Projects and Theses

Neurorobotics Lab
Autonomous Intelligent Systems Lab

Albert-Ludwigs-Universität Freiburg

General Information

- You can work on a project either alone or in small groups of 2-3 students.
- Typically, we only supervise MSc theses after a MSc project in our group.
- If you are interested, talk to us after this session or write us an email.

RL with Hard Constraints

- Pham et al. presented in their recent OptLayer paper an approach to include hard constraints in the reinforcement learning process (ICRA 2018).
- It could be really useful in terms of safety for RL in robotics.
- ➢ Implement OptLayer and evaluate it on some simulated environments (and probably on a real Jaco arm).

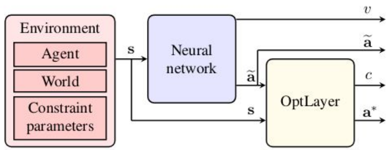

(a) Prediction-correction architecture.

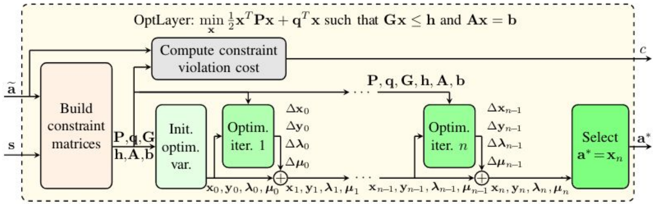

Multi-Task RL for StarCraft II

- DeepMind released a StarCraft II research environment in cooperation with Blizzard last year.
- Super complex task with a huge action space and a long time horizon.
- ➢ Project: Gain experience with the simulator and evaluate recent RL approaches on some simpler tasks.
- ➢ Thesis: Hierarchical/Compositional Multi-Task RL.

Deep Scheduler

- Use RL for deciding which experiments to continue based on their recent performance and costs.
- It gets meta! Agent needs to learn to explore and exploit.

➢ Tackling basic questions in this context:
  - State spaces for generalization over budgets
  - Recursive networks for generalization to many experiments
  - And more

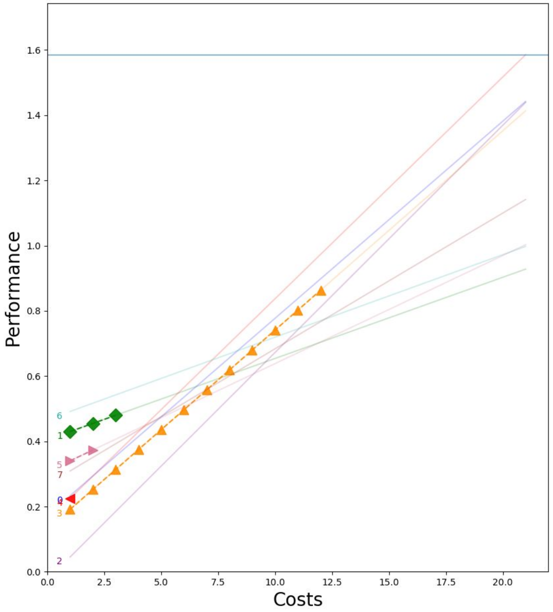

Anomaly Detection in EEG-Data

Automated epileptic seizure detection

- While traditional approaches use supervised learning, you will try to apply unsupervised learning
- Detect anomalies in intracranial EEG data using variational RNNs
- Work in a team with experts of the Epilepsy Center

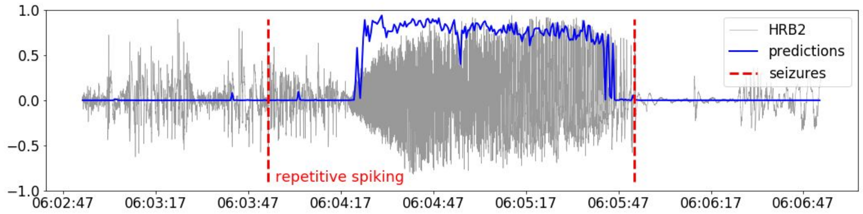

RL for Autonomous Driving

- Our group (in cooperation with AIS) already supervised three teams in the Audi Autonomous Driving Cup (3/10, 5/10, 5/13).
- A miniature car has to drive through a city and solve some tasks autonomously.
- We may apply for it this year again.

