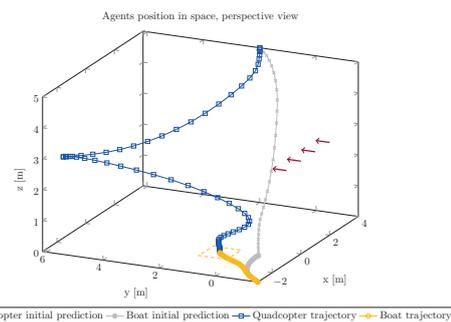
Feasibility in distributed MPC scenario with a common rendezvous location is challenging to guarantee due to

- communication issues and delays,
- disturbances that may affect one or several agents to be unable to reach the previously agreed rendezvous location,
- update law that may propose a new location which is not feasible for the other agent.

### Contributions [1]:

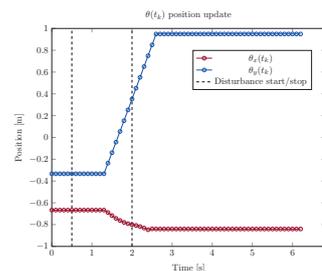
- Distributed rendezvous algorithm with **aperiodic communication** which eliminates unnecessary communication.
- **Time-varying distributed terminal sets** for tracking that depend on the rendezvous point.
- Proof that the proposed algorithm guarantees **recursive feasibility**.

### Simulation example: Autonomous cooperative landing



Arrows show wind direction and the yellow dashed polygon represents the boat landing platform.

Agents initiate the landing according to the initial rendezvous location. Due to the strong wind active for  $t = [0.5s, 2.0s]$ , the initial location becomes infeasible and has been updated using the update law to maintain the feasibility.



Larsson, Martin  
Lund University

## Sensor Node Calibration in Presence of a Dominant Reflective Plane

In this paper we study the problem of sensor network self-calibration in presence of a single reflective plane. We propose a three-step stratified approach utilizing a rank-1 constraint in the measurements: (i) In the case of time difference of arrival (TDOA) measurements, any offsets in the measurements are solved for. (ii) The heights of the receivers and senders relative to the plane are solved for. (iii) The planar receiver and sender positions are solved for. We evaluate our approach on synthetic and real data.

# Sensor Node Calibration in Presence of a Dominant Reflective Plane

Erik Tegler, Martin Larsson, Magnus Oskarsson, Kalle Åström  
Centre for Mathematical Sciences, Lund University



## Abstract

- In this paper we study the problem of sensor network self-calibration in presence of a single reflective plane.
- We propose a three-step stratified approach utilizing a rank-1 constraint in the measurements:
  1. In the case of time difference of arrival (TDOA) measurements, any offsets in the measurements are solved for.
  2. The heights of the receivers and senders relative to the plane are solved for.
  3. The planar receiver and sender positions are solved for.
- We evaluate our approach on synthetic and real data.

## Problem Formulation

Consider the problem of time of arrival (TOA) self-calibration in the presence of a single reflective plane. Then every receiver  $R_\Lambda$  has a virtual mirror receiver  $R_V$ , and there are two distance measurements  $D_\Lambda$  and  $D_V$  to the sender  $S$  given by

$$D_\Lambda^2 = \|R_\Lambda - S\|^2 = d^2 + (g - h)^2,$$

$$D_V^2 = \|R_V - S\|^2 = d^2 + (g + h)^2.$$

See figure to the right for notation. From these we can derive

$$D_\Delta = \frac{D_V^2 - D_\Lambda^2}{4} = gh,$$

$$D_\Sigma = \frac{D_V^2 + D_\Lambda^2}{2} = d^2 + g^2 + h^2.$$

Provided  $m$  receivers and  $n$  senders we get the rank-1 matrix

$$D_\Delta = \begin{pmatrix} g_1 h_1 & \cdots & g_1 h_n \\ \vdots & \ddots & \vdots \\ g_m h_1 & \cdots & g_m h_n \end{pmatrix} = \begin{pmatrix} g_1 \\ \vdots \\ g_m \end{pmatrix} \begin{pmatrix} h_1 & \cdots & h_n \end{pmatrix}$$

### Offset Estimation

In the case of time difference of arrival (TDOA), additional offsets  $o_j$  in the measurements  $Z_{\Lambda ij}$  and  $Z_{V ij}$  need to be estimated.

$$D_{\Lambda ij} = Z_{\Lambda ij} - o_j, \quad D_{V ij} = Z_{V ij} - o_j,$$

where  $i = 1, \dots, m$  and  $j = 1, \dots, n$ .  $D_\Delta$  becomes linear in  $o_j$ . The offsets can be found linearly by utilizing the rank constraint on  $D_\Delta$ .

### Height Estimation

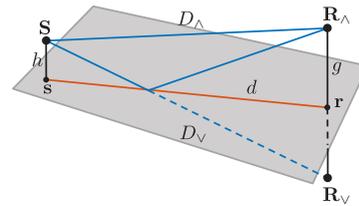
The heights  $g_i$  and  $h_j$  can be retrieved from a robust rank-1 approximation of  $D_\Delta$ , up to some unknown constant  $\lambda$ , since  $D_\Delta = (\lambda g) \left(\frac{1}{\lambda} h^T\right)$ .

### Planar Position Estimation

The planar positions  $r_i$  and  $s_j$  are retrieved by solving a lower dimensional TOA problem, where the distances  $d_{ij}$  depend on  $\lambda$ .

$$d_{ij}^2 = \|r_i - s_j\|^2 = D_{\Sigma ij} - \lambda^2 g_{ij}^2 - h_{ij}^2 / \lambda^2$$

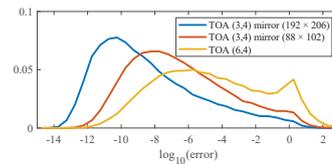
We present two new solvers for the minimal case  $(m, n) = (3, 4)$ .



## Experiments

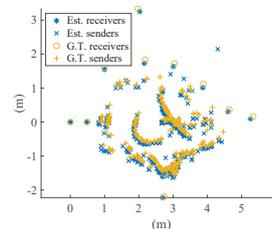
### Experiments on Synthetic Data

Numerical stability of the proposed solvers compared to the existing  $(m, n) = (6, 4)$  solver for noise-less data.

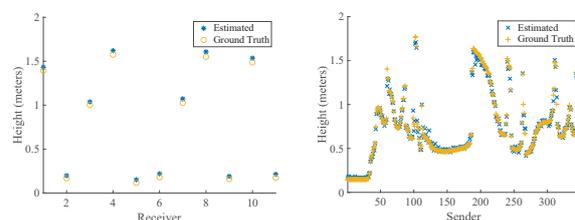


### Experiments on Real Data

Estimated planar positions of the receivers and senders compared to ground truth.



Estimated receiver and sender heights compared to ground truth.



Marta, Daniel  
KTH

## Human-Feedback Shield Synthesis for Perceived Safety in Deep Reinforcement Learning

Despite the successes of deep reinforcement learning (RL), it is still challenging to obtain safe policies. Formal verification approaches ensure safety at all times, but usually overly restrict the agent's behaviors, since they assume adversarial behavior of the environment. Instead of assuming adversarial behavior, we suggest to focus on perceived safety instead, i.e., policies that avoid undesired behaviors while having a desired level of conservativeness. To obtain policies that are perceived as safe, we propose a shield synthesis framework with two distinct loops: (1) an inner loop that trains policies with a set of actions that is constrained by shields whose conservativeness is parameterized, and (2) an outer loop that presents example rollouts of the policy to humans and collects their feedback to update the parameters of the shields in the inner loop.

# Human-Feedback Shield Synthesis for Perceived Safety in Deep Reinforcement Learning

Daniel Marta, KTH Royal Institute of Technology

RPL: Robotics Perception and Learning



## Abstract

Despite the successes of deep reinforcement learning (RL), it is still challenging to obtain safe policies. Formal verification approaches ensure safety at all times, but usually overly restrict the agent's behaviors, since they assume adversarial behavior of the environment. Instead of assuming adversarial behavior, we suggest to focus on perceived safety instead, i.e., policies that avoid undesired behaviors while having a desired level of conservativeness. To obtain policies that are perceived as safe, we propose a shield synthesis framework with two distinct loops: (1) an inner loop that trains policies with a set of actions that is constrained by shields whose conservativeness is parameterized, and (2) an outer loop that presents example rollouts of the policy to humans and collects their feedback to update the parameters of the shields in the inner loop.

## Methods

### Learning safety constraints from humans

- Inner-loop: takes advantage of self-play by sampling from shield distributions of human feedback [1]
- Outer-loop: updates a shield parameter distribution with human feedback [1]

### Human-feedback Shield distribution:

- Computed iteratively from human-feedback datasets in the outer loop.
- Maps high-level human feedback into shield parameter updates.

### Human-feedback Shield distribution:

- Computed iteratively from human-feedback datasets in the outer loop.
- Maps high-level human feedback into shield parameter updates.

$$\text{map}(g_j) = \begin{cases} \mu_h - \frac{|h_j|}{2} & \text{if } g_j = \text{very unsafe,} \\ \vdots & \vdots \\ \mu_h & \text{if } g_j = \text{fine,} \\ \vdots & \vdots \\ \mu_h + \frac{|h_j|}{2} & \text{if } g_j = \text{very safe.} \end{cases}$$

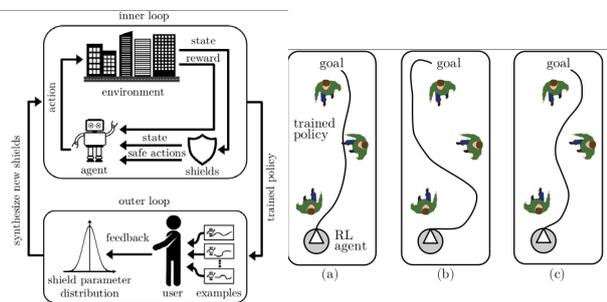
$${}^u \mu_h = \frac{1}{N_{\text{user}}} \sum_{j=1}^{N_{\text{user}}} \text{map}(g_j),$$

$${}^u \sigma_h^2 = \max \left( \frac{1}{N_{\text{user}}} \sum_{j=1}^{N_{\text{user}}} (\text{map}(g_j) - ({}^u \mu_h))^2, \sigma_{\text{min}}^2 \right)$$

$$\forall i: \text{KL}({}^u f_{h_i}, f_{h_i}) \leq \beta$$

$$p(\mathcal{G}|\mu, \sigma^2) = \prod_{j=1}^{|\mathcal{G}|} p(x_j|\mu, \sigma^2), x_j \in \mathcal{G}$$

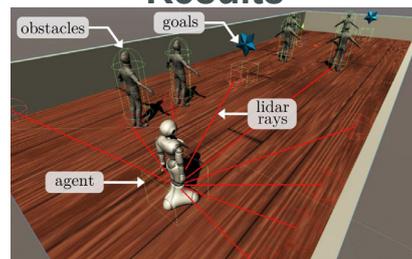
$$\tilde{p}_\theta(\mu|\mathcal{G}, \sigma_\theta^2) \propto p(\mathcal{G}|\mu, \sigma^2) p_\theta(\mu|\mu_\theta, \sigma_\theta^2)$$



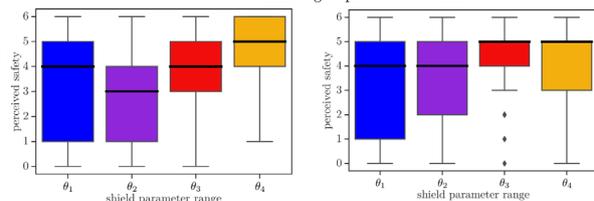
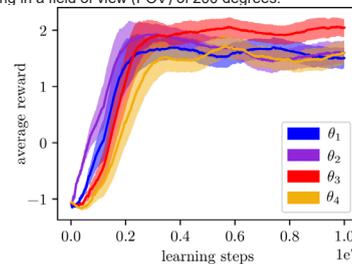
## References

1. D. Marta, C. Pek, G. I. Melsi3n, J. Tumova and I. Leite, "Human-Feedback Shield Synthesis for Perceived Safety in Deep Reinforcement Learning," in IEEE Robotics and Automation Letters, vol. 7, no. 1, pp. 406-413, Jan. 2022, doi: 10.1109/LRA.2021.3128237.

## Results



To evaluate perceived safety, we want to estimate how strong the force field should be, i.e., the shield parameter encodes how much the shield considers the full interaction force of the Social Force Model (SFM). We focus our shield synthesis from human-feedback, to address how our approach could have an impact in a robotic scenario. The state space is comprised by the agent's position and velocity, the velocity of other obstacles in the environment, and nine rays of a lidar-like sensor, commonly used in navigation robots. The agent's actions are composed of the accelerations in x- and y-directions, representing the driving force and a third action proportional to the interaction force of the SFM. For each ray, the agent detects either a goal, one of the humans or walls. The rays are one-hot encoded in addition to the distance between the robot and a specific element. In total, there are 9 rays opening in a field of view (FOV) of 200 degrees.



To access the validity of our approach, we ran a study with real humans. The study was run online using Amazon Mechanical Turk (AMT). In total, there were 92 unique participants (59 males, 33 female and none of other gender identities). Their age ranged from 23 to 65 years old, with a median of 34; the majority were in or had completed college education (N=78) and came from the US (N=71). 62 participants reported to have never or only seen robots in media, 14 to have interacted with one robot before, and 2 to do it on a regular basis.

Mayr, Matthias  
Lund University

## Learning Robot Tasks through Planning, Knowledge Integration and Multi-Objective Optimization

We introduce a framework for integrating planning with targeted learning of scenario-specific parameters. It uses a coarse-to-fine strategy: (1) the user provides a task goal in PDDL, (2) a plan (i.e., a sequence of skills) is generated and the learnable parameters of the skills are automatically identified. An operator then chooses (3) reward functions and hyperparameters for the (4) subsequent learning process. Learning is tightly integrated with a knowledge framework to support planning and to provide priors for learning and using multi-objective optimization, since objectives such as safety and task execution can often affect each other. Our system utilizes extended Behavior Trees for planning, execution. We adopt a multi-objective Bayesian optimization approach to learn the parameters of our tasks statistically efficient. Learning is done entirely in simulation and we use domain randomization techniques to ensure the results work in a real-world scenario.

# Learning Robot Tasks through Planning, Knowledge Integration and Multi-Objective Optimization

Matthias Mayr, Faseeh Ahmad, Konstantinos Chatzilygeroudis, Luigi Nardi and Volker Krueger



LUNDS UNIVERSITET  
Lunds Tekniska Högskola



UNIVERSITY OF PATRAS  
ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΑΣ



Stanford University

## How to safely learn new robot task with explainable policies?

### Motivation

- **Fast adaption** to new tasks
- **Explainable policy representation:** know *what* is performed *when* and *why*
- **Safe learning process**
- **Automated pipeline:** Little user interaction from goal definition to learned policies

### Skill-based systems

- Support planning of tasks
- **Reasoning alone often leads to complex systems**

*How to combine planning and learning to learn the tacit aspects of robot task?*

### Approach

#### Skills with Behavior Trees<sup>1</sup>

- Behavior trees as **reactive** and **parametric** policy representation
- **Human-readable** and **editable** as well as **expandable** and **modular**
- Usable for **planning** and **skill execution**
- **Skills expose parameters** such as:
  - Objects to manipulate
  - Conditions
  - Offsets for motions

#### Planning and Knowledge Integration

- *SkiROS2* as a **platform** for skills and the **world model**
- *PDDL* to formulate the **planning problem**
- Automatically generate planning domain
- The **world model** provides and stores **knowledge** about the given problem
- supplies **information** about the **learnable parameters**

#### Policy Optimization

Dynamical system in the form

$$\mathbf{x}_{t+1} = \mathbf{x}_t + M_{sim}(\mathbf{x}_t, \mathbf{u}_t, \phi_R)$$

with transition dynamics  $M_{sim}(\mathbf{x}_t, \mathbf{u}_t, \phi_R)$  modeled by a simulator.

- **Policy search** with black-box optimization algorithm: *Bayesian optimization*
- **Reward functions** can be select from a **library** by the operator
- **Multiple objectives** can be optimized concurrently

## Simulation-supported Learning

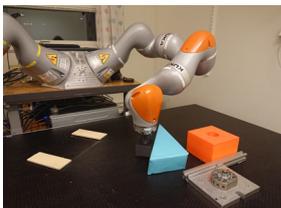
- **Less interaction time** with the robot
- **Safe** for robot and production material
- **Scales** with **cloud resources**
- Allows **object tracking** without a complex setup

### Domain Randomization

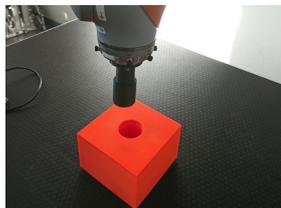
- learn more **robust policies** that generalize to reality
- emulate different configurations

### Examples for Learned Tasks

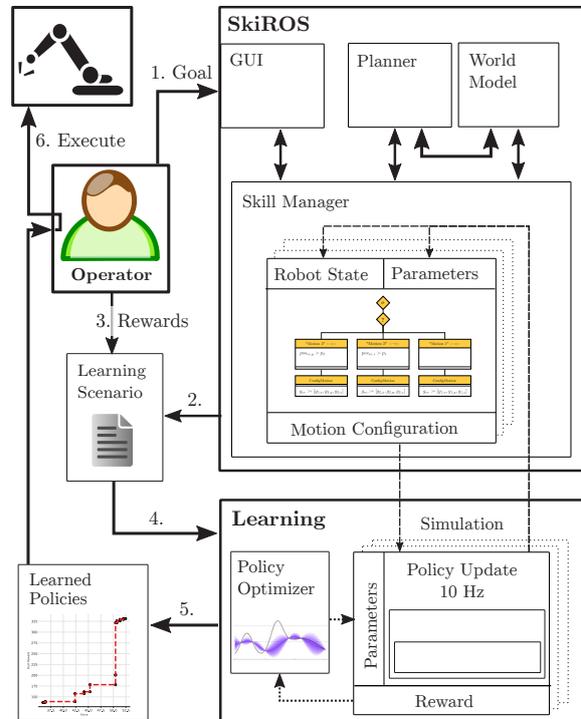
#### Push an Object



#### Peg Insertion



### Learning Pipeline



### Steps of the Learning Pipeline:

1. Goal definition by the operator
2. Plan generation and parameterization; learnable parameters are identified
3. Operator complements scenario with hyperparameters, rewards and objectives
4. Learning in simulation
5. Pareto front with the best policies
6. Operator selects solutions and executes them on the real system

### Future Work

1. **Learning of the task structure** including **recovery behaviors**
2. **Multi-fidelity optimization** that can include the real system
3. **Automatic reasoning** about rewards and objectives

### References

- (1) Rovida, F.; Grossmann, B.; Krueger, V. Extended Behavior Trees for Quick Definition of Flexible Robotic Tasks. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*; 2017; pp 6793–6800.
- (2) Mayr, M.; Chatzilygeroudis, K.; Ahmad, F.; Nardi, L.; Krueger, V. Learning of Parameters in Behavior Trees for Movement Skills. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*; 2021.

## Bootstrapped Representation Learning for Skeleton-Based Action Recognition

In this work, we study self-supervised representation learning for 3D skeleton-based action recognition. We extend Bootstrap Your Own Latent (BYOL) for representation learning on skeleton sequence data and propose a new data augmentation strategy including two asymmetric transformation pipelines. We also introduce a multi-viewpoint sampling method that leverages multiple viewing angles of the same action captured by different cameras. In the semi-supervised setting, we show that the performance can be further improved by knowledge distillation from wider networks, leveraging once more the unlabeled samples.

We conduct extensive experiments on the NTU-60 and NTU-120 datasets to demonstrate the performance of our proposed method. Our method consistently outperforms the current state of the art on both linear evaluation and semi-supervised benchmarks.

# Bootstrapped Representation Learning for Skeleton-Based Action Recognition

Olivier Moliner <sup>1,2</sup> Sangxia Huang <sup>2</sup> Kalle Åström <sup>1</sup>

<sup>1</sup>Lund University <sup>2</sup>Sony R&D Center Lund

SONY



LUND UNIVERSITY

## Introduction

We study self-supervised representation learning for 3D skeleton-based action recognition.

**Motivation** Fully-supervised action recognition algorithms require large datasets of 3D skeleton data with accurate annotations, which are time-consuming and costly to prepare.

**Goal** To learn semantic features from unlabeled 3D skeleton sequence data, making downstream task learning more label-efficient.

## Method

### Self-Supervised Skeleton Sequence Representation Learning with BYOL

Two networks, an online network and a target network, encode two augmented views of the same action sequence captured from different viewing angles. The online network is trained to predict the output of the target network, while the target network is updated with an exponential moving average of the online network.

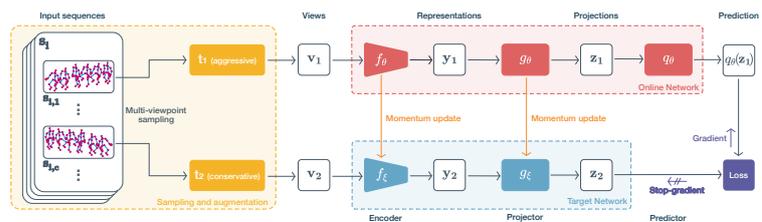


Figure 1: Overview of our proposed method.

## Our contribution

- ▶ A simple framework for self-supervised representation learning for skeleton-based action recognition based on BYOL.
- ▶ A data augmentation strategy for skeleton data based on two distinct transformation pipelines.
- ▶ A multi-viewpoint sampling method that makes better use of action sequences captured simultaneously by different cameras.
- ▶ We show that our method consistently outperforms the current state of the art on linear evaluation and semi-supervised tasks.

### Asymmetric Augmentation Pipelines

We tailored an aggressive data augmentation pipeline to learn semantically-relevant features, while the conservative pipeline reduces the distribution shift between self-supervised pre-training and supervised fine-tuning.

### Multi-Viewpoint Sampling

We leverage recordings of the same action sequences captured simultaneously from different angles by different cameras to learn representations that are invariant to changes of viewpoint and to different camera properties.

