Integrated Project

Robot Learning

TECHNISCHE UNIVERSITÄT DARMSTADT

Today’s agenda!

- Sales Pitch of the IP
- Project Timeline
- Project Proposals
- How to Apply

Sales Pitch

TECHNISCHE UNIVERSITÄT DARMSTADT

20-00-0628 IP Lernende Roboter, Teil 1+2

A Semester in the Life of a PhD!

Among the most important questions ever: continue the research road to a Ph.D. (=Dr.)?

The personal and professional advantages are enormous!

An exciting life:

- follow *your* ideas & dreams...
- actively acquire knowledge and refine it…
- enjoy international conferences and visits with collaborators around the world…

- However, it ain't for everybody!
- Your Master's thesis will already decide on your chances!
- Do you wanna figure out whether there is a researcher in you?

Literature Review

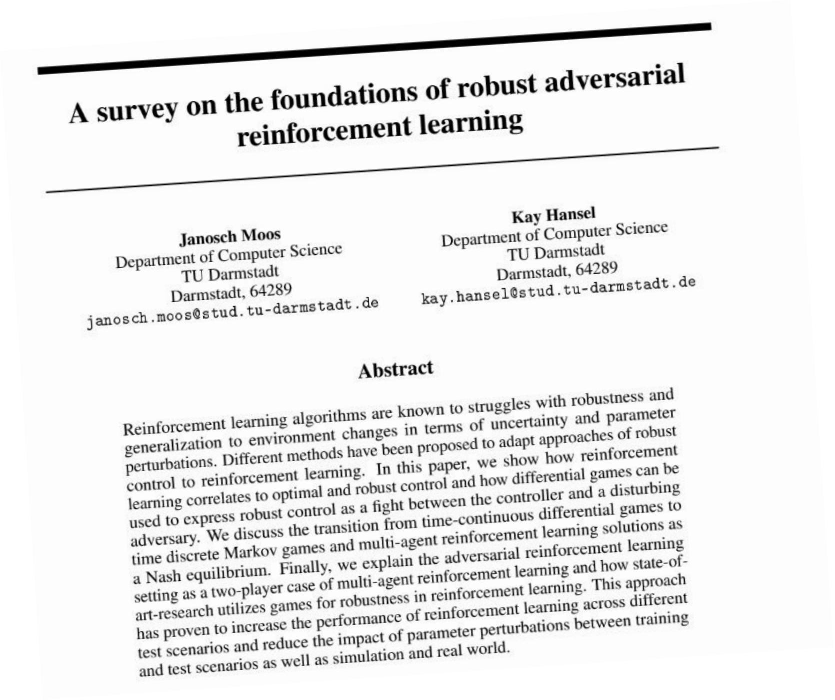

A survey on the foundations of robust adversarial reinforcement learning

Janosch Moos
Department of Computer Science
TU Darmstadt
Darmstadt, 64289
janosch.moos@stud.tu-darmstadt.de

Kay Hansel
Department of Computer Science
TU Darmstadt
Darmstadt, 64289
kay.hansel@stud.tu-darmstadt.de

Abstract

Reinforcement learning algorithms are known to struggles with robustness and generalization to environment changes in terms of uncertainty and parameter perturbations. Different methods have been proposed to adapt approaches of robust control to reinforcement learning. In this paper, we show how reinforcement learning correlates to optimal and robust control and how differential games can be used to express robust control as a fight between the controller and a disturbing adversary. We discuss the transition from time-continuous differential games to time discrete Markov games and multi-agent reinforcement learning solutions as a Nash equilibrium. Finally, we explain the adversarial reinforcement learning setting as a two-player case of multi-agent reinforcement learning and how state-of-art-research utilizes games for robustness in reinforcement learning. This approach has proven to increase the performance of reinforcement learning across different test scenarios and reduce the impact of parameter perturbations between training and test scenarios as well as simulation and real world.