➢ We would like to achieve a more adaptive and robust behaviour and DL and RL techniques could be of help.
➢ Project based on a real system.
➢ AD is currently a hot topic in the real world.
➢ You will work in a team of five.

K-FAC in DDPG and PPO

- K-FAC is an optimizer that considers curvature.
- It has been employed in TRPO, leading to ACKTR (Wu et al., NIPS 2017).
- Evaluating ACKTR on a simulated robot environment was a recent MSc Project at AIS.

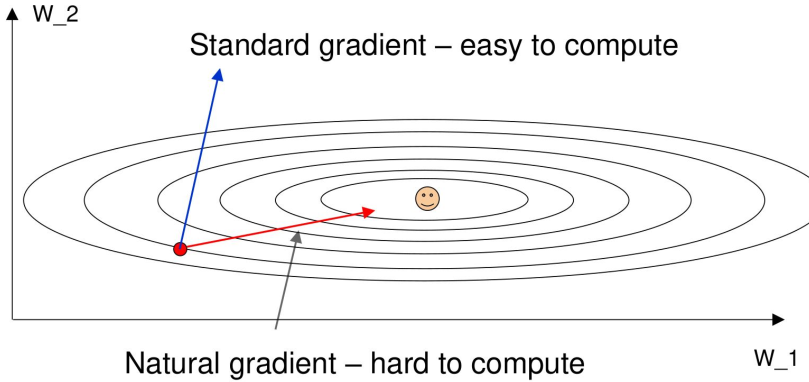

➢ Combine K-FAC with other algorithms, such as DDPG or PPO.

Time-Contrastive Networks: Self-Supervised Learning from Video

- Learning manipulation policies from video
- Learn embedding space for robotic imitation

➢ Collect data, train triplet network and evaluate method on real robot

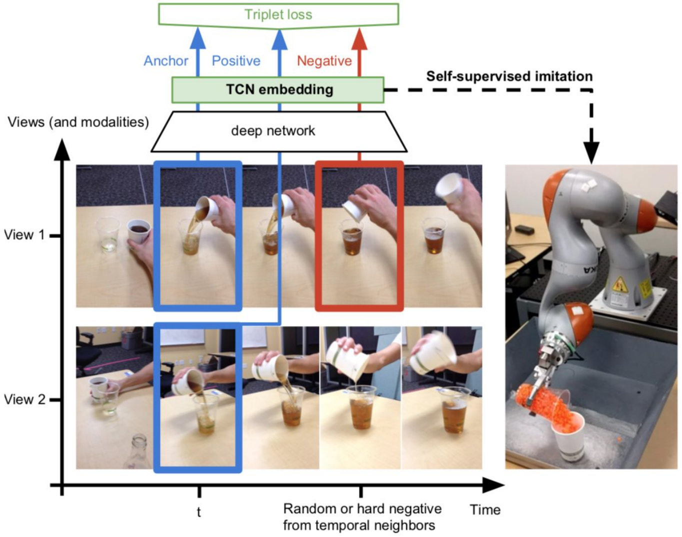

Contact: Jingwei Zhang zhang@informatik.uni-freiburg.de,
Oier Mees meeso@informatik.uni-freiburg.de

One-Shot Imitation from Observing Humans via Meta-Learning

- Learn from human video demonstrations, one-shot domain adaptation to infer robot policy
- Learn embedding space for robotic imitation
- ➢ Get familiar, with meta-learning method, train policies and test on real robot

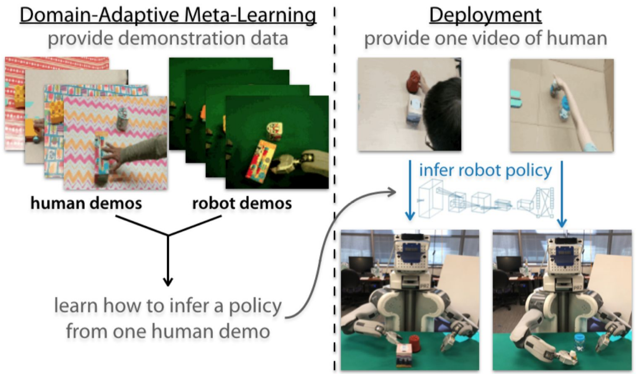

Fig. 1. After meta-learning with human and robot demonstration data, the robot learns to recognize and push a new object from one video of a human.