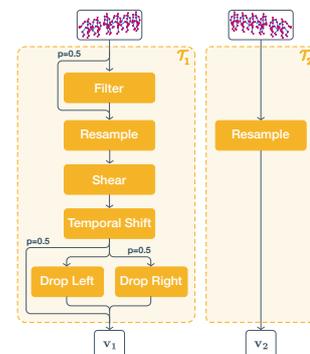


Figure 2: Asymmetric Augmentation Pipelines.

## Experimental Results

Method	NTU-60		NTU-120	
	CS	CV	CSub	CSet
ST-GCN (supervised)	88.5	94.3	83.0	85.1
LongT GAN	39.1	48.1	-	-
MS <sup>2</sup> L	52.6	-	-	-
PCRP	53.9	63.5	-	-
AS-CAL	58.5	64.8	48.6	49.2
Thoker et al.	76.3	85.2	67.1	67.9
3s-CrosSCLR	77.8	83.4	67.9	66.7
Ours	86.8	91.2	77.1	79.2

Table 1: Linear evaluation protocol on NTU-60 and NTU-120.

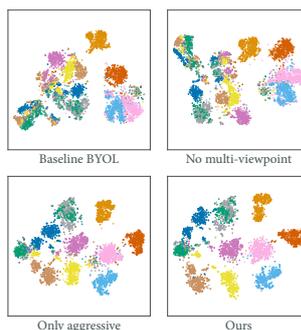


Figure 3: t-SNE projection of learned representations.

Method	Label fraction		
	1%	5%	10%
ST-GCN (supervised)	19.3	59.1	71.7
SESAR-KT	48.1	55.0	58.2
MS <sup>2</sup> L	33.1	-	65.2
ASSL	-	57.3	64.3
3s-CrosSCLR	51.1	-	74.4
Thoker et al.	35.7	59.6	65.9
Ours 1×, distilled	79.4	83.6	84.6
Ours 2×	79.3	84.5	86.0

Table 2: Semi-supervised learning on NTU-60 (Cross-Subject).

## Representing Temporal Data in Semantic Graphs

Semantic parsing is the process of taking input data and translating it into a structured representation of its meaning, for example a graph. The input data has traditionally been in the form of text, but in later years semantic parsing of video content has been an active research topic. Semantic parsing of video is useful for automating several otherwise manual tasks. Examples include automatic trailer creation, automatic caption generation, compliance checking, and knowledge extraction. We introduce the research project, which aims to design a suitable formalism for semantic representation of video and multimodal content, as well as developing tools to generate said representations.

Mollevik, Iris  
Umeå University / Codemill AB

## Representing Temporal Data in Semantic Graphs

Iris Mollevik

### Research project

- Semantic parsing = to translate input data into a representation, suitable for further processing
- The traditional input has been text
- Our focus: Semantic parsing of video and multimodal data

### Applications in video industry

- Automatic trailer creation
- Caption generation
- Compliance checking
- Advanced search among video content

WASP | WALLEBERG AI,  
AUTONOMOUS SYSTEMS  
AND SOFTWARE PROGRAM

CODEMILL

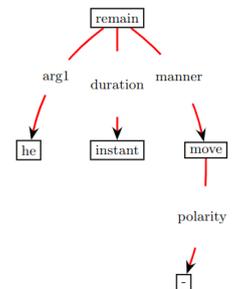


Making computers understand  
video content can help the video  
industry automate manual tasks.

For example trailer creation and  
compliance checking.

It can also facilitate advanced  
search among video material.

Abstract Meaning Representation (AMR) is a graph format used to capture the meaning of single text sentences.



AMR graph for the sentence  
"He remained motionless for an instant".  
A very simple example.

AMR has desirable properties, however it was developed for single sentences of text. We are working on extending this formalism to be able to capture longer sequences of text, including temporal information. In the longer term, we would like to extend it to capture video and other data as well.

If the AMR format proves unsuitable, we will instead develop a different format.

This research is still in a very early stage.

Narri, Vandana  
KTH / Scania AB

## Set-Membership Estimation in Shared Situational Awareness for Automated Vehicles in Occluded Scenarios

The objective of this project is to model, formalize, and analyse a shared situational awareness framework for the ego-vehicle and extended vehicles, i.e., connected vehicles and infrastructure. Shared situational awareness is the ability to perceive and comprehend the traffic situation and to predict the intent of vehicles and road users in the surrounding of the ego-vehicle using local and connected sensors. This framework will allow to orchestrate the utilization of shared resources in complex and crowded environments and to define which kind of information each Connected and Autonomous Vehicle (CAV) and the infrastructure should share. Safety-critical application such as these require robust guarantees for the estimation of the road users.

# Set-Membership Estimation in Shared Situational Awareness for Automated Vehicles in Occluded Scenarios



Vandana Narri

ATS Research, Scania CV AB : vandana.narri@scania.com  
Division of Decision and Control Systems, KTH : narri@kth.se



## Motivation & Research goals

- The objective of this project is to model, formalize, and analyse a shared situational awareness framework for the ego-vehicle and extended vehicles, i.e., connected vehicles and infrastructure.
- Shared situational awareness is the ability to perceive and comprehend the traffic situation and to predict the intent of vehicles and road users in the surrounding of the ego-vehicle using local and connected sensors.
- This framework will allow to orchestrate the utilization of shared resources in complex and crowded environments and to define which kind of information each Connected and Autonomous Vehicle (CAV) and the infrastructure should share.
- Safety-critical application such as these require robust guarantees for the estimation of the road users.

## Background

- Local CAV sensors typically provide a limited understanding of the environment due to limited sensor range, blind spots, and occlusions in the environment.
- Vehicle to vehicle (V2V) communication and vehicle to infrastructure (V2I) communication based on 5G or IEEE 802.11p standards, can help gather more information about the environment, and address the shortcomings of CAV sensors.
- CPM (Collaborative Perception Message) service supports sharing information between ITS-Ss (Intelligent Transportation System – Stations) [1].
- The main research areas are connectivity (enabled by V2I and V2V), cooperative driving, situational awareness, set-based estimation and traffic flow optimization.

## Problem Formulation

- The problem considered in this research work is formulated around a scenario as shown in Figure 1.
- This scenario consist of two-lane road with a sidewalk on each side of the road and a pedestrian crossing.
- The ego-vehicle (blue bus) is traveling from left to right and is approaching the pedestrian crossing. The ego-vehicle is equipped with a sensor having a field of view represented by the blue-shaded circle segment.
- In this scenario, two additional sensors are included. One on the approaching CAV represented by yellow-shaded circle segment and other on the connected road-side sensor units represented by green-shaded circle segment.

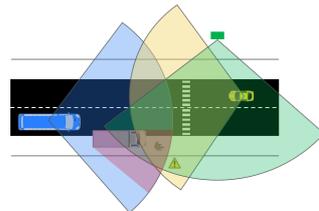


Figure 1 : The ego-vehicle with one local sensor, an additional V2V sensor and an additional V2I sensor.

## References

- Draft ETSI TS 103 324 V0.0.22 Collective Perception Service.
- Vandana Narri, A. Alanwar, J. Mårtensson, C. Norén, L. Dal Col and K. H. Johansson, "Set-Membership Estimation in Shared Situational Awareness for Automated Vehicles in Occluded Scenarios," 2021 IEEE Intelligent Vehicles Symposium (IV).

## Architecture of Share Situational Awareness

- The proposed architecture is presented in Figure 2. It consists of three parts: (i) Local and extended sensor network, (ii) Algorithms for shared situational awareness, and (iii) Decision-making.
- Measurement data from the sensors are collected and fused to perform state estimation.
- Based on these estimates, decisions are made, and actions are planned. In this paper, the focus is on (i) and (ii).

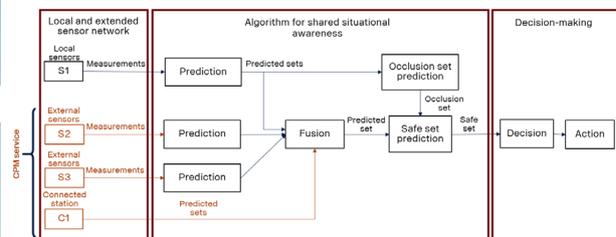


Figure 2 : Proposed architecture for set-based estimation for shared situational awareness.

## Discussion

- In this work, set-based approach is considered, which models the noise and disturbance as unknown variables with known bounds.
- One of the most popular set-based approach is set-membership estimator, which is implemented in this project. And in this approach set of states are considered instead of a single state for estimations which will help in providing robust guarantees and safety margins.
- The set are mathematically represented using zonotopes as shown in Figure 3.

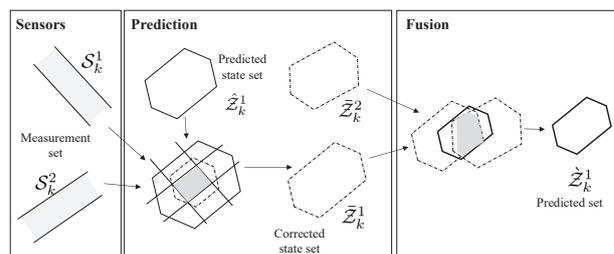


Figure 3 : Illustration of set-membership estimations for shared situational awareness [2].

Nelson, Christian  
Lund University

## A Multilink Channel Measurement System

Wireless channels in vehicular environments are highly dynamic. To evaluate the propagation channel at any given time, they require measurements of all radio channels between several radios simultaneously. Additionally, vehicles and roadside units are distributed at different locations which means the radios cannot be synced using a shared reference source via cables. To this end, a new channel sounder was developed using software-defined radios. This allows each radio to connect to a host computer and a stable rubidium clock. Another source of reference that can be used – when under the open sky – is the pulse-per-second (PPS) transmitted in the GNSS signals. The system has been implemented in the National Instruments software suit LabVIEW.

Nelson, Christian  
Lund University

WASP | WALLENBERG AI,  
AUTONOMOUS SYSTEMS  
AND SOFTWARE PROGRAM

## A Multilink Channel Measurement System

Christian Nelson, PhD Student, Lund University  
Dept. of Electrical Engineering and Information Technology  
Supervisors: Prof. Fredrik Tufvesson



LUND UNIVERSITY

### Motivation & Research Goals

Wireless channels in vehicular environments are highly dynamic. To evaluate the propagation channel at any given time, they require measurements of all radio channels between several radios simultaneously. Additionally, vehicles and roadside units are distributed at different locations which means the radios cannot be synced using a shared reference source via cables. To this end, a new channel sounder was developed using software-defined radios which allows each radio to connect to a host computer and to a stable rubidium clock. Another source of reference that can be used – when there is a clear view of the sky – is the pulse-per-second (PPS) transmitted in the GNSS signals. The system has been implemented in the National Instruments software suite LabVIEW 2021.

### Methods

The radios used have a maximum instantaneous bandwidth of 40 MHz and can be tuned in the range of 1.2 GHz to 6 GHz. For vehicular communication, the frequency of interest is around 5.9 GHz, which is within the operational range. Vehicles and roadside units are distributed in space, so the radios cannot be synced using cables from a shared reference source. Rather, each radio is connected to its host computer for control and a rubidium clock for a stable reference clock. Before each measurement, the rubidium clocks are connected to each other to be synchronized, and can thereafter be separated and maintain coherent for a long enough time to perform needed measurements. If the view of the sky is unobstructed, the internal clocks on the radios can be disciplined using the pulse-per-second (PPS) transmitted in the GPS signal.

<sup>[1]</sup>The sounding technique used is a correlative type, which means that the transmitted signal inhibits good autocorrelation properties. The signal used is the Zadoff-Chu sequence. It has the good autocorrelation properties that we desire and has a flat frequency response. Each radio access the channel in a predefined order using a time-division multiple access (TDMA) scheme.

Additionally, a RF front-end have been designed and built to control the signal paths and to amplify the transmitted signal.

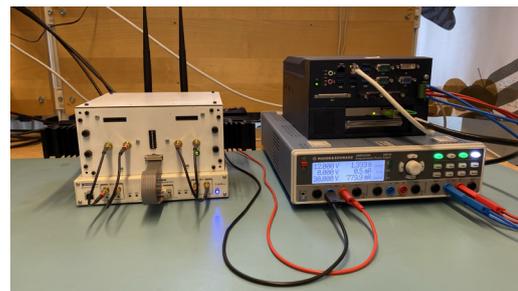
<sup>[2]</sup>The system is implemented using LabVIEW, on a Windows 10 industrial-grade computer. Some of the requirements on the computer are that it needs to be portable for field measurements, and it needs to draw power from car batteries. These requirements limit the performance. The computer connects to the radio via an external PCIe interface which allows for transfer rates up to 200 MSamples/s.

### References

- [1] Wireless Communication, 2nd Edition  
Andreas F. Molisch  
John Wiley & Sons, 2012
- [2] National Instruments LabVIEW  
<https://www.ni.com/labview>  
Accessed 2022-01-06

### Assembled System and Future Work

Below is two figures showing NI LabVIEW 2021 running the host user interface, and the hardware. The system has not yet been tested in the field, but there are measurements planned for spring 2022.



New radios have been acquired, NI-USRP X410 (shown below), with 10x larger instantaneous bandwidth (400 MHz) and 4 RF chains per radio. The work have been initiated to refactor parts of the code to use the new radios. Then radios pose some new challenges regarding data management since it can generate up to 8 GB of data per second.



WASP | WALLENBERG AI,  
AUTONOMOUS SYSTEMS  
AND SOFTWARE PROGRAM

## Multi-Map SLAM

An environment that changes in between visits causes problems for long-term positioning. Can SLAM algorithms handle non-static environments by considering multiple hypotheses of landmark positions?

# Multi-Map SLAM

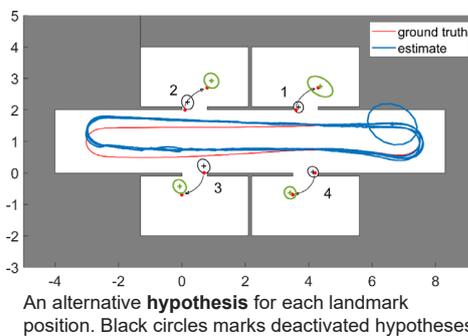
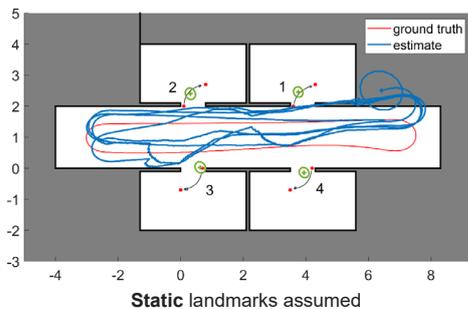
Kristin Nielsen, Linköping University  
Automatic Control  
Main supervisor: Gustaf Hendeby



An environment that changes in between visits causes problems for **long-term positioning**.

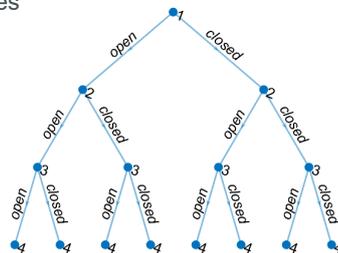
## Can SLAM algorithms handle non-static environments by considering **multiple hypotheses** of landmark positions?

A robot is travelling along a corridor with **unique markers** visible to the robot attached to the door handles. As the robot moves the doors are **opened or closed**, changing how the robot perceives the environment.



A **static world** assumption gives inconsistent estimates of the robot position. By allowing **multiple hypotheses** of landmark positions and statistically decide the most likely hypothesis, the accuracy of the estimate is improved.

Hypotheses



A **hypothesis tree** is maintained where each branch represents one hypothesis.

Multi-Map



In the changing environment of an **underground mine**, different version of the map can be maintained and activated/deactivated online.

Conclusion

A non-static environment that changes in between visits can be represented by a multiple hypothesis map. Decisions and inference in such a map representation has strong connections to theory used in multiple hypothesis tracking.

## Developing Tools and Analyzing Methods for Secure Software Update

Secure Over-The-Air (OTA) software upgrade or update is an important maintenance aspect of any network specifically IoT networks. As a result it is important to figure out which security configuration can affect the security or energy consumption of the devices in the network, and which key sharing and key management scheme, or the actual upgrade procedure is more applicable to the network in case of energy efficiency and security.

To handle these issues, in one of our work, we tried to do an actual OTA update in an IoT environment using CoAP and MQTT protocols to see how security can affect the energy consumption of IoT devices. In another work, we have designed RoSym, a robust, secure and pure symmetric based software upgrade solution for IoT networks. Managing and provisioning of symmetric keys is difficult, as a result in another work, we present Flowrider, a novel key provisioning mechanism for cloud networks that unlocks scalable use of symmetric keys and significantly reduces the related computational load on network endpoints with the use of SDN model. Flowrider makes key distribution agnostic of the network topology and communication patterns, of which it does not require any early knowledge.

Nikbakht Bideh, Pegah  
Lund University

## DEVELOPING TOOLS AND ANALYZING METHODS FOR SECURE SOFTWARE UPDATE

Pegah Nikbakht Bideh, Martin Hell, Nicolae Paladi

Department of Electrical and Information Technology, Lund University, Sweden



### Introduction

Secure Over-The-Air (OTA) software upgrade or update is an important maintenance aspect of any network specifically IoT networks. As a result it is important to figure out which security configuration can affect the security or energy consumption of the devices in the network, and which key sharing and key management scheme, or the actual upgrade procedure is more applicable to the network in case of energy efficiency and security. To handle these issues, in one of our work, we tried to do an actual OTA update in an IoT environment using CoAP and MQTT protocols to see how security can affect the energy consumption of IoT devices. In another work, we have designed RoSym, a robust, secure and pure symmetric based software upgrade solution for IoT networks. Managing and provisioning of symmetric keys is difficult, as a result in another work, we present Flowrider, a novel key provisioning mechanism for cloud networks that unlocks scalable use of symmetric keys and significantly reduces the related computational load on network endpoints with the use of SDN model. Flowrider makes key distribution agnostic of the network topology and communication patterns, of which it does not require any early knowledge.

### Background

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

**OTA upgrade:** In IoT environments, software update and other critical maintenance operations need to be performed over the network and when the device is wireless this often referred to as Over-The-Air (OTA) update/upgrade.

**SDN and OpenFlow:** SDN emerged in response to the increasing complexity of network deployments, facilitating operation and management of virtualized networks. SDN components are:

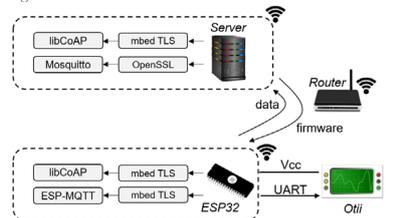
- Data plane
- Control plane
- Network functions

### Energy Consumption for Securing Lightweight IoT Protocols

The ubiquitous nature of IoT devices often requires them to run on batteries, making energy efficiency a primary concern. The large number of devices make it costly to replace batteries, and it will also make the total energy consumption considerable. At the same time, adding security to the communication will add additional overhead. Thus, it is important to not only develop lightweight security protocols, but also to understand to which extent security affects the energy consumption of the devices. The main contribution of our work is a thorough analysis of CoAP and MQTT and the investigation of their energy footprint in different scenarios:

- How added security at the transport layer (TLS/DTLS) affects the energy consumption?
- How important design choices, such as cipher suite, PKI vs. PSK, and client authentication impact the energy consumption?

In our experiments we use ESP32 with libcoap, MQTT, and mbed TLS libraries and conduct real-world measurements using Oti.



We find that aggregating data to larger packets can significantly reduce the energy consumption. We also find that AES-CCMS seems slightly more efficient than other modes of operation.

### RoSym: Robust Symmetric Key Based IoT Software Upgrade Over-the-Air

Despite the fact that the software upgrade problem have been extensively studied, there is still a need for investigating new design of completely symmetric key-based upgrade solutions with post-quantum resistant. Our new scheme RoSym is such a solution. Furthermore, the upgrade process should not be sensitive to Denial-of-Service (DoS) attacks. In particular, it should not be easy for an attacker to force IoT units to waste extensive power battery resources on false software upgrade requests. These properties should be achieved while simultaneously offering a high level of confidentiality and integrity protection, working well for resource constraint units, and avoiding public key operations during software upgrade.

Our solution is consists of two phases:

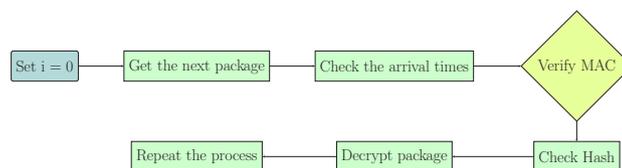
**Key Establishment and Parameter Setup:** During this phase the following information will be transferred to IoT devices:

1. Two randomly generated symmetric keys,  $IK_{sw}$  and  $K_{sw}$ .
2. Timing information,  $T_1$  and  $T_2$  determines the validity of update packets and validity of  $IK_{sw}$  and  $K_{sw}$  keys respectively.
3. A software image one-way hash root value,  $h_0$ .
4. The number of packages ( $n$ ) in the new software distribution.

**Upgrade procedure phase:** Before the upgrade procedure starts, the DMS package the software update image into suitable size packages.

```
File name { index: 4B | Enc software: 944B | hash: 32B | MAC: 32B }
index: Package index,  $i$ 
Enc software: Encrypted software block:  $E(K_{sw}, I_i)$ 
hash $_i$ : One-way hash,  $h_i = H(\{i, E(K_{sw}, I_i), h_{i-1}\})$ ,  $0 \leq i \leq n-2$ 
MAC $_i$ : Message Authentication Code,  $MAC_i = MAC(TK_{sw}, \{i, E(K_{sw}, I_i), h_{i+1}\})$ ,  $0 \leq i \leq n-2$ 
```

The following corresponding upgrade procedure applies on the IoT side:



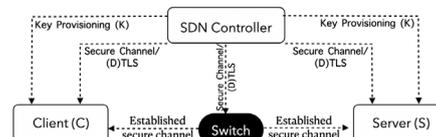
Our security analysis showed that RoSym offers the expected security level.