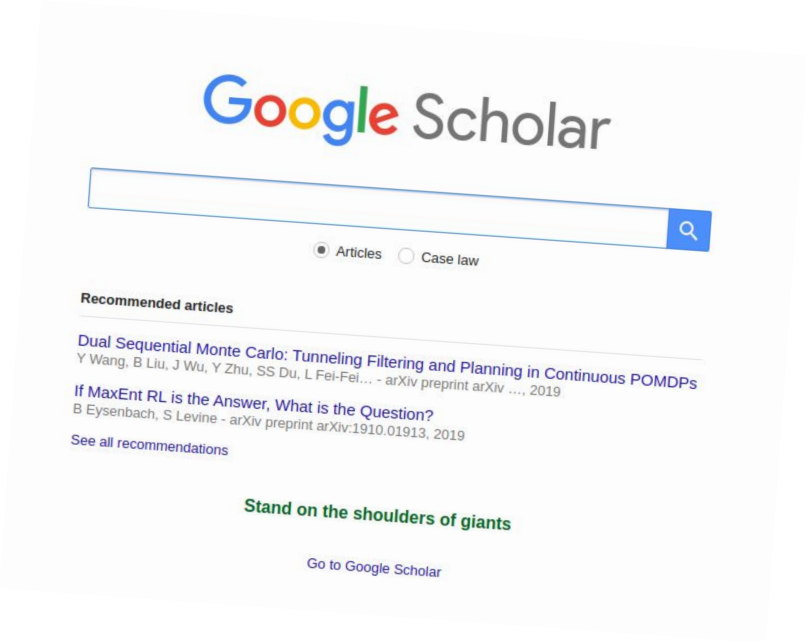

Theoretical Investigations

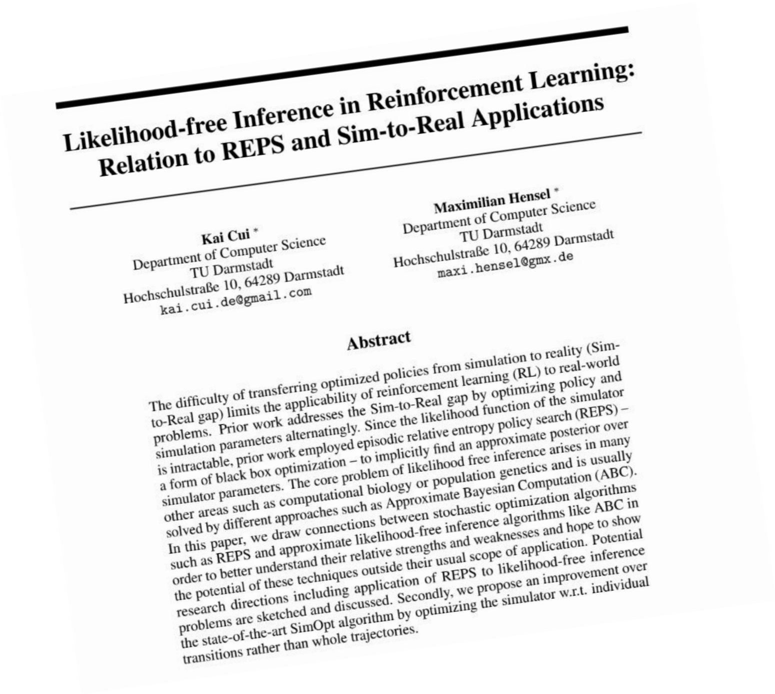

Likelihood-free Inference in Reinforcement Learning: Relation to REPS and Sim-to-Real Applications

Kai Cui *
Department of Computer Science
TU Darmstadt
Hochschulstraße 10, 64289 Darmstadt
kai.cui.de@gmail.com

Maximilian Hensel *
Department of Computer Science
TU Darmstadt
Hochschulstraße 10, 64289 Darmstadt
maxi.hensel@gmx.de

Abstract

The difficulty of transferring optimized policies from simulation to reality (Sim-to-Real gap) limits the applicability of reinforcement learning (RL) to real-world problems. Prior work addresses the Sim-to-Real gap by optimizing policy and simulation parameters alternatingly. Since the likelihood function of the simulator is intractable, prior work employed episodic relative entropy policy search (REPS) – a form of black box optimization – to implicitly find an approximate posterior over simulator parameters. The core problem of likelihood free inference arises in many other areas such as computational biology or population genetics and is usually solved by different approaches such as Approximate Bayesian Computation (ABC). In this paper, we draw connections between stochastic optimization algorithms such as REPS and approximate likelihood-free inference algorithms like ABC in order to better understand their relative strengths and weaknesses and hope to show the potential of these techniques outside their usual scope of application. Potential research directions including application of REPS to likelihood-free inference problems are sketched and discussed. Secondly, we propose an improvement over the state-of-the-art SimOpt algorithm by optimizing the simulator w.r.t. individual transitions rather than whole trajectories.