Contact: Jingwei Zhang zhang@informatik.uni-freiburg.de

UNI FREIBURG

Model Identification of Objects for Learning Manipulation Policies

- Learn mechanical properties of objects through physical interaction
- Identify motion-models
- Use model to learn manipulation policies
- ➢ Closed-loop physics simulation, model-based reinforcement learning

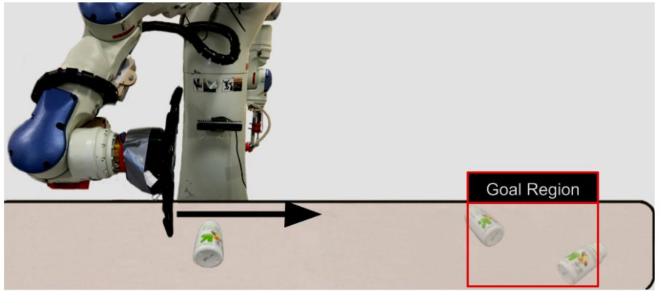

Fig. 8: The task in IV-C to push the object to the other side of the table.

Contact: Andreas Eitel eitel@informatik.uni-freiburg.de

Learning from Demonstration for Furniture Assembly

- Group of 2 students
- Learning from Demonstration with e.g. VR device, 3D joysticks
- Learn manipulation policies (e.g. peg-insertion, drilling) for assembly of IKEA furniture

➢ Collect human demonstrations, Sim2real transfer to perform manipulation actions on real robot, domain randomization

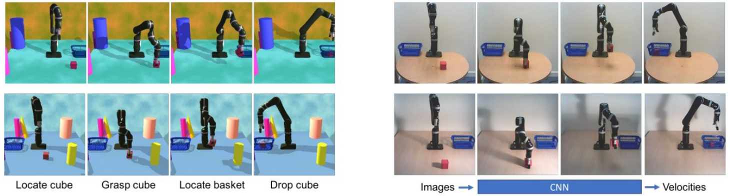

Figure 1: Our approach uses simulation to collect a series of control velocities to solve a multi-stage task. This is used to train a reactive neural network controller which continuously accepts images and joint angles, and outputs motor velocities. By using domain randomisation, the controller is able to run in the real world without having seen a single real image.

Contact: Andreas Eitel eitel@informatik.uni-freiburg.de

Estimating Odometry from LiDAR scans

- Learn the relative pose between two consecutive LiDAR scans
- Investigate different architectures with different input modalities
- Evaluate on KITTI benchmark and test on real robot

Contact: Ayush Dewan dewan@informatik.uni-freiburg.de, Noha Radwan radwann@informatik.uni-freiburg.de

UNI FREIBURG

Loop Closure Detection using LiDAR scans

- Detecting Loop closures is essential for any SLAM system
- The task is to identify if the robot has returned to a previously visited location
- Problem is challenging because of the symmetries in the environment and the change in viewing angles
- ➢ Train a Siamese network and evaluate on a real robot

Contact: Ayush Dewan dewan@informatik.uni-freiburg.de, Noha Radwan radwann@informatik.uni-freiburg.de

What Lies Behind?

Contact: Abhinav Valada
valada@cs.uni-freiburg.de

- In mobile robotics, generation of accurate maps is encumbered by the presence of ephemeral objects such as vehicles or pedestrians
  - ➢ Develop an approach based on Cycle-Consistent Adversarial networks to
    - Detect ephemeral objects in 3D scans
    - Remove detected objects
    - Predict what could be behind and fill gaps in a semantically reasonable way

- Check out Adobe DeepFill which does something similar for 2D images

3D scan | Detected cars | Removed cars | Predict what could be behind?