### Flowrider: Fast On-Demand Key Provisioning for Cloud Networks

Generating symmetric keys requires less computational power and has firmware support on many platforms, the use of symmetric keys leads to challenges such as secure key provisioning and key authentication. This introduces the research question:

**Can the SDN model be leveraged to conveniently provision symmetric keys and reduce computational resource consumption?**

Yes with the use of Flowrider, a novel key provisioning mechanism for network endpoints in SDN deployments that considers the practicalities of cloud systems deployment.



## Planning for minimal uncertainty

A belief-space planning problem for GNSS-denied areas is studied where the location and number of landmarks available are unknown when performing the planning.

To be able to plan an informative path in this situation, an algorithm using virtual landmarks to position the platform during the planning phase is studied.

The virtual landmarks are selected to capture the expected information available in different regions of the map, based on the beforehand known landmark density.

The approach is tested in a simulated environment, in conjunction with an extended information filter, with successful results.

# Planning for minimal uncertainty

Jonas Nordlöf, Gustaf Hendeby and Daniel Axehill

## Summary

- Objective:** Reduce position uncertainty in landmark-based SLAM by using motion planning.
- Challenge:** Landmark positions are unknown.
- Solution:** Introduce approximation using virtual landmarks based on landmark densities.
- Outcome:** Ability to plan minimal uncertainty path and predict position uncertainty.

## Stochastic optimization problem

Future position estimates depend on the actual motion and the obtained measurements of the landmarks positions, which are both unavailable at the planning stage. This leads to a stochastic optimization problem:

$$\begin{aligned}
 & \underset{\pi, T}{\text{minimize}} && \mathbb{E}(J(\mathcal{I}_{|t})) && \text{expected performance} \\
 & \text{subject to} && x_{t+1} = f(x_t, u_t, w_t), && \text{platform state dynamics} \\
 & && y_i^j = h^j(x_t) + \epsilon_i^j, \forall i \in \mathcal{M}_t, && \text{landmark measurements} \\
 & && \mathcal{I}_{t+1|t+1} = \Lambda(\mathcal{I}_{t|t}, y_t, x_t, u_t), && \text{information model} \\
 & && u_t = \pi(x_t, \mathcal{I}_{t|t}), && \text{control policy} \\
 & && \epsilon_i^j \sim \mathcal{N}(0, R_t), \quad w_t \sim \mathcal{N}(0, Q_t)
 \end{aligned}$$

This problem is generally not solvable.

## Deterministic approximation

The stochastic problem can be replaced by a deterministic problem using the following approximations:

**Unknown noise realizations:** Certainty-equivalent control

**Unknown landmark positions:** Virtual landmarks based on landmark densities  $\rho(\cdot)$

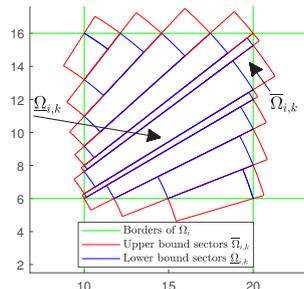
- Each virtual landmark represents a subregion  $\Omega_i$ .
- Information gained from observing a subregion  $\Omega_i$  can be calculated as

$$\mathcal{I}_i^j = \int_{\tilde{m} \in \Omega_i} \rho(\tilde{m}) (H_t(p_t, \tilde{m}))^\top R_t^{-1} H_t(p_t, \tilde{m}) d\tilde{m}. \quad (1)$$

yielding the information update:

$$\mathcal{I}_{t+1|t+1} = (F_t \mathcal{I}_{t|t}^{-1} F_t^\top + G_t Q_t G_t^\top)^{-1} + \sum_{i \in \mathcal{M}_t} \mathcal{I}_i^j \quad (2)$$

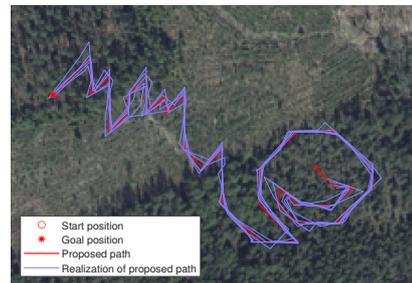
- For range-bearing measurements (1) has closed form solution for circle sectors.
- Approximate  $\Omega_i$  with circle sectors  $\Omega_{i,k}$ .



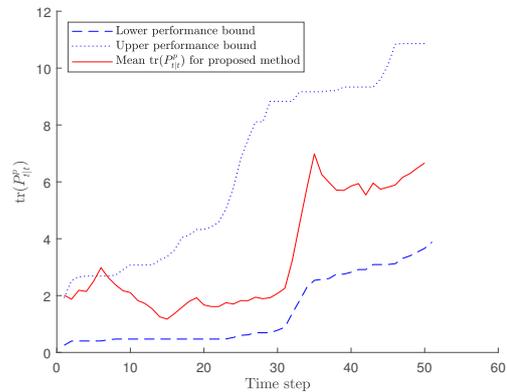
## Simulation study

The approach is evaluated using Monte Carlo simulations in a forest environment. Trees are used as landmarks.

- Planned path using the proposed approach (red) and Monte Carlo realizations of path (blue)



- Estimated performance bounds (blue) and Monte Carlo estimate of performance measure (red)



## Future work

- Apply the method in a real environment
- Investigate impact of terrain properties on position estimate
- Add known landmarks and visual sensors
- Position correction using receding horizon control

## References

J. Nordlöf, G. Hendeby, and D. Axehill, "Belief space planning using landmark density information", in 2020 Proc. IEEE 23rd Int. Conf. Inf. Fusion (FUSION), Rustenburg, South Africa, Jul. 2020, pp. 1–8.

J. Nordlöf, G. Hendeby, and D. Axehill, "Improved Virtual Landmark Approximation for Belief-Space Planning", in 2021 Proc. IEEE 24rd Int. Conf. Inf. Fusion (FUSION), Rustenburg, South Africa, Nov. 2021, pp. 1–8.

## Acknowledgment & Collaborations

This work is partially supported by the WASP Affiliated PhD Student Program. The project is being funded by the division of C4ISR within the Swedish Defence Research Agency (FOI). The project follows the Swedish Armed Forces Research and Tech program for the area of autonomous systems and the area of sensors and signature management.

## Foresee the Unseen: Sequential Reasoning about Hidden Obstacles for Safe Driving

Safe driving requires autonomous vehicles to anticipate unseen objects, such as a cyclist hidden behind a large vehicle, or an object on the road hidden behind a building.

Existing methods are usually not able to consider all possible shapes and orientations of such obstacles.

They also typically do not reason about observations of hidden obstacles over time, leading to conservative anticipations.

We overcome these limitations by (1) modeling possible hidden obstacles as a set of states of a point mass model and (2) sequential reasoning based on reachability analysis and previous observations. Based on (1), our method is safer, since we anticipate obstacles of arbitrary unknown shapes and orientations.

In addition, (2) increases the available drivable space when planning trajectories for autonomous vehicles, which we demonstrate can give rise to significant reductions in time when traversing various intersection scenarios.



# Foresee the Unseen



Truls Nyberg\*, José Manuel Gaspar Sanchez\*, Christian Pek, Jana Tumova, Martin Törngren  
Affiliated with KTH Royal Institute of Technology and Scania CV AB

\*Contributed equally to this research. Article currently under review.

## Sequential Reasoning about Hidden Obstacles for Safe Driving

Safe driving requires autonomous vehicles to anticipate **potential hidden traffic participants and objects**. Existing methods typically do not consider **arbitrary shapes** of hidden obstacles and do not reason about **observations over time**. We overcome these limitations by (1) **modeling possible hidden obstacles as a set of states** of a point mass model and (2) **sequential reasoning based on reachability analysis** and previous observations.

### The Problem

Autonomous vehicles need to model possible hidden obstacles **conservatively enough**, such that any possible unseen obstacle is represented and considered, regardless of their size or orientation, such as the motorcycle in **Figure 1**.

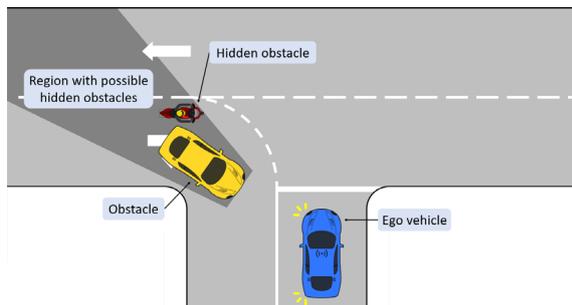


Figure 1. A dangerous situation where defensive driving is needed.

However, modeling possible hidden obstacles **too conservatively** limit autonomous vehicles from finding safe and efficient paths.

Given **past observations** and assumed constraints on driving (e.g., maximum speed and other traffic rules), currently unseen regions can still be concluded free from obstacles, such as the checkered region in **Figure 2**.

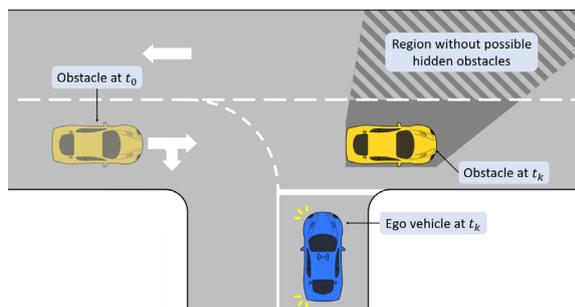


Figure 2. A typical situation where autonomous vehicles generally are too conservative. The checkered region was previously seen, and it can be concluded that no object could have reached there.

### The Solution

By **initially** considering the complete unseen region as potentially occupied, our method captures **any** hidden obstacle (**Figure 3a**). By **iteratively** updating which regions possibly can be occupied, we avoid being **too conservative**.

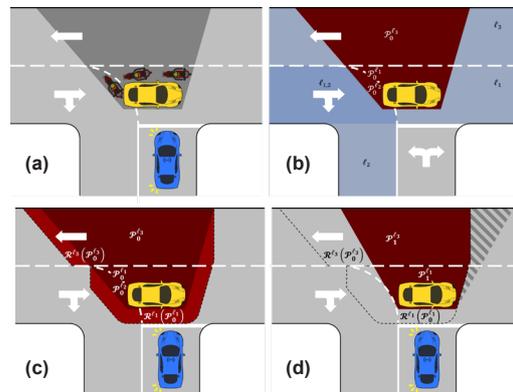


Figure 3. Algorithmic steps for reasoning.

For each lane (**Figure 3b**), the reachability is computed for the possible hidden obstacles (bright red in **Figure 3c**). New unseen regions are deemed free if they cannot have been reached since the last observation (the checkered region in **Figure 3d**). The result can be seen in **Figure 4**, where the time to traverse the intersection is **greatly reduced** by reasoning about possible hidden obstacles over time.

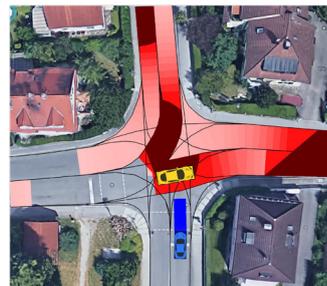


Figure 4. Simulation in CommonRoad of intersection in Fürstenfeldbruck, using the proposed algorithm.

## Towards data-driven Storage Location Assignment: New benchmarks and optimization model

The Storage Location Assignment Problem (SLAP) is concerned with the choice of locations for products in a warehouse. It is of primary significance for operational quality since the travel cost of order-picking vehicles is strongly related to where and how far they have to travel. Unfortunately, a generalized model of the SLAP poses a highly intractable problem. State-of-the-art optimization methods tend to be usecase-specific and there exists no standard benchmark dataset format. In this study we introduce new benchmark data on a modified TSPLIB format and demonstrate how instances can be approximately solved using a state-of-the-art Order Batching Problem (OBP) optimizer aided by a Quadratic Assignment Problem (QAP) surrogate. Our results show that the OBP optimizer can yield significantly better performance when it is aided by the surrogate.

# Towards data-driven storage assignment: New benchmarks and optimization model



Johan Oxenstierna, Jacek Malec, Volker Krueger  
Computer Science Dept., Lund University

The Storage Location Assignment Problem (SLAP) is concerned with the choice of locations for products in a warehouse. It is of primary significance for operational quality since the travel cost of order-picking vehicles is strongly related to where and how far they have to travel. Unfortunately, a generalized model of the SLAP poses a highly intractable problem. State-of-the-art optimization methods tend to be usecase-specific and there exists no standard benchmark dataset format. In this study we introduce new benchmark data on a modified TSPLIB format and demonstrate how instances can be approximately solved using a state-of-the-art Order Batching Problem (OBP) optimizer aided by a Quadratic Assignment Problem (QAP) surrogate. Our results show that the OBP optimizer can yield significantly better performance when it is aided by the surrogate.

## Contributions

1. Introduction of SLAP benchmark data on the TSPLIB format.
2. Evaluation of the usage of QAP surrogates within SLAP optimization.

## Problem formulation

The general SLAP objective is to assign locations for products such that the distance to pick the products, using order batching, is minimized [1]:

$$\min \sum_{b \in \mathcal{B}} D(b)x_{mb}, m \in \mathcal{M}$$

where  $\mathcal{B}$  denotes generated batches, where  $D(b)$  denotes the distance to pick batch  $b$  (a solution to a Traveling Salesman Problem where the node-location pointers are mutable), where  $m$  denotes a vehicle and where  $x_{mb}$  denotes a binary variable that is 1 if vehicle  $m$  is assigned to pick  $b$  and 0 otherwise.

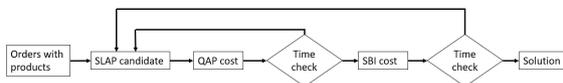
An optimizer would require significant computational time to find a value close to the minimum above since the problem is NP-hard. For a faster but less precise quality evaluation of a product to location assignment, a Quadratic Assignment Problem (QAP) surrogate is used [2]. The QAP uses distances and weights between products. The weights between two products is here defined as the number of times two products appear in the same order. The QAP function solution value is computed by:

$$\sum_{p_1 \in \mathcal{P}} \sum_{p_2 \in \mathcal{P}} \sum_{l_1 \in \mathcal{L}} \sum_{l_2 \in \mathcal{L}} w_{p_1 p_2} d_{l_1 l_2} x(p_1, l_1) x(p_2, l_2)$$

where  $w$  denotes weight, where  $d$  denotes distance and  $x(p, l)$  a function which returns 1 if product  $p$  is located at location  $l$  and 0 otherwise.

## Optimization model

An order batching optimizer (Single Batch Iterated) with a Quadratic Assignment Surrogate (SBI-QAS):



This model is delimited in two major ways: 1. No bootstrapping is used to generate SLAP candidates. 2. The QAP stopping criterion is time-based rather than convergence based. These are motivated since they are bias-imposing techniques [3] and they should only be implemented after the accuracy of the surrogate can be established through experimentation.

## References

1. Gademann and Velde: Order batching to minimize total travel time in a parallel-aisle warehouse, IIE Transactions, 2005.
2. Kubler, Glock and Bauernhansl: A new iterative method for solving the joint dynamic storage location assignment, order batching and picker routing problem in manual picker-to-parts warehouses, Computers & Industrial Engineering, 2021.
3. Sutton and Barto: Reinforcement Learning: An introduction, MIT Press, 2018.
4. Busa-Fekete, Róbert, et al. "An apple-to-apple comparison of learning-to-rank algorithms in terms of normalized discounted cumulative gain." *ECAI 2012-20<sup>th</sup>, 2012*.
5. Oxenstierna et al: [https://github.com/johanoxenstierna/OBP\\_instances](https://github.com/johanoxenstierna/OBP_instances), [https://github.com/johanoxenstierna/L09\\_251](https://github.com/johanoxenstierna/L09_251), 2021.

## Surrogate evaluation metric

As shown in the above diagram, the slower OBP optimizer is run on the most promising set of candidate solution(s), whose selection is determined by ordering the QAP surrogate prediction values. Hence, to evaluate the surrogate quality, we are not interested in the QAP estimates themselves, but rather the surrogate's ability to correctly rank those estimates. We therefore use the Normalized Discounted Cumulative Gain [4]:

$$NDCG(k) = 1 - \frac{DCG(k)}{IDCG(k)}, \quad DCG(k) = \sum_{i=1}^k \frac{G_i}{\log_2(i+1)}, \quad IDCG(k) = \sum_{i=1}^{|I(k)|} \frac{G_i}{\log_2(i+1)}$$

where  $k$  denotes an index and  $G$  denotes gain. NDCG compares the ranking quality of a candidate ranker against a ground truth ranker.

## Experiment

The aim of the experiment is to empirically test the QAP surrogate predictive strength against a ground truth ranker (SBI).

## Benchmark data and format

The public TSPLIB format datasets in L6\_203 and L09\_251 are modified for the SLAP [5]. For analysis, the instances are divided into 7 classes according to number of orders in the test data (10 – 1000). The experiment is set up such that 60 randomly generated SLAP candidate solutions are generated for each instance. For each of these SLAP candidates the QAP and OBP estimates and CPU-times are recorded. In total there are 2274 candidates and corresponding predictions for the QAP and OBP modules.

## Experiment result

On average, the QAP surrogate predictive error ranges from 35 - 44% (the data is normalized such that anything below 50% means the proposed algorithm is successful), with standard deviations in the range 12-16%. A slight pattern for larger predictive errors for larger instance sizes can be observed, but this is inconclusive.

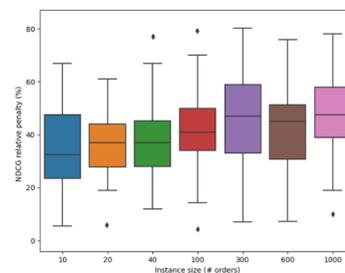


Figure 1: Box-plot showing instance sizes on the x-axis (in terms of number of orders) and NDCG on the y-axis (how wrong the QAP ranking is on average). The inner and outer boxes represent 95% and 99% confidence intervals respectively.

This result provides evidence that the QAP surrogate can be successfully utilized within the larger algorithm. For future work the distribution of work between the QAP surrogate and the SBI optimizer could be investigated. It was found that the QAP surrogate requires around 60x less CPU-time than SBI, so an experiment could be set up with various number of surrogate solution candidates.

## Optimal Input Design Through Infinity Norm Minimization Using Proximal Mapping

To avoid non-convexity of the criterion, various relaxations are typically used in input design. For example, the input may be assumed to be stationary and the design problem may be formulated in terms of the correlation coefficients. Now, we instead propose a method to directly design the input sequence. This allows to maximize the information obtained from short-time (transient) experiments using non-stationary inputs. We do this by fitting the achieved Fisher matrix to a desired target matrix in a matrix sense, using the infinity norm. The target matrix can either be the desired Fisher matrix, obtained from quality considerations of the intended use of the model, or a matrix directly representing the performance of the application. An often used quantity is the Hessian of the so called the application cost. Thus, the method is formulated as a time domain optimization problem that is non-convex. This optimization problem is solved by alternative minimization and proximal mapping.

# Optimal Input Design through Infinity Norm Minimization

Javad Parsa, PhD student, KTH University  
Division of Decision and Control Systems  
Supervisor: Prof. Håkan Hjalmarsson



## Abstract

To avoid non-convexity of the criterion, various relaxations are typically used in input design. For example, the input may be assumed to be stationary and the design problem may be formulated in terms of the correlation coefficients. In this contribution, we instead propose a method to directly design the input sequence. This allows to maximize the information obtained from short-time (transient) experiments using non-stationary inputs. We do this by fitting the achieved Fisher matrix to a desired target matrix in a matrix sense, using the infinity norm. The target matrix can either be the desired Fisher matrix, obtained from quality considerations of the intended use of the model, or a matrix directly representing the performance of the application. An often used quantity is the Hessian of the so called the application cost. Thus, the method is formulated as a time domain optimization problem that is non-convex. This optimization problem is solved by alternating minimization and proximal mapping.

## Methods

Linear regression model:

$$\mathbf{y} = \Phi\theta + \mathbf{e}, \quad \mathbf{e} \sim \mathcal{N}(0, \lambda\mathbf{I})$$

where  $\theta \in \mathbb{R}^{n_\theta \times 1}$ ,  $\Phi \in \mathbb{R}^{N \times n_\theta}$  and  $\mathbf{e} \in \mathbb{R}^{N \times 1}$ .

Fisher Information Matrix (FIM):

$$\mathbf{I}_F = \frac{1}{\lambda} \Phi^T \Phi.$$

Identification set:

$$\varepsilon_{si} = \{\theta : (\theta - \theta_0)^T \mathbf{I}_F (\theta - \theta_0) \leq \chi_\alpha^2(n_\theta)\}$$

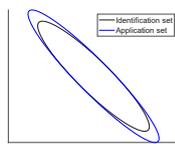
where  $\chi_\alpha^2(n_\theta)$  is the  $\alpha$ -percentile of the chi-square distribution with  $n_\theta$  degrees of freedom and  $\theta_0$  denotes true parameter vector.

Application performance cost  $V_{app}$ : measures the performance degradation when the model parameter  $\theta$  is used in a model based design.

$$\varepsilon_{app} = \{\theta : (\theta - \theta_0)^T V_{app}''(\theta_0) (\theta - \theta_0) \leq \frac{1}{\gamma}\}.$$

To guarantee that the estimates are inside the application set with high probability, we need to ensure the identification set is inside the application set [1], i.e.:

$$\mathbf{I}_F = \frac{1}{\lambda} \Phi^T \Phi \geq \gamma \chi_\alpha^2(n_\theta) V_{app}''(\theta_0).$$



State-of-the-art: Per sample Fisher matrix (means that the input is stationary) designed using convex optimization. Results in a desired Fisher matrix  $\mathbf{I}_F^d$ . Will be denoted the Frequency Domain Method (T-FDM).

## Proposed method:

To find the optimal regressor, the following optimization problem is proposed:

$$\min_{\Phi \in \mathcal{D}} \|\Phi^T \Phi - \mathbf{T}\|_\infty.$$

in which the  $\mathcal{D}$  denotes the set of  $N \times n_\theta$  Toeplitz matrices and the infinity norm of a matrix means maximum of the absolute values of the elements of this matrix. To solve the above optimization problem, we use proximal mapping and a closed form solution to update the regressor  $\Phi$  can be found [2].