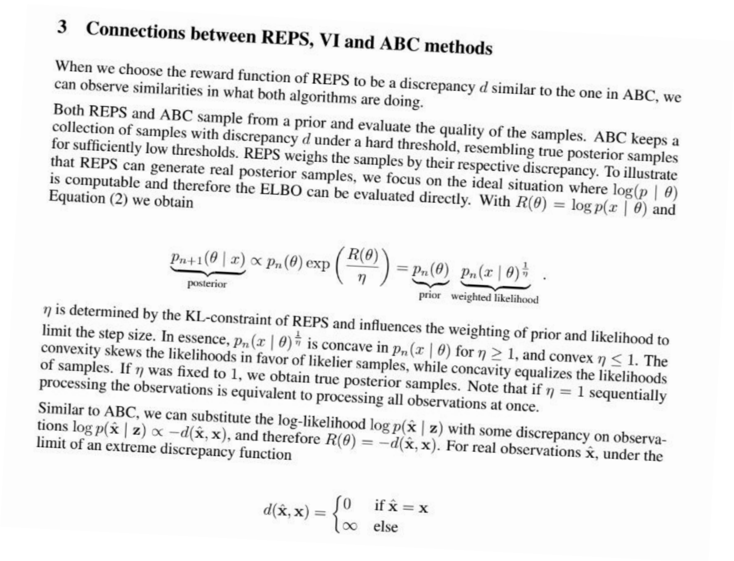

3 Connections between REPS, VI and ABC methods

When we choose the reward function of REPS to be a discrepancy d similar to the one in ABC, we can observe similarities in what both algorithms are doing.

Both REPS and ABC sample from a prior and evaluate the quality of the samples. ABC keeps a collection of samples with discrepancy d under a hard threshold, resembling true posterior samples for sufficiently low thresholds. REPS weighs the samples by their respective discrepancy. To illustrate that REPS can generate real posterior samples, we focus on the ideal situation where $\log(p \mid \theta)$ is computable and therefore the ELBO can be evaluated directly. With $R(\theta) = \log p(x \mid \theta)$ and Equation (2) we obtain

$$\underbrace{p_{n+1}(\theta \mid x)}_{\text{posterior}} \propto p_n(\theta) \exp\left(\frac{R(\theta)}{\eta}\right) = \underbrace{p_n(\theta)}_{\text{prior}} \underbrace{p_n(x \mid \theta)^{\frac{1}{\eta}}}_{\text{weighted likelihood}} .$$

η is determined by the KL-constraint of REPS and influences the weighting of prior and likelihood to limit the step size. In essence, $p_n(x \mid \theta)^{\frac{1}{\eta}}$ is concave in $p_n(x \mid \theta)$ for $\eta \geq 1$, and convex $\eta \leq 1$. The convexity skews the likelihoods in favor of likelier samples, while concavity equalizes the likelihoods of samples. If η was fixed to 1, we obtain true posterior samples. Note that if $\eta = 1$ sequentially processing the observations is equivalent to processing all observations at once.

Similar to ABC, we can substitute the log-likelihood $\log p(\hat{\mathbf{x}} \mid \mathbf{z})$ with some discrepancy on observations $\log p(\hat{\mathbf{x}} \mid \mathbf{z}) \propto -d(\hat{\mathbf{x}}, \mathbf{x})$, and therefore $R(\theta) = -d(\hat{\mathbf{x}}, \mathbf{x})$. For real observations $\hat{\mathbf{x}}$, under the limit of an extreme discrepancy function

$$d(\hat{\mathbf{x}}, \mathbf{x}) = \begin{cases} 0 & \text{if } \hat{\mathbf{x}} = \mathbf{x} \\ \infty & \text{else} \end{cases}$$

Implement Algorithms

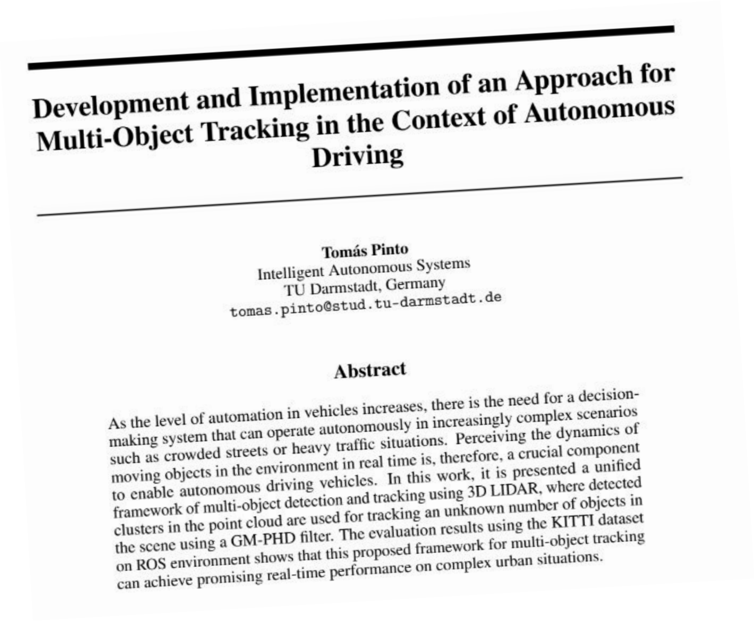

Development and Implementation of an Approach for Multi-Object Tracking in the Context of Autonomous Driving