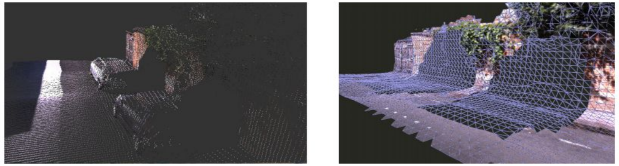

Existing approaches use geometry-based methods to fit triangular meshes but do not predict texture

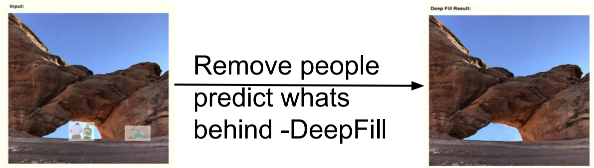

Self-Supervised Terrain Learning

Contact: Abhinav Valada
valada@cs.uni-freiburg.de

- Predicting the type of terrain is a prerequisite for mobile robot navigation
  - Acquiring labeled training data is a very tedious task!!
  - Utilize our audio-based terrain classification CNN to provide labels for online learning of a visual classifier

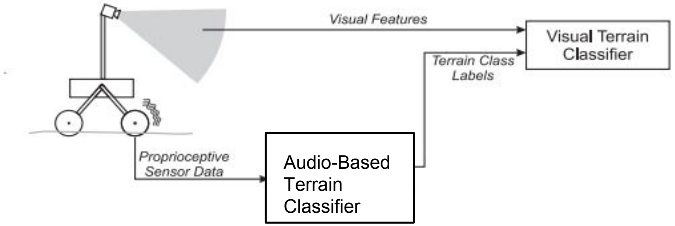

During Training

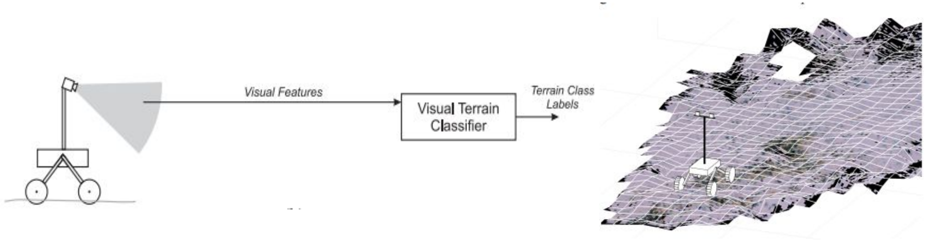

During Inference

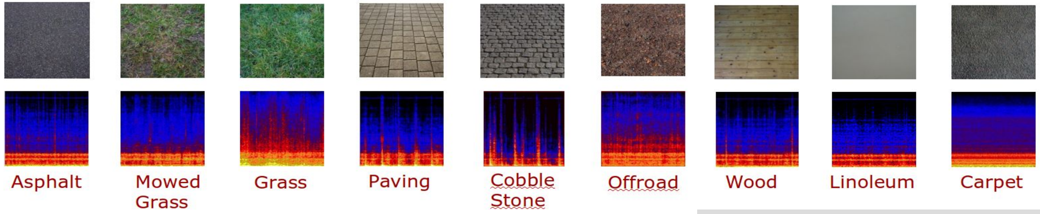

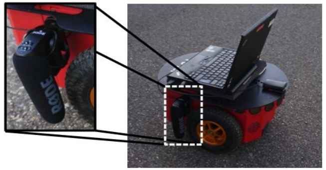

P3-DX

Behavior prediction in LiDAR scans

- Learn to track vehicles in LiDAR scans
- Learn motion-models from previous time-steps
- Use the information to predict the behavior for future time-steps

➢ Evaluate on the KITTI benchmark and r robot data

(a) Camera (b) LiDAR

Contact: Noha Radwan radwann@informatik.uni-freiburg.de, Daniel Buescher buescher@informatik.uni-freiburg.de

Robot Bring Me (the leftmost) Beer

- Locate objects based on natural language
- Use semantic and spatial relations
- Evaluate on PR2 robot
- Combine Deep Learning methods
  - ➢ Image captioning
  - ➢ Object detection
  - ➢ Spatial Grounding with LSTMs

Oier Mees meeso@informatik.uni-freiburg.de