• Target matrix:

- 1)  $\mathbf{T} = \lambda \mathbf{I}_F^d$  (T-FDM)
- 2)  $\mathbf{T} = \lambda \chi_\alpha^2(n_\theta) \gamma V_{app}''(\theta_0)$  - the scaled Hessian of Application cost (T-APP).

## Selected Results

### Feedforward control problem:

State Space Model:

$$\begin{bmatrix} x_1(t+1) \\ x_2(t+1) \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} u(t)$$

$$y(t) = \begin{bmatrix} \theta_1 & \theta_2 \end{bmatrix} \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix} + d(t), \quad d(t) = \begin{cases} 1, & t \geq 0 \\ 0, & t < 0 \end{cases}$$

Feedforward controller:

$$u(t) = \frac{-d(t)}{\theta_1 + \theta_2}$$

Application Cost:

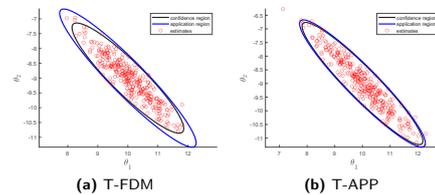
$$V_{app}(\theta) = \frac{1}{B} \sum_{t=1}^B (y(t, \theta) - y(t, \theta_0))^2.$$

Optimal input design:

$$\min_{\phi_u(\omega)} \left( \frac{1}{2\pi} \int_{-\pi}^{\pi} \phi_u(\omega) d\omega \right)$$

$$\text{st } \mathbf{I}_F \geq \chi_\alpha^2(n_\theta) \gamma V_{app}''(\theta_0),$$

The outcome of above optimization problem (using the MOOSE2 toolbox) is the desired per sample matrix  $\mathbf{I}_F^d$ . Then, we take  $\mathbf{T} = \mathbf{I}_F^d$  and T-FDM (Target matrix using per sample FIM from FDM) denotes this case.



From these figures we see that approximately 99% and 98% of the estimates are inside the application cost for T-FDM and T-APP, respectively.

## References

- [1] System identification of complex and structured systems H. Hjalmarsson, European Journal of Control, 2009.
- [2] Optimal Input Design Through Infinity Norm Minimization Using Proximal Mapping, J. Parsa and H. Hjalmarsson, 60th IEEE Conference on Decision and Control conference, CDC 2019.

Peng, Haorui  
Lund University

## Mission-critical Applications at the Edge of 5G

In this work, we performed the evaluation of real-time HTTP applications that can be deployed at the edge of a complete mid-band stand alone 5G base station. We showed the advantages and impacts of a real 5G network system on mission-critical processes, which are envisioned to benefit from the ultra-low latency and extreme-high bandwidth of the 5G.

Peng, Haorui  
Lund University



## Value-Maximizing Combinatorial Assignment

We investigate value-maximizing combinatorial assignment—i.e., the problem of partitioning items into bundles among a set of alternatives to maximize some notion of social welfare. This problem is a major research challenge in computer science with many applications in for example operations research, economics, and artificial intelligence. Unfortunately, combinatorial assignment problems are in general both NP-hard and inapproximable.

# Value-Maximizing Combinatorial Assignment

Fredrik Präntare (fredrik.prantare@liu.se)

Department of Computer Science (AIICS/RealL), Linköping University

## Background & Motivation

We investigate *value-maximizing combinatorial assignment*—i.e., the problem of partitioning items into bundles among a set of alternatives to maximize some notion of social welfare. This problem is a major research challenge in computer science with many applications in for example operations research, economics, and artificial intelligence. Unfortunately, combinatorial assignment problems are in general both NP-hard and inapproximable.

## Contributions

Our contributions include developing the state-of-the-art algorithms for combinatorial assignment, as well as theoretical and empirical advances. Our algorithms have also been applied to (and are being used in) a commercial real-world application with more than 1 million users (Fig. 1).



**Fig. 1:** Our algorithms are used in the commercial strategy game *Europa Universalis 4* (a game with more than 1 million players) to coordinate and deploy armies to different regions.

## Publications

[1] Präntare, F., Ragnemalm, I., & Heintz, F. (2017).

**Lilla Polhemspriset**

**An Algorithm for Simultaneous Coalition Structure Generation and Task Assignment.** In *International Conference on Principles and Practice of Multi-Agent Systems*.

[2] Präntare, F., & Heintz, F. (2018).

**Best Student Paper Award**

**An Anytime Algorithm for Simultaneous Coalition Structure Generation and Assignment.** In *International Conference on Principles and Practice of Multi-Agent Systems*.

[3] Präntare, F., & Heintz, F. (2019).

**An Anytime Algorithm for Optimal Simultaneous Coalition Structure Generation and Assignment.** In *Journal of Autonomous Agents and Multi-agent Systems*.

[4] Präntare, F., & Heintz, F. (2020).

**Best Student Paper Award**

**Dynamic Programming for Optimal Simultaneous Coalition Structure Generation and Assignment.** In *International Conference on Principles and Practice of Multi-Agent Systems*.

[5] Präntare, F., Appelgren, H., & Heintz, F. (2021).

**Anytime Heuristic and Monte Carlo Methods for Large-Scale Simultaneous Coalition Structure Generation and Assignment**

In *AAAI Conference on Artificial Intelligence*.

[6] Präntare, F., Tiger, M., Bergström, D., Appelgren, H., & Heintz, F. (2022).

**Learning Heuristics for Combinatorial Assignment Problems by Optimally Solving Subproblems.**

In *International Conference on Autonomous Agents and Multi-Agent Systems*.

[7] Präntare, F., Osipov, G., Eriksson, L. (2022).

**Concise Representations and Complexity of Combinatorial Assignment Problems.**

To be published.

## Power Efficient Multi DNN Accelerator for Future IoT Devices

Future IoT devices and autonomous systems would have widespread use of deep learning. It is increasingly important to do those heavy computations on the device itself without offloading them to the cloud. Those devices often use specialized domain-specific DNN accelerators to achieve high computation efficiency. However, those devices are battery-powered and power and area constrained. On the other hand, future applications would require simultaneous execution of multiple network models on the same device concurrently. Further due to being chip area constrained those devices could not be provisioned with large compute and memory resources. Limited on-chip memory leads to frequent off-chip memory accesses which dominates the power consumption in such devices and multi DNN acceleration aggravates this problem where the on-chip memory should be shared between multiple DNNs. Therefore this research aims at exploring architectural techniques for achieving better power efficiency for multi DNN acceleration such as stream caching and efficient memory management for such accelerators which could potentially reduce the number of off-chip accesses.

Ranawaka, Piyumal  
Chalmers



## Topological Method for fMRI Analysis

Functional magnetic resonance imaging (fMRI) is used to measure brain activity due to tasks or stimuli. Resting-state measurements are used to provide a subject's baseline. The signal is prone to noise from various sources. Random brain activity and noise, from the scanner, can reach a strength comparable to the signal itself. Thus, extracting the underlying signal is a challenging process typically approached by applying statistical methods. The goal of this study is to investigate possibilities to recover information from the signal using topological feature vectors directly based on the raw signal without any medical pre-knowledge. The goal is to recover, the temporal development of brain activations, connectivity between these activations, and their relation to these cognitive tasks

# Topological Method for fMRI Analysis

Farhan Rasheed, Linköping University

Department of Science and Technology

Main advisor: Ingrid Hotz



## Motivation & Research goals

Functional magnetic resonance imaging (fMRI) is used to measure brain activity due to tasks or stimuli. Resting-state measurements are used to provide a subject's baseline. The signal is prone to noise from various sources. Random brain activity and noise, from the scanner, can reach a strength comparable to the signal itself. Thus, extracting the underlying signal is a challenging process typically approached by applying statistical methods. The goal of this study is to investigate possibilities to recover information from the signal using topological feature vectors directly based on the raw signal without any medical pre-knowledge. The goal is to recover, the temporal development of brain activations, connectivity between these activations, and their relation to these cognitive tasks

## Data and Methods

### Functional MRI

fMRI measures BOLD signals that are coupled with neural activation. Participants were scanned while solving 4 different task in two 3 min block for each task[1].

Session time = 25 min 24 sec  
Each Block = 3 min

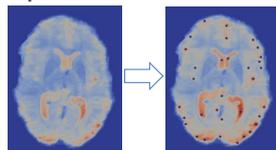
4D (3D + time)

### Topological feature vector – requirements:

- 1) represents the main characteristics of brain activity, 2) acquire directly from raw data, 3) reduce dimensionality significantly 4) has a fixed size so it is comparable across time steps.

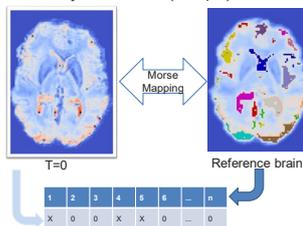
### Proposed feature vector per time step

- Activity vector derived from the merge tree[2] (Req 1,2)
- Use persistence simplification as controlled tree simplification (Req 3)

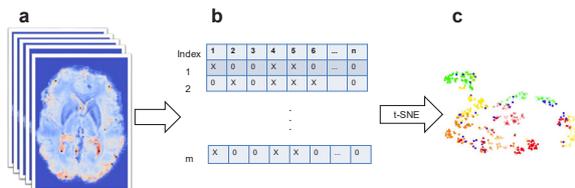


### Reference brain (RF) – Feature correspondence (Req 4)

- RF combines all activity sites of one subject from all time-steps
- RF introduces a **global index set** that allows
- Morse-mapping[3] is used to establish temporal correspondence.

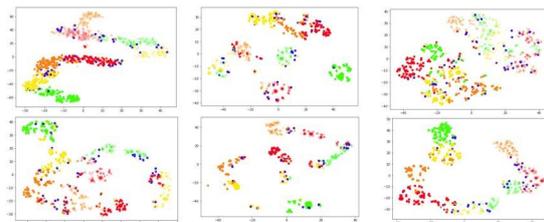


## Selected Results



**Pipeline:** (a) raw data (m time steps) (b) Feature matrix representing the activity levels at the n dominant activity sites (c) 2d embedding, color coding according to the tasks, blue refers to instruction times.

Feature vectors mapping into two dimensions space reveals the different activity tasks and their transition

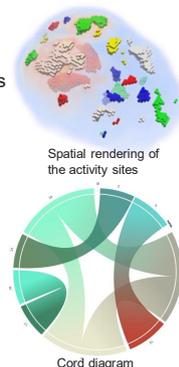


**Fig:** Six subjects in comparison, it can be observed that the feature vector coherency is differently strongly expressed for the subjects. For subject 1 one can observe an almost continuous change of the brain activity over time. In contrast subject 4 shows clear clusters per task.

### Activation sites connectivity

For a brain connectivity analysis we focus on the correlation of individual activation sites and relate their behavior to each other. This includes analyzing the overall activation level.

Chord diagram visualize the relation between the activation sites (indicated by index). The ribbon connects regions with correlation value higher than 60%



## References

1. Javier Gonzalez-Castillo et al, Tracking ongoing cognition in individuals using brief, whole-brain functional connectivity patterns, *Proceedings of the National Academy of Sciences* (2015)
2. Carr et al, Computing contour trees in all dimensions, *Computational Geometry* (2003)
3. W. Engelke et al, Topology-Based Feature Design and Tracking for Multi-Center Cyclones, *Topological Methods in Data Analysis and Visualization VI*, Eds. Springer, (2021)

Rodriguez-Deniz, Hector  
Linköping University

## Robust Real-Time Delay Predictions in a Network of Urban Buses

Providing transport users and operators with accurate forecasts on travel times is challenging due to a highly stochastic traffic environment. In this paper we develop a robust model for real-time bus travel time prediction that depart from Gaussian assumptions by using Student-t errors. The proposed approach uses spatiotemporal characteristics from the route and previous bus trips to model short-term effects, and date/time variables and Gaussian processes for long-run forecasts. The model allows for flexible modeling of mean, variance and kurtosis spaces. Experiments are performed using data from high-frequency buses in Stockholm, Sweden.

Rodriguez-Deniz, Hector  
Linköping University



# Robust Real-Time Delay Predictions in a Network of Urban Buses

Hector Rodriguez-Deniz<sup>1</sup> and Mattias Villani<sup>1,2</sup>

<sup>1</sup>Linköping University - Statistics and Machine Learning (STIMA)

<sup>2</sup>Stockholm University - Statistics



## DESCRIPTION

Providing transport users and operators with accurate forecasts on travel times is challenging due to a highly stochastic traffic environment. In this paper we develop a robust model for real-time bus travel time prediction that depart from Gaussian assumptions by using Student- $t$  errors. The proposed approach uses spatiotemporal characteristics from the route and previous bus trips to model short-term effects, and date/time variables and Gaussian processes for long-run forecasts. The model allows for flexible modeling of mean, variance and kurtosis spaces. Experiments are performed using data from high-frequency buses in Stockholm, Sweden.



## METHODOLOGY

The conditional distribution of the delay at time  $t$  for a bus arriving to stop  $j$ ,  $y_j(t)$ , given a feature vector  $\mathbf{x}_j(t)$ , is modeled as

$$y_j(t)|\mathbf{x}_j(t) \sim \mathcal{T}[\mu(\mathbf{x}_j(t)), \sigma^2(\mathbf{x}_j(t)), \nu(\mathbf{x}_j(t))],$$

where  $\mathcal{T}(\mu, \sigma^2, \nu)$  is the Student- $t$  distribution with degrees of freedom  $\nu > 0$ , location  $\mu$  and scale  $\sigma^2 > 0$ . The Student- $t$  model parameters  $\mu$ ,  $\sigma^2$  and  $\nu$  are all functions of the time indexed feature vector  $\mathbf{x}_j(t)$ , which contains functions of time and time-discounted past delays in the network to capture the dynamics in time and space. For each of  $\mu$ ,  $\sigma^2$  and  $\nu$  we use a short-term and a long-term component, e.g. the mean is modeled by

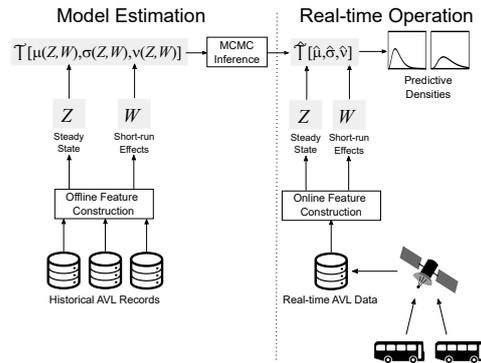
$$\mu(\mathbf{x}_j(t)) = f_j(t) + g_j(t) \tag{1}$$

with long-run part

$$f_j(t) = \mathbf{z}_j(t)^\top \boldsymbol{\alpha}_j \tag{2}$$

and short-run effects modeled by  $g_j(t)$ . Importantly, the short-run effects will be modeled so that  $g_j(t) \rightarrow 0$  when forecasting at longer horizons. The term  $f_j(t)$  will therefore be referred to as the steady-state and will be the determinant of forecasts at longer horizons, e.g. predictions much later in the day or even several days ahead for route planning. The steady-state feature vector  $\mathbf{z}_j(t)$  models additive effects for the within-day congestion peaks (hourly), within-week differences (day of the week) and seasonal patterns.

We model some of these effects non-parametrically using GP priors on the parameters  $\boldsymbol{\alpha}_j$  to encode smoothness over e.g. the hours of the day. We account for heteroskedasticity by modeling the logarithm of the scale and degrees of freedom as linear regressions in a similar fashion as for the mean.

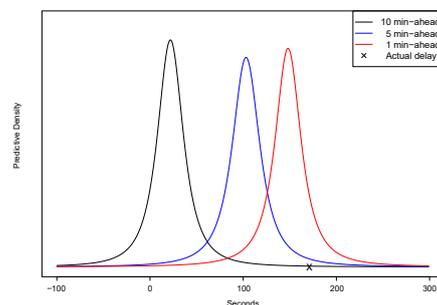
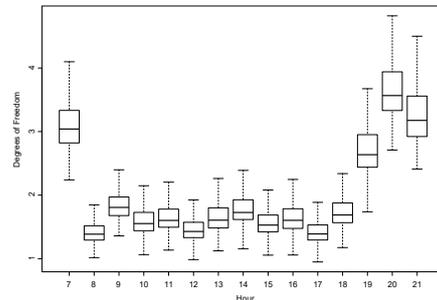


## RESULTS

Results show that Student- $t$  models outperform Gaussian ones in terms of log-posterior predictive density to forecast bus delays at specific stops. Even though there is little variation in mean predictions, which are below 1 minute in most cases, our results illustrate the importance of accounting for predictive uncertainty in model selection. Importantly, the flexibility induced by the inclusion of covariates on the scale and degrees-of-freedom spaces has clearly paid off in terms of out-of-sample predictive performance. Estimates from linear regressions in the most elaborate model capture various patterns in the delay distribution, such as rush/non-rush hours and weekday/weekend peaks, etc. For mean delay prediction, the stronger spatiotemporal effects are relative to incoming buses from immediately previous stops, which is line many recently developed models. The estimated regression on the degrees-of-freedom reveals a higher probability of extreme observations during working hours. Bayesian inference naturally allows for generating predictive intervals, and calculating the probability of e.g. incoming buses coming late, by using the posterior predictive distributions.

The results are promising, but extensive simulations must be carried out with larger databases to obtain more consistent evidence, investigate wider spatiotemporal network effects and contrast results between high and low-frequency services. Variational inference could be also an alternative to MCMC, in order to scale up to massive datasets.

Model	LPPD <sub>train</sub>	LPPD <sub>test</sub>	MAE <sub>train</sub>	MAE <sub>test</sub>
Hist. average	-85,429	-22,607	61.79	75.55
Random walk	-82,637	-22,150	33.48	<b>41.08</b>
Gauss-Homosk.	-79,167	-21,656	31.83	43.35
Gauss-Heterosk.	-77,801	-21,183	31.01	42.10
$t$ -Homosk.	-71,644	-19,399	28.79	42.64
$t$ -Heterosk.	-71,452	-19,356	<b>28.78</b>	42.58
$t$ -Full	<b>-71,380</b>	<b>-19,305</b>	28.80	42.45



## Dual Control by Deep Reinforcement Learning using a Deep Hyperstate Transition Model

A method is proposed for performing dual control using a deep reinforcement learning algorithm in combination with a neural network model trained to represent hyperstate transitions. The method is evaluated on a simple nonlinear system suggested as a suitable benchmark for such problems, but can scale to high-dimensional systems.

The hyperstate is compactly represented as the parameters of a mixture model, which is fitted to Monte Carlo samples from the hyperstate using the Expectation Maximization algorithm. This compact representation is then used to train a hyperstate transition model, which is used by a standard reinforcement learning algorithm to find a dual control policy. It is demonstrated that the method is able to learn a probing technique that reduces hyperstate uncertainty, yielding improved control performance.

Rosdahl, Christian  
Lund University



# Dual Control by Deep Reinforcement Learning using a Deep Hyperstate Transition Model

Christian Rosdahl, Anton Cervin, Bo Bernhardsson  
Department of Automatic Control, Lund University



## Abstract

A method is proposed for performing **dual control** using a **deep reinforcement learning algorithm** in combination with a **neural network model trained to represent hyperstate transitions**. The method is evaluated on a simple nonlinear system suggested as a suitable benchmark for such problems, but can scale to high-dimensional systems.

The **hyperstate** is compactly **represented as the parameters of a mixture model**, which is fitted to Monte Carlo samples from the hyperstate using the Expectation Maximization algorithm. This compact representation is then **used to train a hyperstate transition model**, which is used by a standard **reinforcement learning algorithm to find a dual control policy**. It is demonstrated that the method is able to learn a probing technique that reduces hyperstate uncertainty, yielding improved control performance.

## The Dual Control Problem

We consider a system with state and measurement equations

$$\begin{aligned} x_{k+1} &= f(x_k, u_k, \theta_k^f, w_k), \\ y_k &= g(x_k, \theta_k^g, e_k), \end{aligned}$$

with parameters  $\theta_k = (\theta_k^f, \theta_k^g)$  and disturbances  $w_k$  and  $e_k$ . The state  $x_k$  and parameter vector  $\theta_k$  are not deterministically known, but their initial probability distribution  $P(\alpha_0, \theta_0)$  is. The goal is to **determine a control policy  $\pi$  which minimizes some cost function  $J$ , given past inputs and outputs**  $\mathcal{D} = \{y_k, u_{k-1}, y_{k-1}, u_{k-2}, \dots\}$ . We define the **hyperstate** as

$$\xi_k = \mathbb{P}\{x_k, \theta_k \mid y_k, u_{k-1}, y_{k-1}, u_{k-2}, \dots\},$$

i.e., the probability distribution of the state and parameter vector given past data. The goal can then be described as finding a policy  $\pi(\xi_k) = \text{argmin}_u J(\pi, \xi_k)$  that minimizes a cost

$$J(\pi, \xi_k) = \mathbb{E} \left\{ \sum_{i=k}^{k+T-1} c_i(x_{i+1}, u_i = \pi(\xi_i)) \mid \xi_k \right\}.$$

## Hyperstate Transition Model

The probability density function (PDF) for the hyperstate  $\xi_k$  is approximately represented by a **mixture model**

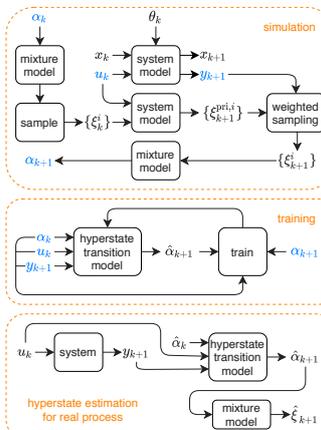
$$f(\xi) = \sum_{i=1}^c \lambda_i^k f_i(\xi; \beta_i^k),$$

where  $f_i$  are some simple PDFs, such as normal distributions,  $\sum_i \lambda_i = 1$ , and the distribution parameter vector

$$\alpha_k = [\lambda_1^k \dots \lambda_{c-1}^k \ (\beta_1^k)^T \dots (\beta_c^k)^T]^T$$

defines the approximate PDF. Using a system model for simulation, we collect data for how the **hyperstate parametric representation  $\alpha_k$  evolves for different inputs  $u_k$ , and the corresponding measurement signals  $y_{k+1}$** .

This data is used for training a hyperstate transition model, which yields a prediction of the new hyperstate representation  $\alpha_{k+1}$ , given the one at the previous timestep  $\alpha_k$ , the applied input  $u_k$ , and the measurement signal  $y_{k+1}$ . The hyperstate predictions from the model are then **used as the state in a reinforcement learning procedure**, which is used for determining an approximately optimal policy  $\pi(\xi_k)$ .



## Example System

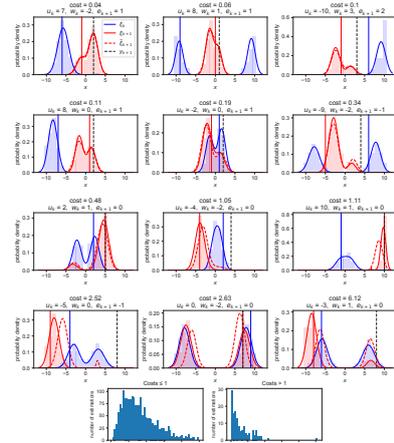
We try out the method on systems

$$x_{k+1} = x_k + u_k + w_k, \quad y_k = |x_k| + e_k,$$

with a linear state equation and a **simple but nonlinear measurement equation**. The hyperstate transition model is first tested for a case where the state  $x_k$ , input  $u_k$  and disturbances  $w_k$  and  $e_k$  are **discrete-valued**, since the hyperstate in this case can be computed exactly, and used for comparison with the model output. Then, we allow the signals to be **continuous-valued**, train a new model, and apply reinforcement learning with the estimated hyperstates, to find a control policy  $\pi(\xi)$ .

## Hyperstate Transition Model Results

Some **predictions (---)** of the next hyperstate representation  $\xi_{k+1}$  (**—**) by the hyperstate transition model, given the current hyperstate  $\xi_k$  (**—**), the measurement  $y_{k+1}$  (**---**), and the input  $u_k$  for the discrete-valued case.

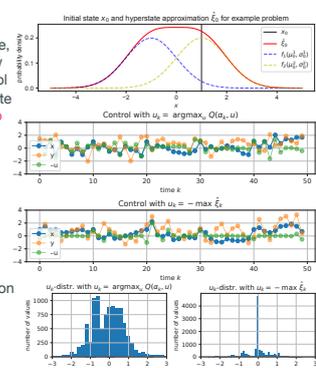


The real hyperstates are, for the illustrated example, discrete, and can therefore be represented exactly, as the bars in the diagrams. The vertical lines represent the real states  $x_k$  and  $x_{k+1}$ .

The evaluation cost is the Wasserstein  $L_1$ -distances between the predicted and real hyperstate representations.

## Control Results

We study a case where the system is run for episodes of length 10. Each episode is started at  $x_0 = 1$ . In this case, it is impossible for a controller to know whether a positive or a negative control signal should be applied to reduce state error, and **small controls do not lead to large enough changes** to resolve the issue, due to the measurement noise. Therefore, some **exploration**, by applying some larger control signal, could be useful.



Using  $u_k = -x$  as input, where  $x$  is the maximum-likely state, i.e. the state corresponding to the maximum of the probability density for the approximation of the hyperstate  $\xi_k$ , does not give better performance than using  $u_k = 0$ . However, using the **hyperstate approximation in combination with a simple reinforcement learning algorithm**, to get an action-value function  $Q$ , leads to a control policy which gives a smaller cost.

Control policy	Cost $J = \frac{1}{T} \sum_{t=0}^{T-1} \ x_t\ ^2$	Mean and standard deviations for the cost
$u_k = 0$	2.45 ± 0.04	
$u_k^{ML} = -\max(\xi)$	2.46 ± 0.04	
$u_k^{RL} = \text{arg max}_u Q(\xi_k, u)$	1.97 ± 0.05	
$u_k^{optimal} = -x_k$	1.20 ± 0.02	

## Fluid Models for Cloud Service Graphs

Resource management in cloud computing is a difficult problem, as one is often tasked with balancing between adequate service to clients and cost minimization in dynamic environments of many interconnected components. To make correct decisions in these environments, good performance models are necessary. A common modeling methodology is to use networks of queues, but as these are prohibitively expensive to evaluate for many real-time applications, different approximation methods for important metrics are frequently employed. Here we build on one such method, the fluid model, to generate a time-dynamic model for mixed networks with general phase-type services times and show how these can be extracted from tracing data of a service graph.



LUNDS  
UNIVERSITET



# Fluid Models for Cloud Service Graphs

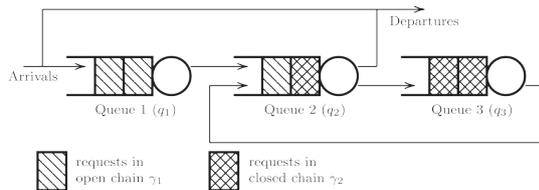
Johan Ruuskanen, Lund University  
Department of Automatic Control

## Modeling performance metrics in modern cloud applications

Resource management in cloud computing is a difficult problem, as one is often tasked with balancing between adequate service to clients and cost minimization in dynamic environments of many interconnected components. To make correct decisions in these environments, good performance models are necessary. A common modeling methodology is to use networks of queues, but as these are prohibitively expensive to evaluate for many real-time applications, different approximation methods for important metrics are frequently employed. Here we build on one such method, the fluid model, to generate a time-dynamic model for mixed networks with general phase-type services times and show how these can be extracted from tracing data of a service graph.

### Mean-field fluid model

Below follows an illustration of the mixed multiclass network.



It is possible to show that such a network, assuming that the queues follow either the processor sharing and/or delay disciplines, and where the service times have a general phase-type distribution, fulfills the so-called Kurtz's Theorem. This implies that a fluid model of the mean queue lengths can be obtained via the mean-field approximation.

By stacking the parametrization matrices of the phase-type distributions into block diagonals and combining with the routing matrix  $P$ , the mean-file fluid model becomes

$$\dot{x} = (\Psi + BPA^T)^T \theta(x) + A\lambda$$

Kurtz's theorem states that the queue lengths converges to the solution of this ODE as the system size scales to infinity. At lower system sizes, this model can however be inaccurate. Instead, it can be improved by using an inverse p-norm smoothing

$$\hat{\theta}_{i,r,a}(x, p) = \frac{x_{i,r,a}}{(1 + (k_i^{-1} \sum_{j,b} x_{i,j,b})^{p_i})^{1/p_i}}$$

Given such a fluid model, it is further possible to retrieve a closed form approximation of the entire response time CDF over almost any subset of classes in the network

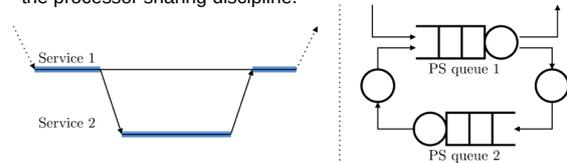
$$\Phi_{C_R}(t | \beta) \approx 1 - \beta^T A^T \exp [D\hat{g}(p^*) W_R t] \mathbb{1}$$

### References

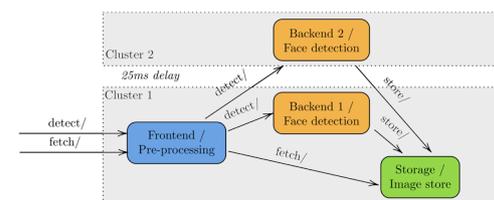
- Johan Ruuskanen, Tommi Berner, Karl-Erik Årzén, Anton Cervin. *Improving the mean-field fluid model of processor sharing queueing networks for dynamic performance models in cloud computing*, Performance Evaluation, 2021.
- Johan Ruuskanen, Haorui Peng, Alfred Åkesson, Lars Larsson, Maria Kihl. *FedApp: a Research Sandbox for Application Orchestration in Federated Clouds using OpenStack*. <https://github.com/JohanRuuskanen/FedApp>. 2021.

### Queuing network from trace

It is possible to retrieve a rudimentary multiclass network model directly from tracing data by assuming that each service follows the processor sharing discipline.

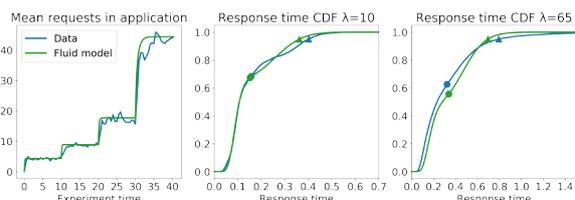


To evaluate such a model, we created the following example application performing face detection in the cloud.



Using the sandbox FedApp, the example application was then deployed in two Kubernetes clusters with Istio to handle routing and trace logging, and intercluster delay emulated by TC Neterm. The two call types were then loaded simultaneously with both open and closed connections, yielding the following results.

Arrival rate / client wait time	10 / 1	20 / 0.5	35 / 0.25	65 / 0.125
$\mathbb{E}$ [utilization]	0.19	0.36	0.59	0.84
max [utilization]	0.30	0.57	0.85	0.98



Due to modeling errors, the service time distributions will shift slightly depending on the load. In this example, we have dealt with this by refitting the distributions at each of the four operating points to demonstrate the ability to capture system metrics.

## Self-Driving Microservices

Recently, there has been a paradigm shift in software architectures from large monolithic applications into graphs of hundreds of loosely coupled microservices. The combination of this architectural transition and the DevOps movement with CI/CD has blurred the border between software development and IT operations. For IT operations, the new microservice paradigm results in a constantly evolving infrastructure landscape of software components. Ensuring of performance, reliability, and cost efficient operations in such dynamic environments is too complex for human operators, but autonomic computing mechanisms are required to make the systems increasingly manage themselves. The use of a service mesh enables outstanding observability without imposing any particular implementation costs during the development process, which suggests that it may be beneficial to develop methods for autonomous control of traffic management policies in the service mesh.

# Self-driving Microservices

Mohammad Reza Saleh Sedghpour, Umeå University  
Department of Computing Science

## Motivation

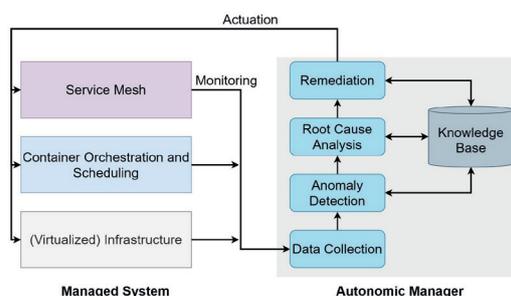
Recently, there has been a paradigm shift in software architectures from large monolithic applications into graphs of hundreds of loosely coupled microservices. The combination of this architectural transition and the DevOps movement with CI/CD has blurred the border between software development and IT operations. For IT operations, the new microservice paradigm results in a constantly evolving infrastructure landscape of software components. Ensuring of performance, reliability, and cost efficient operations in such dynamic environments is too complex for human operators, but autonomic computing mechanisms are required to make the systems increasingly manage themselves. The use of a service mesh enables outstanding observability without imposing any particular implementation costs during the development process, which suggests that it may be beneficial to develop methods for autonomous control of traffic management policies in the service mesh.

## State of the Art

The herein proposed research project on autonomic self-management for service mesh clusters, extends on early efforts on service mesh. Since the incarnation of autonomic computing [1], there have been substantial efforts within autonomic management of cloud infrastructures. Selected references include surveys by Jennings and Stadler [2] and Mani and Shyam [3]. The rather recent microservice concept has been studied from an architectural perspective e.g., by Jamshidi et al. [4] and Mendonça et al. [5] who propose architecture based self-adaption. Garriga defines a taxonomy of microservice architectures and identifies several challenges for autonomic management of microservices [6]. Toffetti et al. suggest an architectural framework for microservice self-management [7].

## Methodology

The management system gathers monitoring information from the whole stack. The gathered information, stored in the Data collection, consists mostly of time series data, but also discrete event data. Similarly, Actuation of management decisions may occur at different levels, e.g., circuit breaking and retry mechanism in the service mesh. In comparison with traditional resource management research in clouds, this microservice software stack gives both increased observability and additional actuators. The service meshes may be combined with distributed tracing that allows individual service invocations to be tracked across the microservice topology. This enables breakdown analysis of response times and thus greatly simplifies root cause analysis, at the expense of having to modify applications.



## Selected Results

Preliminary results related to the proposed project include two recent papers on service mesh. In the first paper [8], we proposed a controller to manage the circuit breaker adaptively in order to maximize throughput while maintaining response time of single service in service mesh. In the second paper [9], We studied the impact of various tuning parameters for circuit breaking and retry mechanisms empirically.

## References

1. Kephart, J., & Chess, D. (2003). The vision of autonomic computing. *Computer*, (1), 41-50
2. Jennings, B., & Stadler, R. (2015). Resource management in clouds: Survey and research challenges. *Journal of Network and Systems Management*, 23(3), 567-619
3. Manvi, S. S., & Shyam, G. K. (2014). Resource management for Infrastructure as a Service (IaaS) in cloud computing: A survey. *Journal of network and computer applications*, 41, 42444.
4. Jamshidi, P., Pahl, C., Mendonça, N. C., Lewis, J., & Tilkov, S. (2018). Microservices: The journey so far and challenges ahead. *IEEE Software*, 35(3), 24-35
5. Mendonça, N. C., Garlan, D., Schmerl, B., & Cámara, J. (2018). Generality vs. reusability in architecture-based self-adaptation: the case for self-adaptive microservices. *European Conference on Software Architecture*: (p. 18-24)
6. Garriga, M. Towards a taxonomy of microservices architectures. In *International Conference on Software Engineering and Formal Methods* (pp. 203-218).
7. Toffetti, G., Brunner, S., Blöchlinger, M., Dudouet, F., & Edmonds, A. (2015, April). An architecture for self-managing microservices. In *Proceedings of the 1st International Workshop on Automated Incident Management in Cloud* (pp. 19-24). ACM
8. Saleh Sedghpour, M.R., Klein, C., Tordsson, J. (2021). Service mesh circuit breaker: From panic button to performance management tool. *1st Workshop on High Availability and Observability of Cloud Systems (HAOC '21)*: (p. 4-10)
9. Saleh Sedghpour, M.R., Klein, C., Tordsson, J. (2021). An empirical study of service mesh traffic management policies for microservices. Submitted to 13th ACM/SPEC International Conference on Performance Engineering.

Salt Ducaju, Julian M.  
Lund University

## Joint Stiction Avoidance with Null-Space Motion in Real-Time Model Predictive Control for Redundant Collaborative Robots.

Null-space motion has been used in a Franka Emika Panda robot, a redundant collaborative robot, to ensure a continuous movement of all joints during an entire trajectory execution as an approach to avoid joint stiction and allow accurate kinesthetic teaching. As is conventional for collaborative and industrial robots, the Panda robot is equipped with an internal controller, which allows to send position and velocity references directly to the robot. Therefore, null-space motion can be added directly to the velocity references, which we generate using Model Predictive Control. The observed trajectory deviation caused by discretization approximations of the Jacobian matrix when implementing null-space motion has been corrected experimentally using sensor feedback for the real-time velocity-reference recalculation and by performing a fast sampling of the null-space vector. Null-space motion has been experimentally seen to contribute to reducing the friction torque dispersion present in static joints.

# Joint Stiction Avoidance with Null-Space Motion in Real-Time Model Predictive Control for Redundant Collaborative Robots [1]



LUNDS  
UNIVERSITET

J. M. Salt Ducaju, B. Olofsson, A. Robertsson, R. Johansson  
Department of Automatic Control, LTH, Lund University, Sweden

## Motivation

In kinesthetic teaching applications, the human operator should be comfortable with the physical interaction with the robot. Thus, it is important to be familiar with the force/torque required for leading the robot. Since the necessary force should not vary greatly between different human interventions, joint stiction should be avoided.

## Problem Formulation

This research experimentally analyzed the use of null-space motion to avoid joint stiction in a redundant robot. By adding null-space motion to a generated trajectory reference we ensured that no joint remained still during the trajectory execution and facilitate kinesthetic teaching.

Moreover, we evaluated the use of sensor feedback from joint angular position when online recalculating a point-to-point time-constrained trajectory using Model Predictive Control (MPC) [2] to address possible undesired side-effects of the null-space motion addition.



Franka Emika robot (7-DoF collaborative robot) used in the experiments.

## Null-Space Motion Addition

The null-space unitary vector,  $\hat{q}_{NSU}$ , must be scaled before being added to the MPC-generated reference,  $\hat{q}_{MPC}$ , and sent to the robot as a velocity reference,  $\hat{q}_{ref}$ .

$$\hat{q} = J^1(q)\gamma + N(q)\hat{q}_a \quad \hat{q}_{NSU} = \frac{N(q)\hat{q}_a}{\|N(q)\hat{q}_a\|} \quad \hat{q}_{ref} = \hat{q}_{MPC} + \hat{q}_{NS}$$

$$N(q) = I_7 - J^1(q)J(q) \quad \hat{q}_{NS} = \hat{q}_{NSU} \alpha \sin\left(\frac{2\pi t}{T_F}\right)$$

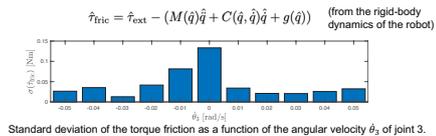
(where the matrix  $N(q)$  projects the additional arbitrary joint angular velocity,  $\hat{q}_a$ , into the null-space so that it is independent of the end-effector Cartesian motion)

## References

[1] J. M. Salt Ducaju, B. Olofsson, A. Robertsson and R. Johansson, "Joint Stiction Avoidance with Null-Space Motion in Real-Time Model Predictive Control for Redundant Collaborative Robots," in IEEE International Conference on Robot & Human Interactive Communication (RO-MAN), 2021, pp. 307-314.  
 [2] M. Ghazaei Ardakani, B. Olofsson, A. Robertsson, and R. Johansson, "Model predictive control for real-time point-to-point trajectory generation", *IEEE Transactions on Automation Science and Engineering*, vol. 16, no. 2, pp. 972-983, Apr. 2019.  
 [3] M. Linderoth, A. Stolt, A. Robertsson, and R. Johansson, "Robotics force estimation using motor torques and modeling of low velocity friction disturbances", *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Tokyo, Japan, Nov. 3-7, 2013, pp. 3550-3556.

## Results

A reduction of friction-torque dispersion has been experimentally observed thanks to the addition of null-space motion in a joint that otherwise would be static.



We investigated trajectory deviations caused by discretization approximations of the Jacobian matrix when implementing null-space motion.

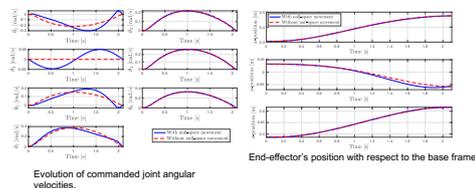
END-EFFECTOR'S FINAL CARTESIAN POSITION ERRORS [MM]

	CL - NS	CL	OL - NS	OL
Mean	1.43	0.48	5.97	1.18
Std. Dev.	0.78	0.25	0.39	0.11

(CL = closed-loop, OL = open-loop, NS = null-space)

The effects of these approximations have been reduced by:

- Performing a fast sampling of the null-space vector.
- Using sensor feedback for the online trajectory-reference recalculation.



## Conclusion

Null-space motion addition can avoid joint stiction, while also avoiding the appearance of robot vibrations that may appear with dithering [3]. However, such a null-space approach is only applicable to redundant robots. Moreover, the use of sensor feedback for online trajectory recalculation has shown to address the side-effects caused by null-space motion addition.

Schuppe, Georg  
KTH

## Decentralized Multi-Agent Strategy Synthesis via Exchange of Least-Limiting Advisers

We propose a decentralized solution to a highlevel task-planning problem for a multi-agent system under a set of possibly dependent LTLf specifications. We propose an approach where the problem is turned into a number of individual two and a half player stochastic games with reachability objectives. If almost-surely winning strategies cannot be found for them, we deploy so-called least-limiting advisers to restrict agents' behaviours. A key step is treating safety and liveness separately, by synthesizing necessary safety and fairness assumptions and iteratively exchanging them in the form of advisers between the agents. We avoid the state-space explosion problem by computing advisers locally in each game, independently of the model and specification of other agents.

The solution is sound, but conservative. We demonstrate its scalability in a series of simulated scenarios involving cleaning of an office-like environment.

# Decentralized Multi-Agent Strategy Synthesis via Exchange of Least-Limiting Advisers

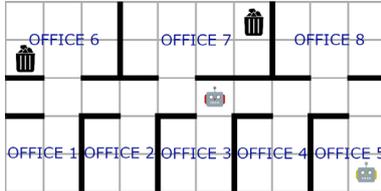
Georg Friedrich Schuppe and Jana Tumova  
KTH Royal Institute of Technology  
{schuppe, tumova}@kth.se



## Motivation

Heterogeneous robots in shared environment might occasionally be required to collaborate, even though they were originally not deployed to operate as a team.

Example:



**Figure 1:** A partitioned office-like environment. A state of a robot is determined by its orientation and the cell it occupies. In each state, a robot can choose to stay, move forward, or turn 90°.

- Bin-emptying robots are tasked to empty  $k$  specific bins in the offices
- Cleaning robots need to clean detected spillages in certain offices and guarantee that none of the bin-emptying robots enter the affected office in order to prevent further damage

How do we express their interdependent tasks?  
How do we ensure that the tasks are accomplished?

## Safety Advisers

**Definition 1** (Minimality). A safety assumption  $E_s$  is minimal if  $|E'_s| \leq |E_s|$  for all safety assumptions  $E'_s \in E_2$ . The unique, minimal safety assumption can be computed as

$$E_s = \{(s, s') \in E_2 \mid s \in \langle(1, 2)\rangle\psi \text{ and } s' \notin \langle(1, 2)\rangle\psi\}$$

The assumption  $E_s$  cannot be directly communicated as an adviser to the other agents. Instead, we communicate the advice in the form of an adviser:

**Definition 2** (Safety Adviser). A safety adviser is a set of tuples:

$$SafeAdv = \{(pre, \sigma) \mid pre \in AP_i, \sigma \in \widehat{\Sigma}_i\}$$

Given that agent  $i$  satisfies  $pre$ , other agents should not satisfy  $\sigma$  in their next state.

Safety Advisers are implemented by expanding the specification formula of affected agents:

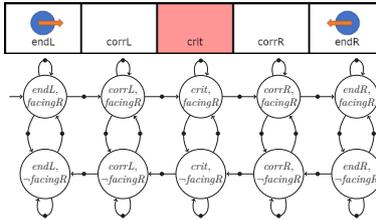
$$\phi_{(pre, \sigma), i} = G(pre \rightarrow \neg Xproj_i(\sigma)),$$

We can incorporate all advisers from all agents into the specification of the agent  $i$  through conjunction:

$$\phi_{s, i} = \bigwedge_{(pre, \sigma) \in SafeAdv_j, j \in N} \phi_{(pre, \sigma), i}$$

## Problem Formulation

- Each agent modelled an MDP  $\mathcal{M}_i$



**Figure 2:** A small example of an MDP modeling the left robot in the corridor illustrated above. A state of a robot is determined by its orientation and the cell it occupies, the actions are to move, turn around, or stay.

- Each agent given an LTL<sub>f</sub> specification  $\phi_i$

$$\rightarrow \phi_i = \bigwedge_{k \in \{1, \dots, \ell\}} Fbin_{i,k},$$

$$\rightarrow \phi_j = Foff_{j, o_j} \wedge G(\bigwedge_{i \in \{1, \dots, n\}} \neg off_{i, o_j})$$

- Develop an efficient procedure to synthesize reactive strategies for all  $\mathcal{M}_i$  such that all  $\phi_i$  are satisfied, i.e. avoiding to construct a centralized Product MDP

## Fairness Advisers

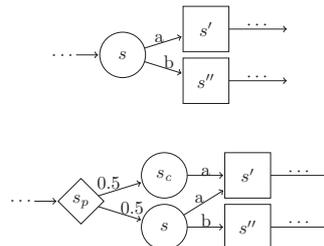
Since computing a minimal fairness assumption is NP-hard, we compute *locally minimal* fairness assumptions instead. Similarly to safety assumptions, we transform fairness assumptions into Fairness Advisers:

**Definition 3** (Fairness Adviser). A fairness adviser is a set of tuples:

$$FairAdv = \{(pre, \sigma) \mid pre \in AP_i, \sigma \in \widehat{\Sigma}_i\}$$

Given that agent  $i$  satisfies  $pre$ , other agents should satisfy  $\sigma$  with non-zero probability in their next state.

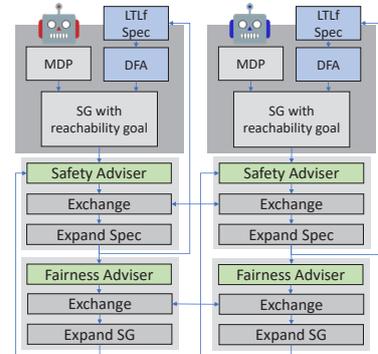
Fairness Advisers are implemented through explicit modification of the stochastic games.



**Figure 4:** Enforcing fairness on  $(s, a, s')$  by prepending a probabilistic state.

## Contributions and Approach

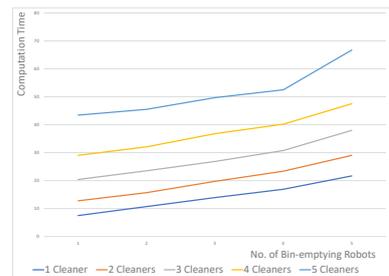
- A *reactive synthesis*-based approach for multi-agent high-level task planning
- A novel, decentralized approach via exchange of least-limiting advisers
- Demonstrating the scalability of the approach on selected use-cases



**Figure 3:** Schema of the approach for two agents. Each agent constructs their stochastic game locally and computes minimal, necessary assumptions on the behaviour of other agents. In an iterative process, agents incorporate advice from each other and compute additional advisers, if necessary.

## Results

- The solution is sound, but conservative
- Conservativeness stems from the information gap between agents and the implementation of fairness advisers
- When dependencies between agent specifications are low, the computation time depending on the number of agents behave almost linearly
- Conveying advisers to humans poses an interesting line of future research



## Acknowledgement

This work is partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation and the Swedish Research Council (VR) (project no. 2017-05102).

## Complexity certification of Mixed-Integer Quadratic Programming

In hybrid model predictive control (MPC), a non-convex optimization problem has to be solved at each time step, which in real-time applications makes it important to solve these efficiently and to have good upper bounds on worst-case solution time. For linear hybrid MPC problems, the optimization problem is often a multi-parametric mixed-integer quadratic program (mp-MIQP) that depends on parameters such as system states and reference signals.

The aim of the research is to certify the complexity of MIQPs by computing which sequence of subproblems are required to solve in the branch and bound (B&B) method for every parameter of interest. These sequences can be used to compute the worst-case bounds on how many iterations, floating-point operations and, ultimately, the maximum solution time, the B&B algorithm would require to converge online.

# Complexity certification of Mixed-Integer Quadratic Programming

Shamisa Shoja, Linköping University  
Division of Automatic Control, Department of Electrical Engineering (ISY)  
Supervisor: Daniel Axehill



## Motivation & Research Goals

In hybrid model predictive control (MPC), a non-convex optimization problem has to be solved at each time step, which in real-time applications makes it important to solve these efficiently and to have good upper bounds on worst-case solution time. For linear hybrid MPC problems, the optimization problem is often a multi-parametric mixed-integer quadratic program (mp-MIQP) that depends on parameters such as system states and reference signals. The aim of the research is to certify the complexity of MIQPs by computing which sequence of subproblems are required to solve in the branch and bound (B&B) method for every parameter of interest. These sequences can be used to compute the worst-case bounds on how many iterations, floating-point operations and, ultimately, the maximum solution time, the B&B algorithm would require to converge online.

## Methods

### Problem Formulation

– mp-MIQP

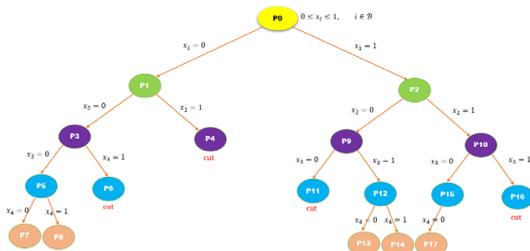
$$\begin{aligned} \min_x \quad & \frac{1}{2}x^T Hx + f^T x + \theta^T f_\theta^T x, \\ \text{s.t.} \quad & Ax \leq b + W\theta, \\ & x_i \in \{0, 1\}, \quad \forall i \in \mathcal{B} \end{aligned}$$

- \*  $x = [x_c^T, x_b^T]^T \in R^{n_c} \times \{0, 1\}^{n_b}$ : state vector
- \*  $\theta \in \Theta_0 \subset R^{n_\theta}$ : parameter vector

### B&B method

Solving a sequence of relaxed convex mp-QP problems by fixing a binary variable to 0 and 1, forming nodes in the B&B search tree and cut a node if the solution of a relaxation is

- infeasible
- does not provide better solution
- integer feasible



### Contribution:

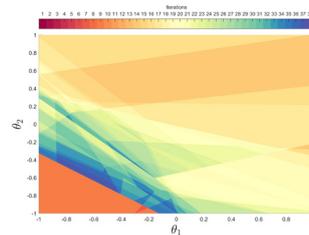
An algorithm for computing a useful upper bound of the worst-case computational complexity for solving any possible MIQP that can arise from a specific parameter in a polyhedral parameter set

## References

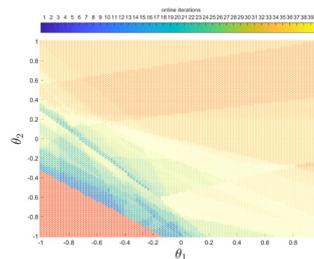
- [1] A parametric branch and bound approach to suboptimal explicit hybrid MPC  
Daniel Axehill, Thomas Besselmann, Davide Martino Raimondo, Manfred Morari  
Automatica, 2014
- [2] A unifying complexity certification framework for active-set methods for convex quadratic programming  
Daniel Axehill, Daniel Axehill  
IEEE Transactions on Automatic Control, 2022
- [3] Integer programming  
Laurence A. Wolsey  
John Wiley & Sons, 2020

## Results

Partitioning the parameter space based on the total number of QP iterations, i.e., the total number of linear system of equations solved, for a random example with  $n_c = 2$ ,  $n_b = 4$ ,  $n_\theta = 2$ , from the proposed certification algorithm. Points with the same color share the same number of complexity measure.



The total QP iteration number for 10000 samples specified by \* in the parameter space derived by applying online B&B to the same example.



The complexity certification result coincides with the online algorithm in all sample points, despite that the conservative upper bound is used in the certification method, using depth-first search strategy

### Ongoing & Future works:

- Exact complexity certification of mixed-integer linear programming (MILP) (ongoing)
- Complexity certify the B&B method for different node selection strategies such as best-first strategy
- Certification of the warm-started algorithm to decrease the computational complexity

Song, Qunying  
Lund University

## Critical Scenario Identification for Realistic Testing of Autonomous Driving Systems

Testing of autonomous vehicles involves enormous challenges for the automotive industry. The number of real-world driving scenarios is extremely large, and choosing effective test scenarios is essential, as well as combining simulated and real-world testing. We present an industrial workbench of tools and workflows to generate critical test scenarios for active safety and autonomous driving functions in an efficient way. The workbench is based on existing engineering tools and helps smoothly integrate simulated testing, with real vehicle parameters and software. We validate the workbench with real autonomous driving systems and demonstrate its effectiveness for the realistic testing of such systems.

# Software Testing of Autonomous Systems

Qunying Song, Lund University  
Department of Computer Science



LUNDS  
UNIVERSITET

## Concepts in Testing of Autonomous Systems

Testing of autonomous systems is extremely important as many of them are both safety-critical and mission-critical, yet it is still an open challenge on how to test such systems effectively and efficiently. To gain a better understanding of autonomous systems practice and facilitate testing of different autonomous systems, we conduct an exploratory study [1] by synthesizing existing academic literature with a focus group discussion and interviews with industry practitioners. As a result, we present a conceptualization of autonomous systems, classifications of challenges and current practices as well as of available techniques and approaches for testing of autonomous systems.

### Critical Scenario Identification for Realistic Testing of Autonomous Driving Systems

The number of real-world operational scenarios for autonomous systems is extremely large, and choosing effective test scenarios is essential, as well as combining simulated testing and real-world testing. We focus on a common area within autonomous systems – autonomous vehicles and establish an industrial workbench to generate efficient and effective scenarios for testing such systems [2]. The workbench consists three existing engineering tools and a workflow, and helps smoothly integrate simulated testing, with real vehicle parameters and software. We also demonstrate the effectiveness of the workbench by using two real autonomous driving systems from industry by collaborating with Volvo Cars.

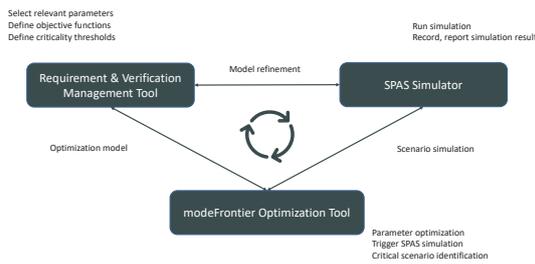


Figure-1. Tools and workflow involved in the workbench for critical scenario identification for autonomous driving systems

### References

1. Song, Qunying, Emelie Engström, and Per Runeson. "Concepts in Testing of Autonomous Systems: Academic Literature and Industry Practice." In *WAIN'21 1st Workshop on AI Engineering*. IEEE Computer Society, 2021.
2. Song, Qunying, Kaige Tan, Per Runeson, and Stefan Persson. "An Industrial Workbench for Test Scenario Identification for Autonomous Driving Software." In *2021 IEEE International Conference on Artificial Intelligence Testing (AITest)*, pp. 81-82. IEEE, 2021.

### A Vehicle-Pedestrian Time-To-Collision Model for Testing of Autonomous Driving Systems

While autonomous driving systems are expected to reduce road accidents and improve traffic safety, understanding of the intensive and complex traffic situations is prerequisite to enable testing of such systems in a realistic traffic setup. We propose a model that predicts the worst-case distribution of TTC (Time-to-Collision) for vehicle-pedestrian interactions at unsignalized crossings, based on the traffic density. We validate the model using real traffic data collected in Sweden. We also demonstrate its use for testing of autonomous driving systems by connecting the model to critical test scenario identification for an autonomous emergency braking function from the industry.

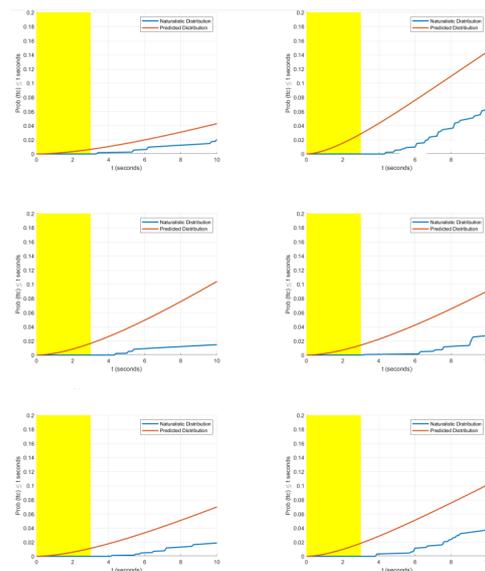


Figure-2. Model validation with naturalistic traffic data collected by Viscando in Linköping, Sweden

## Bayesian Prediction with Covariates Subject to Detection Limits

Missing values in covariates due to censoring by signal interference or lack of sensitivity in the measuring devices are common in industrial problems. We propose a full Bayesian solution to the prediction problem with an efficient Markov Chain Monte Carlo (MCMC) algorithm that updates all the censored covariate values jointly in a random scan Gibbs sampler. We show that the joint updating of missing covariate values can be at least two orders of magnitude more efficient than univariate updating. This increased efficiency is shown to be crucial for quickly learning the missing covariate values and their uncertainty in a real-time decision making context, in particular when there is substantial correlation in the posterior for the missing values. The approach is evaluated on simulated data and on data from the telecom sector. Our results show that the proposed Bayesian imputation gives substantially more accurate predictions than naïve imputation, and that the use of auxiliary variables in the imputation gives additional predictive power.

Svahn, Caroline  
Linköping University / Ericsson AB



# Bayesian Prediction with Covariates Subject to Detection Limits

Caroline Svahn<sup>†,‡</sup>, Mattias Villani<sup>†,§</sup>

<sup>†</sup>IDA, Linköping University: {firstname.lastname}@liu.se

<sup>‡</sup>Ericsson Research, Linköping: {firstname.lastname}@ericsson.com

<sup>§</sup>Dept of Statistics, Stockholm University: {firstname.lastname}@stat.su.se



## DESCRIPTION

Missing values in covariates due to censoring by signal interference or lack of sensitivity in the measuring devices are common in industrial problems. We propose a full Bayesian solution to the prediction problem with an efficient Markov Chain Monte Carlo (MCMC) algorithm that updates all the censored covariate values jointly in a random scan Gibbs sampler. We show that the joint updating of missing covariate values can be at least two orders of magnitude more efficient than univariate updating. This increased efficiency is shown to be crucial for quickly learning the missing covariate values and their uncertainty in a real-time decision making context, in particular when there is substantial correlation in the posterior for the missing values. The approach is evaluated on simulated data and on data from the telecom sector. Our results show that the proposed Bayesian imputation gives substantially more accurate predictions than naïve imputation, and that the use of auxiliary variables in the imputation gives additional predictive power.

## BACKGROUND & MOTIVATION

While frequentist approaches generally have the advantage of being relatively fast, Bayesian methods can quantify the uncertainty for both parameters and predictions in a way that is directly usable for decision making under uncertainty. This is clearly crucial in safety critical scenarios where faulty decisions have severe consequences, but also in less dramatic but often occurring decisions, such as in wireless telecommunications where a faulty decision may disconnect users from the network [1]. The existing Bayesian literature use Gibbs sampling algorithms to simulate from the joint posterior of the model parameters and the missing values. The proposed Gibbs samplers update the missing covariate values in an observation conditional on all other missing values, see e.g. [2] and [3]. This can be highly inefficient when the missing values are highly correlated in the posterior.

## METHODS & RESULTS

The complete model can be written

$$y_i = \beta_0 + \tilde{\beta}^\top \mathbf{x}_i + \varepsilon_i$$

$$\mathbf{x}_i = \Gamma^\top \mathbf{w}_i + \mathbf{v}_i,$$

with  $\varepsilon_i \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2)$ ,  $\mathbf{v}_i \stackrel{iid}{\sim} \mathcal{N}(0, \Omega)$  and where  $x_{ij}$  is unobserved if  $x_{ij} < c_{ij}$ .

We take a Bayesian approach and assume the following prior with independence between the

blocks of parameters

$$\beta_0 \sim \mathcal{N}(0, \tau_{\beta_0}^2)$$

$$\tilde{\beta} \sim \mathcal{N}(0, \tau_{\tilde{\beta}}^2 \mathbf{I}_p)$$

$$\sigma^2 \sim \text{IG}(a, b)$$

$$\gamma | \Omega \sim \mathcal{N}(0, \Omega \otimes \tau_\gamma^2 \mathbf{I}_r)$$

$$\Omega \sim \text{IW}(\mathbf{A}, \kappa),$$

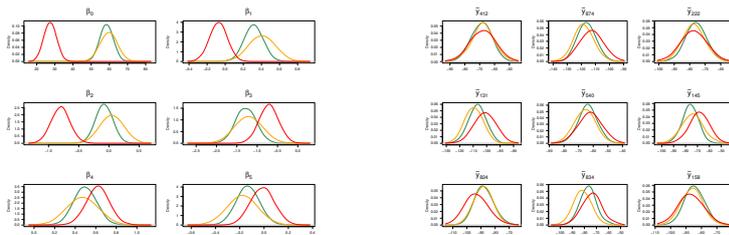
where  $\gamma = \text{vec} \Gamma$  stacks the columns of  $\Gamma$  in a vector,  $\text{IG}(a, b)$  is the inverse Gamma

distribution and  $\text{IW}(\mathbf{A}, \kappa)$  is the inverse Wishart distribution with  $\kappa$  degrees of freedom.

The missing values in a given observation,  $\mathbf{x}_i^{(m)}$ , are conditionally independent of the missing values in all other observations. Each  $\mathbf{x}_i^{(m)}$  vector can therefore be drawn in parallel from truncated multivariate normal distributions.

Dataset	Quantiles				
	0%	25%	50%	75%	100%
1	0.52	1.03	1.25	6.61	43.8
2	0.51	1.04	1.29	9.08	65.6
3	0.54	1.03	1.22	6.62	61.1
4	0.48	1.03	1.24	8.00	57.7
5	0.55	1.03	1.25	10.13	141.7

**Table 1:** Ratios of the effective sample size (ESS) comparing joint sampling of missing values to univariate updates. Each row represents the quantiles of the ratio  $\text{ESS}_{\text{multi}}/\text{ESS}_{\text{uni}}$  for a simulated dataset with  $n = 1000$  observations on  $p = 40$  covariates and approximately 40 % censored values.



**Figure 1:** Posterior densities for  $\beta$  and a subset of predictive distributions. We compare our Bayesian imputation (yellow) to two baselines. The green line is an idealized model using the uncensored (complete) data, the red represent a naïve imputation strategy (red) where

$$\mathbf{x}_i^{(m)} = \max(\mathbf{x}_i^{(o)}) - \Delta,$$

meaning that all missing values are imputed to the lowest limit of detection.

## RESEARCH GOAL & QUESTION

Efficient sampling of missing values is particularly important in the prediction phase where the missing covariates for a new observation must be learned quickly in a real-time context. We therefore develop a fast and efficient Markov Chain Monte Carlo (MCMC) algorithm that samples all missing covariates jointly. The joint sampling is performed using the recently proposed and highly efficient truncated multivariate normal simulation algorithm in [4] and we additionally propose a random scan implementation [5] to further increase the speed of the missing covariate updating step. The data are censored according to the following principle, which aims to mimic the censoring due to interference from the strongest signal among a set of signals:

$$x_{ij} = \begin{cases} x_{ij} & \text{if } x_{ij} \geq \max(\mathbf{x}_i^{(o)}) - \Delta \\ \max(\mathbf{x}_i^{(o)}) - \Delta & \text{otherwise,} \end{cases} \quad (1)$$

where  $\Delta$  is a known distance from the strongest signal for which covariates of lower amplitude are still detectable.

## BIBLIOGRAPHY

- [1] Rydén, H., Berglund, J., Isaksson, M., Cöster, R. and Gunnarsson, F., Predicting Strongest Cell on Secondary Carrier using Primary Carrier Data, IEEE Wireless Communications and Networking Conference Workshops (WCNCW): 7th International Workshop on Self-Organizing Networks (IWSON), pp. 137–142, 2018.
- [2] Yue, Yu Ryan and Wang, Xiao-Feng, Bayesian inference for generalized linear mixed models with predictors subject to detection limits: An approach that leverages information from auxiliary variables, Statistics in medicine, vol. 35(10), 2016, pp. 1689–1705.
- [3] Wu, Huiyun and Chen, Qingxia and Ware, Lorraine and Koyama, Tatsuki, A Bayesian Approach for Generalized Linear Models with Explanatory Biomarker Measurement Variables Subject to Detection Limit - an Application to Acute Lung Injury, Journal of applied statistics, vol. 39, 2012, pp. 1733–1747.
- [4] Botev, Zdravko I, The normal law under linear restrictions: simulation and estimation via minimax tilting, Journal of the Royal Statistical Society, B, vol. 79, 2016, pp. 1–24.
- [5] Amit, Yali and Grenander, Ulf, Comparing sweep strategies for stochastic relaxation, Journal of multivariate analysis, vol. 39(2), 1991, pp. 197–222.

Varnai, Peter  
KTH

## Revisiting Path Integral Policy Improvement

Path Integral Policy Improvement (PI2) is a reinforcement learning algorithm developed for solving stochastic optimal control problems. The main idea is to linearize the stochastic Hamilton-Jacobi-Bellman equations underlying the control problem to allow optimal feedback controls to be calculated from path integrals, ie., open-loop sample trajectories and corresponding costs, of the dynamical system. The computation is highly parallelizable and can be accelerated using modern hardware, such as GPUs, for real-time implementation. This work revisits the method, simplifying and expanding upon its previous theory, in hopes of paving way for future potential applications.

## Revisiting Path Integral Policy Improvement

P. Varnai, PhD, KTH Royal Institute of Technology  
Division of Decision and Control Systems  
Supervisor: D. V. Dimarogonas



## Motivation &amp; Research Goals

Policy improvement with path integrals (PI<sup>2</sup>) is a reinforcement learning algorithm developed for solving stochastic optimal control problems. The main idea is to linearize the stochastic Hamilton–Jacobi–Bellman (HJB) equations underlying the control problem to allow optimal feedback controls to be calculated from path integrals, *i.e.*, open-loop sample trajectories and corresponding costs, of the dynamical system. The computation is highly parallelizable and can be accelerated using modern hardware, such as GPUs, for real-time implementation. This work revisits the method, simplifying and expanding upon its previous theory, in hopes of paving way for future potential applications.

## Methods

Policy improvement with path integrals (PI<sup>2</sup>)<sup>[1]</sup> is a strategy for controlling dynamical systems of the form:

$$\dot{\mathbf{x}}_t = \mathbf{f}(\mathbf{x}_t, t) + \mathbf{g}(\mathbf{x}_t)\mathbf{u}_t + \Sigma_{x,t}^{1/2}\epsilon_{x,t}, \quad (1)$$

where  $\mathbf{x}_t \in \mathbb{R}^n$  is the system state,  $\mathbf{u}_t \in \mathbb{R}^m$  is the control input,  $\epsilon_{x,t} \in \mathbb{R}^n$  is zero-mean Gaussian white noise with covariance  $\Sigma_{x,t} \geq 0$ ,  $\mathbf{f}(\cdot)$  describes the autonomous system dynamics, and  $\mathbf{g}(\cdot)$  describes the influence of the input. The goal is to find the controls  $\mathbf{u}_t$  which minimize a pre-defined trajectory cost in the form of a feedforward / feedback strategy

$$\mathbf{u}_t = \mathbf{k}_t + \delta\mathbf{u}_t, \quad (2)$$

where the feedforward  $\mathbf{k}_t$  is parameterized by some  $\mu_{0t}$  through its derivative as  $\dot{\mathbf{k}}_t = \mu_{0t}$ , and  $\delta\mathbf{u}_t$  is the feedback.

Let  $\tau_T$  denote the trajectory of the combined state  $\mathbf{z}_t := (\mathbf{x}_t, \mathbf{k}_t)$  during a time horizon  $T > 0$ , *i.e.*,  $\tau_T := \{\mathbf{z}_t \mid 0 \leq t \leq T\}$ . A cost  $S_t(\tau_T)$  is then assigned to each trajectory in the form

$$S_t(\tau_T) \equiv S(\tau_T, t) := \phi(\tau_T) + \int_t^T q(\mathbf{z}_s, s) ds, \quad (3)$$

where  $\phi(\tau_T)$  is a terminal cost and  $q_t(\mathbf{z}_t) \equiv q(\mathbf{z}_t, t)$  is an instantaneous running cost. The PI<sup>2</sup> theory then chooses the feedback  $\delta\mathbf{u}_t$  in a specific way such that the expected closed-loop cost can be calculated using the open-loop system dynamics with  $\mathbf{u}_t = \mathbf{k}_t$ , *i.e.*, with  $\delta\mathbf{u}_t$  set to zero:

$$V(\tau_t) = -\lambda \log \mathbb{E}_{OL} \left[ \exp \left( -\frac{1}{\lambda} S_t(\tau_T) \right) \mid \tau_t \right]. \quad (4)$$

Here,  $\lambda > 0$  is a controller parameter representing the trade-off between control effort and correction for system noise. Note that  $V(\tau_t)$  depends on  $\mathbf{k}_t$  through the expectation, and thus its parameterization  $\mu_{0t}$ .

A practical realization of the PI<sup>2</sup> control strategy then consists of the following two stages:

- I. The parameters  $\mu_{0t}$  of the feedforwards  $\mathbf{k}_t$  are optimized to minimize the expected closed-loop cost  $V(\tau_t)$  under the PI<sup>2</sup> feedback control strategy.
- II. The closed-loop feedbacks  $\delta\mathbf{u}_t$  are calculated and implemented during real-time operation.

The expression (4) allows the computations necessary for both of these steps to rely on expectations involving noisy open-loop trajectory samples. As the samples can be generated independently of one another, the strategy is highly parallelizable computation-wise, and offers performance guarantees for the stochastic optimal control problem in terms of the expected cost.

## Selected Results

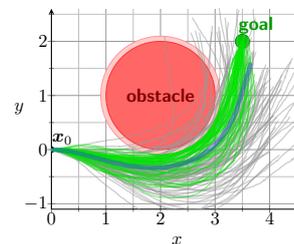
The main theoretical results are summarized as follows:

- 1) The computations needed to find the feedbacks  $\delta\mathbf{u}_t$  in real-time are shown to take a much simpler form than in previous works. Theoretical derivations to arrive at the result are simplified while the results are valid for a broader range of system dynamics and without some previously necessary assumptions<sup>[2]</sup>.
- 2) Our introduction of the feedforward state  $\mathbf{k}$  allows us to connect determining the optimal parameters  $\mu_{0t}$  to natural gradient descent, yielding iterative updates of the form previously only heuristically proposed in the literature<sup>[3]</sup>.

Furthermore, our work is the first to demonstrate the correctness of the theoretical results by showing that the achieved closed-loop cost after implementing the PI<sup>2</sup> control strategy matches the predicted value from (4). Our results also lay a solid theoretical foundation for addressing extensions and further improvements to the PI<sup>2</sup> theory; for example, we have already extended our results to the multi-agent case.

**Table 1:** Achieved average closed-loop costs as a function of the number of  $N$  roll-outs used for feedback calculation.

$\delta\mathbf{u}_t$ calculation	$N = 500$	$N = 5000$	$N = 50000$	theoretical
$\mathbb{E}_{CL}[S], \lambda = 0.01$	2.5	1.21	1.14	1.08
$\mathbb{E}_{CL}[S], \lambda = 0.05$	1.42	1.39	1.37	1.39
$\mathbb{E}_{CL}[S], \lambda = 0.2$	1.86	1.82	1.82	1.85



**Figure 1:** Sample open-loop (gray) and closed-loop (green) trajectories obtained for a reach and avoid control problem involving a unicycle system with  $\lambda = 0.01$  and  $N = 5000$  roll-outs used for feedback calculation. The thicker blue line represents the nominal, noiseless trajectory obtained by following the optimized open-loop feedforward control actions.

## References

- [1] A generalized path integral control approach to reinforcement learning  
E. Theodorou, J. Buchli, and S. Schaal  
Journal of Machine Learning Research (JMLR), 2010
- [2] The Two-Stage PI<sup>2</sup> Control Strategy  
P. Varnai and D. V. Dimarogonas  
IEEE Control Systems Letters (L-CSS), 2021
- [3] Path Integral Policy Improvement: An Information-Geometric Optimization Approach  
P. Varnai and D. V. Dimarogonas  
Submitted to the Journal of Machine Learning Research (JMLR), 2020

Vladu, Emil  
Lund University

## Robust Control of Large-scale Networks

Networks in various domains such as district heating, power systems and transportation have grown increasingly complex over the years. It is therefore of interest to control the system in a manner which scales well with the network size. In this project, we are particularly concerned with scalable means of suppressing the influence of disturbances on the desired output. A recent result features a distributed controller which optimally suppresses worst-case disturbances for a class of nonlinear systems.

Vladu, Emil  
Lund University

## Robust Control of Large-scale Networks

Emil Vladu, Lund University  
Dept. of Automatic Control  
Main supervisor: Anders Rantzer

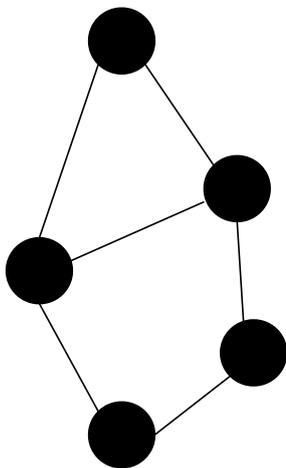


LUND UNIVERSITY

### Motivation & Research Goals

Networks in various domains such as district heating, power systems and transportation have grown increasingly complex over the years. It is therefore of interest to control the system in a manner which scales well with the network size. In this project, we are particularly concerned with developing scalable means of suppressing the influence of disturbances on the desired output.

### Methods



The control of networks becomes more *computationally demanding* as the system increases in size, with many nodes and interconnections as a result. For example, it may not be feasible for each controller in a network to have access to data from all nodes in the network. Such limitations can be circumvented by maintaining good control using only *local* information. In this project, we are particularly concerned with the control of *nonlinear* systems in the face of *disturbances*.

**Goal:** suppress the impact of disturbances on a desired output in a way suitable for large-scale systems.

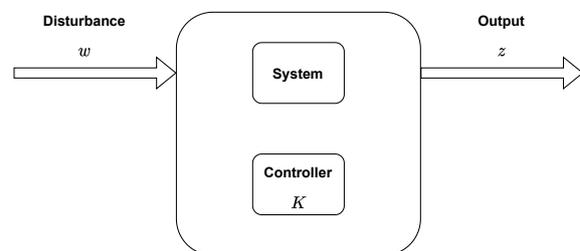
We use techniques from various areas both within and outside of control, including:

1. Robust Control
2. Nonlinear Control
3. Optimization
4. Positive Systems
5. Network Dynamics

### References

1. E. Vladu, C. Bergeling and A. Rantzer, *Global Solution to an H-infinity Control Problem with Input Nonlinearity*, CDC, 2021

### Recent Result

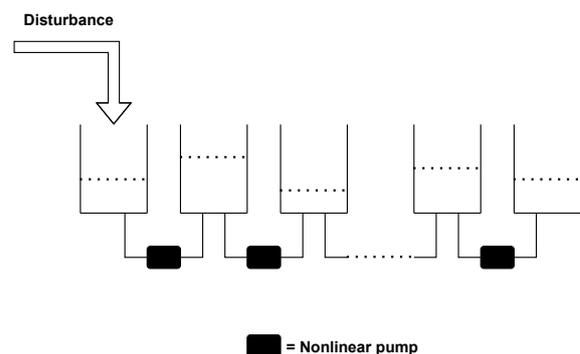


**Problem:** find controller  $K$  to minimize the worst-case disturbance impact on actuator effort and desired output

$$\min_K \max_w \frac{\|z\|}{\|w\|}$$

**Systems considered:** nonlinear systems with state-dependent input matrix and constant, stable and symmetric state matrix, where the actuator influence is the weakest at the desired output.

**Result<sup>1</sup>:** closed-form expression for an optimal controller with potential for scalability.



**Application:** consider a network of buffers with some contents, e.g. water, connected by flow links.

**Goal:** find a controller to minimize worst-case disturbance impact on pump effort and deviations from a desired water level.

**Complication:** nonlinear pumps.

**Application of the result<sup>1</sup> gives:** an optimal controller such that each pump only requires info from its neighboring buffers.  
→ Useful for large-scale networks!

Wingqvist, Birgitta  
Lund University

## Exploring autonomous USVs - Planning and Manoeuvring

The usage of USVs in a search-and-rescue scenario may be to assist in searching an area, either as a stand-alone vessel or in collaboration with other vessels. This collaboration calls for knowledge sharing and situational awareness together with automated planning to cover the search area in an efficient way between possibly heterogeneous agents/vessels. Furthermore, the concept of obstacle avoidance at sea is important in order to perform relocations in a safe manner. This includes performing preventive actions in compliance with maritime rules and regulations (COLREGs), which means that these rules must be incorporated both into the motion planning of the search task and as reactive patterns for replanning with respect to other sea traffic entering the area.

Wingqvist, Birgitta  
Lund University

WASP | WALLENBERG AI,  
AUTONOMOUS SYSTEMS  
AND SOFTWARE PROGRAM

## Exploring autonomous USVs – Planning and Manoeuvring



LUND  
UNIVERSITY

Birgitta Wingqvist, Lund University  
Department of Automatic Control



SAAB

### Motivation & Research goals

The usage of unmanned surface vessels, USVs, in a search-and-rescue scenario may be to assist in searching an area, either as a stand-alone vessel or in collaboration with other vessels. This collaboration calls for knowledge sharing and situational awareness together with automated planning to cover the search area in an efficient way between possibly heterogeneous agents/vessels. The goal of having unmanned vessels assisting in this scenario is to increase efficiency and to relieve the humans involved. Having the vessels collaborating on low-level tasks, leaves the operator out of continuous control and the human-machine system becomes an autonomous hybrid discrete event system. As communication bandwidth is limited, local low-level processing of information is desirable [Lager, 2021] before transmitting as well as automating the continuous control.

Furthermore, the concept of obstacle avoidance at sea is important in order to perform relocations in a safe manner. This includes performing preventive actions in compliance with the traffic rules at sea, COLREGs, which means that these rules must be incorporated both into the motion planning of the search task and as reactive patterns for replanning with respect to other sea traffic entering the area. Previous work in COLREGs-compliant trajectory planning can be found in [Bergman et. al., 2020].



Figure 1. Heterogeneous vessels at WARA-PS 2020.

### Traffic rules at sea

The Convention on the International Regulations for preventing Collisions at Sea, by International Maritime Organization (IMO), COLREGs, defines the traffic rules at sea. Given that the vessel is moving along a straight path, modified manoeuvres can be suggested by using predictive control [Hagen et. al., 2018]. When another vessel is detected the situation and applicable rule is identified. The movement of the own vessel is predicted with a discrete set of modifications in speed and direction, creating a set of trajectories for evaluation. For each trajectory, costs are associated with collision risk, COLREGs compliance as well as on the control signal. The cost is to be minimized. The scenarios were tested experimentally at WARA-PS 2020, see Figure 2.

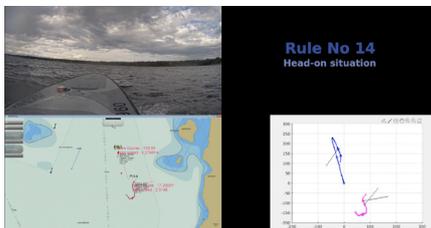


Figure 2. Initial test of COLREGs compliant manoeuvres at WARA-PS 2020 with two vessels involved. Here, a head-on situation where the vessels are both suggested to turn starboard (right). The upper left image shows the situation as seen from the magenta track in the lower right plot.

### Path following

For path following, a Model Predictive Controller, MPC, with a simplified model was developed and compared with a non-model-based PID controller. The model-based controller is using the Serret-Frenet frame. Two versions of the MPC were developed, one using a non-linear model and one using a linearised version for speeding up the optimization. The PID controller is based on a Line-of-Sight, LOS, controller [Fossen, 2021] with a fixed look-ahead distance,  $l$ . This means that the reference heading value is the bearing to a point on the path located a distance  $l$  ahead on the intended path as presented in Figure 3a. Comparisons in simulation show that the path-following performance is similar under the simulated conditions.

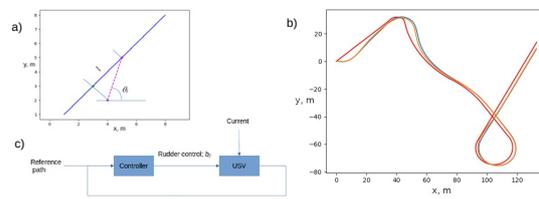


Figure 3. a) Look-ahead point. A point at a look-ahead distance,  $l$ , ahead in the LOS controller the point to aim at each moment. b) and c) Path-following controller and resulting path in simulation. Reference track in red, non-linear MPC in blue and MPC using a simplified linearised version in yellow. The prediction horizon is 10 s and a sea current is added as a disturbance.

### References

1. Lager M., Digital Cognitive Companions for Marine Vessels: On the Path Towards Autonomous Ships, PhD Thesis, Print ISBN 978-91-7895-608-1, Lund University, 2021
2. Bergman, K., Ljungqvist, O., Linder, J., & Axehill, D., A COLREGs-Compliant Motion Planner for Autonomous Maneuvering of Marine Vessels in Complex Environments. arXiv preprint arXiv:2012.12145, 2020
3. Hagen I. B., Kufoalor D. K. M., Brekke E. F., Johansen T. A., MPC-based Collision Avoidance for Existing Marine Vessel Guidance Systems; IEEE International Conference on Robotics and Automation (ICRA), May 21-25, Brisbane, Australia, 2018
4. Fossen, T. I., Handbook of Marine Craft Hydrodynamics and Motion Control, Wiley, Hoboken, NJ, 2nd edition, 2021

Xie, Yiping  
KTH

## Bathymetry Reconstruction from Sidescan Sonar

In recent years there has been increasing interest in the use of sidescan to reconstruct bathymetric maps. Bathymetric maps are usually constructed with high-end multibeam echo sounders (MBES), which are normally mounted on survey vessels or large autonomous underwater vehicles (AUVs). However, such MBES are relatively large and expensive compared to sidescan sonars, thus not suitable for smaller AUVs. Furthermore, sidescans generally have a wider swath range than multibeam and can produce images with a much higher resolution. If information about the seafloor's slope changes can be inferred from sidescan images, a low-cost and efficient method to construct high-resolution bathymetric maps would result and be of great benefit to many applications using smaller AUVs.

Xie, Yiping  
KTH

# Bathymetry Reconstruction from Sidescan Sonar

Yiping Xie, KTH Royal Institute of Technology

Robotics Perception and Learning Lab

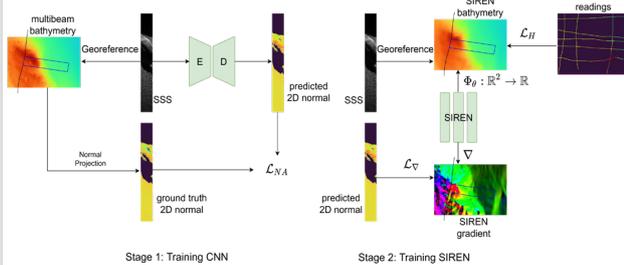
Main supervisor: John Folkesson



## Motivation & Research Goals

In recent years there has been increasing interest in the use of sidescan to reconstruct bathymetric maps. Bathymetric maps are usually constructed with high-end multibeam echo sounders (MBES), which are normally mounted on survey vessels or large autonomous underwater vehicles (AUVs). However, such MBES are relatively large and expensive compared to sidescan sonars, thus not suitable for smaller AUVs. Furthermore, sidescans generally have a wider swath range than multibeam and can produce images with a much higher resolution. If information about the seafloor's slope changes can be inferred from sidescan images, a low-cost and efficient method to construct high-resolution bathymetric maps would result and be of great benefit to many applications using smaller AUVs.

## Methods



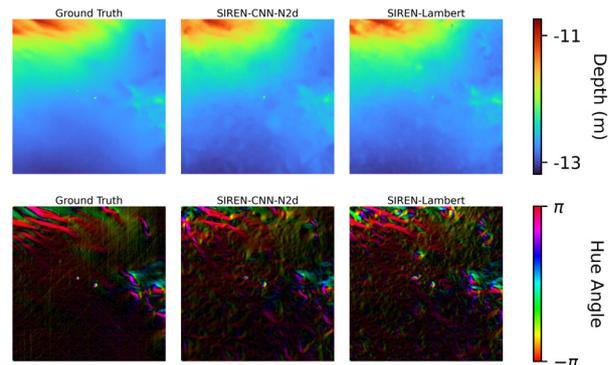
Overview: The method has two stages. Stage 1: Given the sidescan sonar (SSS) data and the mesh formed from MBES data, we can georeference the sidescan images to create training data for learning the inverse sensor model that estimates the surface normal from sidescan with a CNN. Such CNN (with an encoder "E" and a decoder "D") is trained with normal-aware loss, which can focus both the high normal area and the low normal area. Stage 2: Once the CNN is trained, we use the surface normal predicted from the CNN to constrain the gradient of the SIREN MLP, i.e., the to be estimated bathymetry. At the same time, we also use the altimeter readings to constrain the SIREN directly.

Instead of representing the bathymetry with explicit methods (meshes or grids), we use a function  $\Phi_\theta$ , parameterized by a fully connected neural network with parameters  $\theta$  to represent the bathymetry. The fully connected neural network is a variant of MLP with sinusoidal activation functions, known as SIREN [1], mapping 2D spatial coordinates to the corresponding seafloor height,  $\Phi_\theta: \mathbb{R}^2 \rightarrow \mathbb{R}$ , the same as in our previous work [2]. Note that the representation is continuous, differentiable and capable of producing high-quality derivatives with respect to the 2D spatial coordinates, allowing us to supervise the derivatives during training. In addition to constraining the derivatives of the bathymetry, we also need some boundary conditions which come from the altimeter readings. Here we assume to have access to high-quality navigation data.

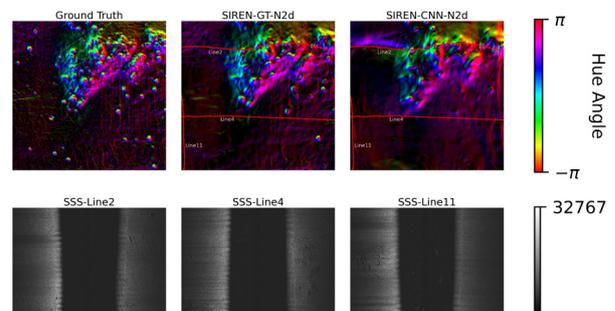
## References

- [1] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," Proc. NeurIPS, vol. 33, 2020.
- [2] N. Bore and J. Folkesson, "Neural shape-from-shading for survey-scale self-consistent bathymetry from sidescan-submitted to ieejoe," IEEE J. Ocean. Eng.

## Selected Results



The results for the final bathymetry maps from Dataset II, another place in freshwater. Top: bathymetry maps (approximately 300m×300m); bottom: gradient. Column 1, 2 and 3: MBES Ground Truth, the SIREN trained with CNN normals from all the sidescan lines, the SIREN trained with Lambertian models with all the sidescan lines. Note that the CNN only has been trained on Dataset I, a different natural environment in seawater. The results show one application of the proposed method, that is, reconstructing survey-scale bathymetry from another place with the same ship and equipment once the CNN is trained in one place. This then shows that once trained the approach will work with only a sidescan sonar.



Zoomed-in areas including a ridge (approximately 50m×50m). Top: gradient of the bathymetric map in HSV; left: ground truth gradient of our bathymetry from MBES data; middle: using normal computed from the MBES data to constrain the SIREN; right: using normal predicted from the CNN to constrain the SIREN. Bottom: corresponding SSS images from different lines.

Xu, Xuechun  
KTH

## End-to-end DNA Basecaller with semiHMM-DNN

DNA-basecalling is a modeling task that aims to solve a sequential classification problem. It has unique challenges: First one is that the data has no alignment information between the nucleotide and current measurements. This requires the model to be in an end-to-end fashion. Other challenges come from the properties of the DNA, e.g. distinguish homopolymers, and this problem sits on a large state space. We propose a suitable solution which is a hybrid structure that combines the neural network with the bayesian graphic model. This model is both trained and decoded end-to-end. According to the current experiments and other publications, this hybrid approach performs around 4% better (with an accuracy of ~94%) compared with using structures like the LSTM or TRANSFORMERS which will rely on the neural network itself to learn everything.



# END-TO-END DNA BASECALLER

## WITH SEMIHMM-DNN

Xuechun Xu

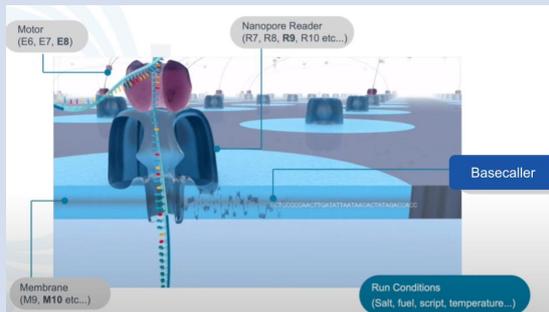
chunx@kth.se

Department of Information Science and Engineering



### DNA Base-calling

#### ► Nanopore base-caller

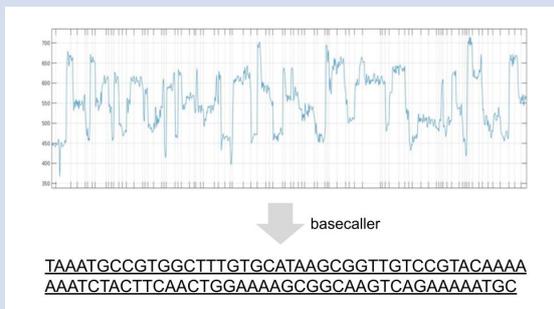


#### ► Basecalling - a challenge for engineers

- “Basecalling, the computational process of translating raw electrical signal to nucleotide sequence, is of critical importance to the sequencing platforms produced by Oxford Nanopore Technologies (ONT).”
- Wicket al., Performance of neural network basecalling tools for Oxford Nanopore sequencing, Genome Biology (2019)

### Problem and challenges

#### ► Translate current measurements into sequence of nucleotide

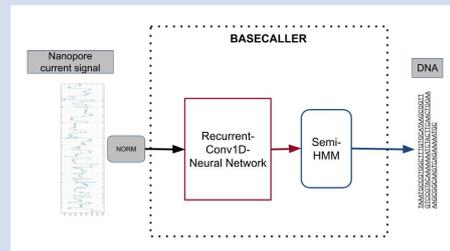


#### ► Challenges

- **End-to-end task**
  - Nucleotide sequence of length L corresponding to N current measurements ( $L > N$ )
  - Alignment between measurements and nucleotides is unknown
  - The basecalling task requires simultaneously decoding and alignment
- **Decode homopolymers**
  - Missing one gene in the genome can cause severe disease (human genome consists of 3 billion nucleotides)
  - In basecalling it is critical to distinguish 'AAAAA' and 'AAAAAAA', unlike in nature language processing where the accent is not important as long as the word is meaningful.
- **Large state space**
  - The measured current is affected by multiple nucleotides, normally consider 5, which leads to number of  $4^5 = 1024$  states

### Method Proposal

#### ► A hybrid model consists of semi-HMM and deep neural network

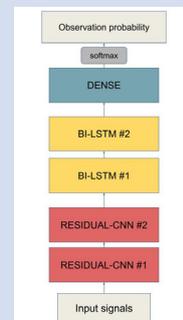


#### ► Semi Hidden Markov Model - for alignment

- **Explicit-duration HMM**
  - Denote 'K-mer' to represent the K consecutive nucleotide bases affecting the measurement
  - Ad hoc duration variable  $d \in 1, 2, \dots, D$ , which represents the number of measurements corresponding to a certain K-mer
  - The total state space extends to of size  $4^K \times D$
  - Algorithms like Viterbi and Baum-Welch (EM) are applicable for decoding and training

#### ► DNN - for decoding

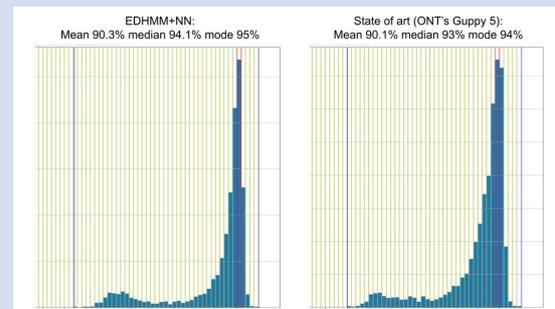
- Consists of two blocks of residual CNN and two blocks of bi-directional LSTM, followed by one dense layer with softmax transition
- 16M parameters
- Can be used as sequential classification network separately
- The output of this DNN is used as emission/observation probabilities in the semi-HMM
- Require to be trained simultaneously with the semi-HMM



### Current results

#### ► Sequence accuracy evaluated with edit distance

- Trained with NSC Berzelius SuperPOD



Zhu, Xiaomeng  
KTH

## Automatic quality inspection based on Computer Vision and Unsupervised Domain Adaptation

Deep learning-based computer vision technologies could offer a possible solution for automatic quality inspection with their outperformance. However, most deep learning methods currently implemented in production are based on supervised learning, which requires a large amount of labeled training data that is time-consuming and expensive to collect in the industry.

This research aims to solve this problem by utilizing unsupervised domain adaptation (UDA) models. The models can be trained on annotated synthetic images generated from CAD models and unannotated images captured from cameras. They achieve promising results on an industry case study of pedal car front-wheel assembly. Furthermore, since the models do not require manually annotated images, they are less time-consuming to implement in production.

# Automatic quality inspection based on Computer Vision and Unsupervised Domain Adaptation



Xiaomeng Zhu, Scania & KTH  
Robotics, Perception, and Learning

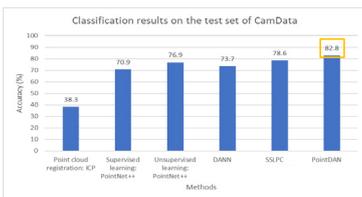


## Motivation & Research goals

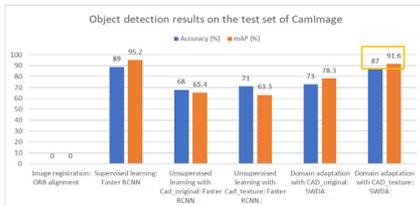
Deep learning-based computer vision technologies with their outperformance could offer a possible solution for automatic quality inspection. However, most deep learning methods currently implemented in production are based on supervised learning, which requires a large amount of labeled training data that is **time-consuming** and **expensive** to collect in the industry. This research aims to solve this problem by utilizing **unsupervised domain adaptation (UDA) models**. The models can be **trained on annotated synthetic images generated from CAD models and unannotated images captured from cameras**. Since the models do not require any manually annotated images, they are less time-consuming to implement in production.

## Selected Results

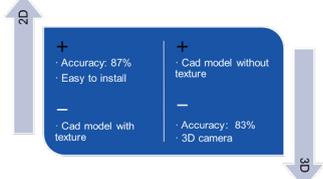
Research in 3D point cloud UDA



Research in 2D image UDA



2D & 3D results comparison



Contribution:

- End-to-end methods for automatic assembly quality inspection based on unsupervised domain adaptation with promising results.
- 2D and 3D.
- Easy to adapt to different production projects.
- Does not require does not require manual labeling work.

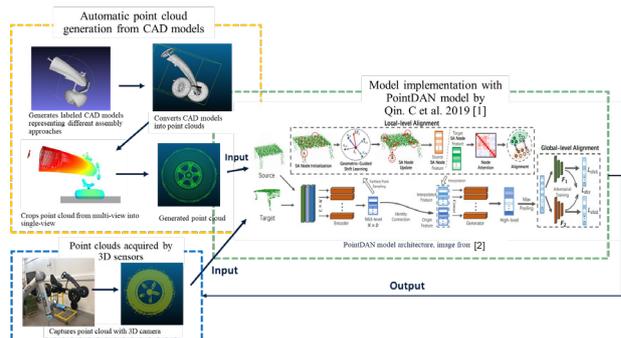
## References

[1] X. Zhu, H. Manamasa, J.L. Jiménez Sánchez, A. Maki, L. Hanson, Automatic assembly quality inspection based on an unsupervised point cloud domain adaptation model, *Procedia CIRP*, 104 (2021) 1801–1806.  
 [2] C. Qin, H. You, L. Wang, C.-C.J. Kuo, Y. Fu, PointDAN: A Multi-Scale 3D Domain Adaption Network for Point Cloud Representation, *ArXiv:1911.02744 [Cs]*. (2019).  
 [3] X. Zhu, A. Maki, L. Hanson, Unsupervised domain adaptive object detection for assembly quality inspection, *CIRP ICME '21 Virtual Conference*, 14-16 July, *Procedia CIRP*, Elsevier, ISSN: 2212-8271  
 [4] K. Saito, Y. Ushiku, T. Harada, K. Saenko, Strong-Weak Distribution Alignment for Adaptive Object Detection, *ArXiv:1812.04798 [Cs]*. (2019).

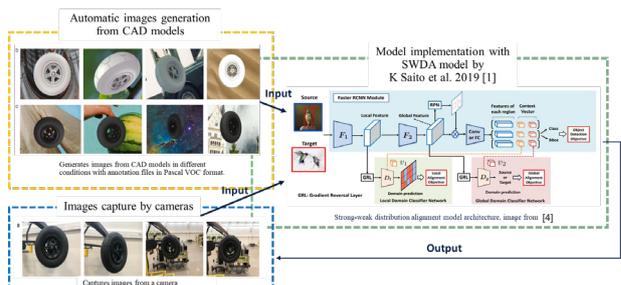
## Method

- Automatic data generation from CAD models**
  - Data can represent difference assembly end positions.
  - Data have label.
- Model implementation**
  - Choose an unsupervised domain adaptation model to implement.
  - Comparing the unsupervised domain adaptation method with:
    - Non-deep learning method;
    - Supervised learning method;
    - Unsupervised learning method.

Research on 3D point cloud UDA [1]



Research on 2D image UDA [3]



## Pose Estimation from RGB Images of Highly Symmetric Objects using a Novel Multi-Pose Loss and Differential Rendering

We propose a novel multi-pose loss function to train a neural network for 6D pose estimation, using synthetic data and evaluating it on real images. Our loss is inspired by the VSD (Visible Surface Discrepancy) metric and relies on a differentiable renderer and CAD models. This novel multi-pose approach produces multiple weighted pose estimates to avoid getting stuck in local minima. Our method resolves pose ambiguities without using predefined symmetries. It is trained only on synthetic data. We test on real-world RGB images from the T-LESS dataset, containing highly symmetric objects common in industrial settings.

We show that our solution can be used to replace the codebook in a state-of-the-art approach. So far, the codebook approach has had the shortest inference time in the field. Our approach reduces inference time further while a) avoiding discretization, b) requiring a much smaller memory footprint and c) improving pose recall.

# Pose Estimation from RGB Images of Highly Symmetric Objects using a Novel Multi-Pose Loss and Differential Rendering

Hampus Åström, Lund University, Computer Science



Adapted from paper by Stefan Hein Bengtson and Hampus Åström, Thomas B. Moeslund, Elin A. Topp, Volker Krueger - IROS 2021



AALBORG UNIVERSITY  
DENMARK

We propose a novel multi-pose loss function to train a neural network for 6D pose estimation, using synthetic data and evaluating it on real images. Our loss is inspired by the VSD (Visible Surface Discrepancy) metric and relies on a differentiable renderer and CAD models. This novel multi-pose approach produces multiple weighted pose estimates to avoid getting stuck in local minima. Our method resolves pose ambiguities without using predefined symmetries. It is trained only on synthetic data. We test on real-world RGB images from the T-LESS dataset, containing highly symmetric objects common in industrial settings.

The code for our project is available at <https://github.com/shbe-aau/multi-pose-estimation>

## Overview

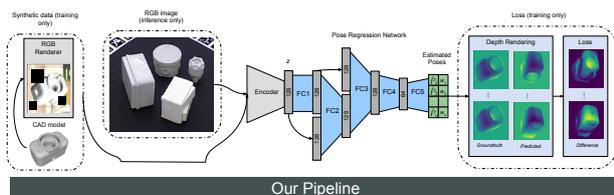
We propose an adaptation of the 6D pose estimation approach in [1,2], that relies on an autoencoder for feature extraction in a codebook-based approach. By replacing their codebook with a neural network and utilizing differential rendering [3], we provide a solution that:

- has improved pose recall when tested on the T-LESS dataset.
- is faster at inference
- has a significantly smaller memory footprint

Our solution does not require discretizing poses and it is therefore easily extendable. It is trained on synthetic RGB images (no depth information required) rendered from CAD models or reconstructions and requires no labelled data or predefined symmetries.

## Method

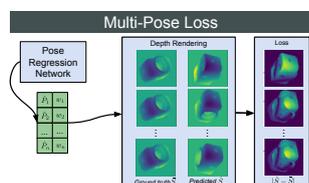
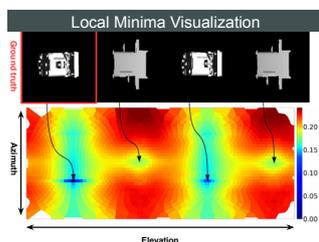
Our pose estimation method utilizes synthetic data (RGB images from CAD models) and a pre-trained encoder to train a regression network. With differential rendering we manage ambiguities with symmetries. Multiple weighted pose estimates overcome problems with local minima that stem from low output dimensionality.



Our Pipeline

Error is measured by a difference in differentially rendered depth maps for each pose versus the ground truth. This inherently handles symmetries as it relies on physical appearance.

Each pose estimate is describes the pose as a rotation matrix. Due to the low dimensionality of the output, steepest descent get problems with local minima. Our network outputs multiple weighted pose estimates to avoid those problems.



The final loss is a weighted average of the depth map differences plus a term that forces the pose estimates to be spread out.

$$= L_{\text{pose}}(\hat{P}) + \sum_{i=1}^n L_{\text{single}}(\hat{S}_i, \bar{S}) \cdot (\gamma + w_i)$$

## Performance

On the highly symmetric dataset T-LESS [4], we achieve better recall by replacing the codebook in [1,2] with our network.



Our method operates at similar inference speed while drastically reducing the amount of GPU memory needed. We can handle cluttered scenes with occlusions as seen on the right.

The novel multi-pose loss significantly improves recall and we perform better on objects with continuous symmetries..

Inference Speed				Recall		
On a GTX 1060 GPU:						
6.2 ms (ours)						
7.0 ms (codebook)						
Memory Usage						
Encoder	Codebook	Pose Regression Network	Total	Obj.	Codebook	Ours
06	15 MB	307.45 MB	333 MB	01	51.88	51.84 ± 2.8
07	15 MB	307.45 MB	333 MB	02	62.87	63.74 ± 1.8
08	15 MB	307.45 MB	333 MB	03	56.00	71.53 ± 3.3
09	15 MB	307.45 MB	333 MB	04	66.00	62.66 ± 3.5
10	15 MB	307.45 MB	333 MB	05	77.18	80.82 ± 0.3
11	15 MB	307.45 MB	333 MB	06	68.04	66.71 ± 4.6
12	15 MB	307.45 MB	333 MB	07	65.18	65.68 ± 4.9
13	15 MB	307.45 MB	333 MB	08	63.11	61.21 ± 0.8
14	15 MB	307.45 MB	333 MB	09	68.96	55.56 ± 0.5
15	15 MB	307.45 MB	333 MB	10	58.55	54.14 ± 2.0
16	15 MB	307.45 MB	333 MB	11	52.15	51.48 ± 2.4
17	15 MB	307.45 MB	333 MB	12	62.19	56.58 ± 1.6
18	15 MB	307.45 MB	333 MB	13	63.56	64.21 ± 5.0
19	15 MB	307.45 MB	333 MB	14	57.29	63.01 ± 1.2
20	15 MB	307.45 MB	333 MB	15	64.91	66.37 ± 3.8
21	15 MB	307.45 MB	333 MB	16	75.82	73.16 ± 2.7
22	15 MB	307.45 MB	333 MB	17	76.62	77.72 ± 0.9
23	15 MB	307.45 MB	333 MB	18	71.26	62.71 ± 2.0
24	15 MB	307.45 MB	333 MB	19	62.97	63.85 ± 1.2
25	15 MB	307.45 MB	333 MB	20	40.71	51.19
26	15 MB	307.45 MB	333 MB	21	43.25	43.31 ± 1.4
27	15 MB	307.45 MB	333 MB	22	38.15	32.03 ± 0.5
28	15 MB	307.45 MB	333 MB	23	39.18	56.68 ± 1.1
29	15 MB	307.45 MB	333 MB	24	58.97	61.93 ± 3.3
30	15 MB	307.45 MB	333 MB	25	69.86	63.98 ± 1.6
31	15 MB	307.45 MB	333 MB	26	57.94	58.87 ± 2.3
32	15 MB	307.45 MB	333 MB	27	68.09	77.62 ± 1.2
33	15 MB	307.45 MB	333 MB	28	68.06	73.33 ± 1.3
34	15 MB	307.45 MB	333 MB	29	76.43	80.67 ± 0.7
35	15 MB	307.45 MB	333 MB	30	77.81	83.41 ± 2.1
36	15 MB	307.45 MB	333 MB	mean	57.47	60.09 ± 0.4
Multi-Pose Improvement				All objects		
1 pose	10 poses	improvement		All	Codebook	Ours
57.37 ± 1.6	67.23 ± 2.7	9.86		mean	60.77	62.34 ± 0.9
50.62 ± 0.9	59.51 ± 0.3	8.89				
53.10 ± 0.6	62.34 ± 0.9	9.24				
mean	62.97	63.85 ± 1.2				

## Current and Future Work

In our paper the regression network determines the rotation from a cropped image. Translation can then be determined from the bounding box. In our current work we are extending the regression network to also determine the translation directly by providing it with bounding box information. We do this by adding additional outputs. In a similar way it could be possible to do pose predictions for flexible objects by adding additional degrees of freedom to the output.

The current version of the code has a shared autoencoder for multiple objects, but individual regression networks for each object type. Those could possibly be merged to improve training time and scalability further.

## References

1. M. Sundermeyer, Z.-C. Marton, M. Durner, M. Brucker, and R. Triebel, "Implicit 3d orientation learning for 6d object detection from rgb images," in ECCV, September 2018.
2. M. Sundermeyer, M. Durner, E. Y. Puang, Z.-C. Marton, N. Vaskevicius, K. O. Arras, and R. Triebel, "Multi-path learning for object pose estimation across domains," in CVPR, June 2020.
3. N. Ravi, J. Reizenstein, D. Novotny, T. Gordon, W.-Y. Lo, J. Johnson, and G. Gkioxari, "Pytorch3d," <https://github.com/facebookresearch/pytorch3d>, 2020.
4. T. Hodaň, P. Haluza, Š. Obdržálek, J. Matas, M. Lourakis, and X. Zabulis, "T-LESS: An RGB-D dataset for 6D pose estimation of texture-less objects," WACV, 2017.

# WASP WINTER CONFERENCE 2022

## POSTER CATALOGUE 3/4

### AUTONOMOUS SYSTEMS [AS]

Author	Pages
Ahmad, Faseeh.....	102 A+B
Alshnakat, Anoud.....	103 A+B
Baravdish, Gabriel.....	104 A+B
Batkovic, Ivo.....	105 A+B
Brissman, Emil.....	106 A+B
Bruns, Leonard.....	107 A+B
Ceylan, Ciwan.....	108 A+B
Charitidou, Maria.....	109 A+B
Eryonucu, Cihan.....	110 A+B
Faris, Muhammed.....	111 A+B
Ferizbegovic, Mina.....	112 A+B
Forough, Javad.....	113 A+B
Fredriksson, Teodor.....	114 A+B
Gyllenhammar, Magnus.....	115 A+B
Hellander, Anja.....	116 A+B
Heskebeck, Frida.....	117 A+B
Hynén Ulfsjö, Carl.....	118 A+B
Iovino, Matteo.....	119 A+B
Jakobsson, Erik.....	120 A+B
Jensen, Maarten.....	121 A+B
Johansson, Simon.....	122 A+B
Johnander, Joakim.....	123 A+B
Jonnarth, Arvi.....	124 A+B
Kaalen, Stefan.....	125 A+B
Kampik, Timotheus.....	126 A+B
Khosravi, Hedieh.....	127 A+B
Krook, Jonas.....	128 A+B
Kullberg, Anton.....	129 A+B
Lapandić, Dženan.....	130 A+B
Larsson, Martin.....	131 A+B
Marta, Daniel.....	132 A+B
Mayr, Mattias.....	133 A+B
Moliner, Olivier.....	134 A+B
Mollevik, Iris.....	135 A+B
Narri, Vananda.....	136 A+B