Tomás Pinto
Intelligent Autonomous Systems
TU Darmstadt, Germany
tomas.pinto@stud.tu-darmstadt.de

Abstract

As the level of automation in vehicles increases, there is the need for a decision-making system that can operate autonomously in increasingly complex scenarios such as crowded streets or heavy traffic situations. Perceiving the dynamics of moving objects in the environment in real time is, therefore, a crucial component to enable autonomous driving vehicles. In this work, it is presented a unified framework of multi-object detection and tracking using 3D LIDAR, where detected clusters in the point cloud are used for tracking an unknown number of objects in the scene using a GM-PHD filter. The evaluation results using the KITTI dataset on ROS environment shows that this proposed framework for multi-object tracking can achieve promising real-time performance on complex urban situations.

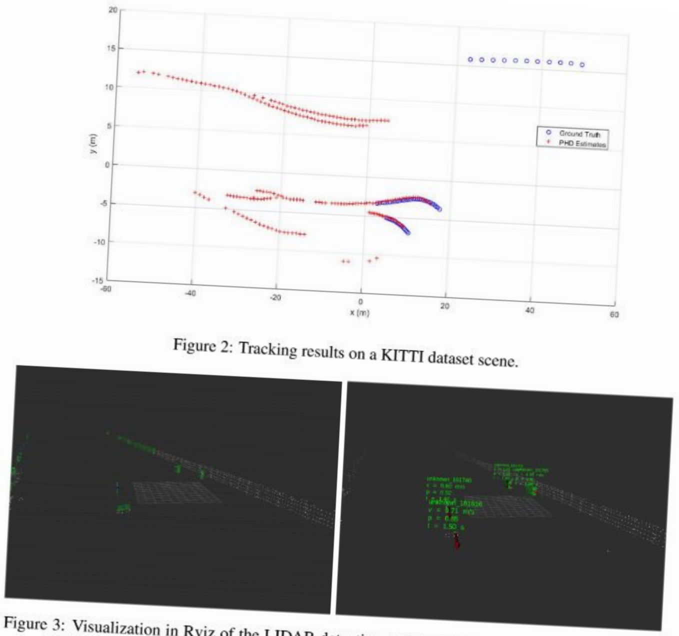

Figure 2: Tracking results on a KITTI dataset scene.

Figure 3: Visualization in Rviz of the LIDAR detection node output (Left) and the Tracker node output (Right).

Experimentation

Enhancing Intention Aware Movement Primitives

Artur Kruk, Yanhua Zhang
Department of Computer Science
Technology Universtiy of Darmstadt
Karolinenplatz 5, Darmstadt 64289
artur.kruk@stud.tu-darmstadt.de
yanhua.zhang@stud.tu-darmstadt.de

Abstract

In the future, robots will not only operate at a safe distance from humans but also have closer contact with them. To provide a comfortable human-robot cooperation, the robot is desirable to dynamically adapt its movements to not disrupt the workflow of its partner. The goal of this work is hereby to allow the robot to infer human's intention and predict its future movements based on a few early observations of human motion. Probabilistic movement primitives (ProMPs) provide a theoretical framework to model variability and inherent correlation of human motions. In this work, we are therefore learning ProMPs to predict human intentions in terms of most likely future trajectories, based on prior observed motion data. Compared to prior work that relied on linear motion models, the learned ProMPs can hereby improve the accuracy and stability of real-time motion predictions. In particular, we compare two different approaches of incorporating ProMPs for the prediction of trajectories with partial observations: First, a method which conditions the weight distribution of learned ProMPs to match the early observed data points and, second, an Expectation-Maximization based algorithm which allows for learning from trajectory with missing data and accounts for the spatial-temporal variability of the demonstrations. Experimental evaluations on recorded human motion data of 25 subjects show a better performance of the first method. A trajectory prediction close to ground truth can already be made after observing 30 percent of trajectory points.

keywords: *human-robot cooperation, trajectory prediction, intention-aware ProMPs*

4 Experimental evaluation

In this section, we first describe the setup of experiments and the task performed by a human and the robot in the experiments of a user study conducted by Koert et al. [9]. After that, we present and discuss the results of different methods for phase estimation and trajectory prediction evaluated on the recorded data collected in the user study.

4.1 Experimental setup

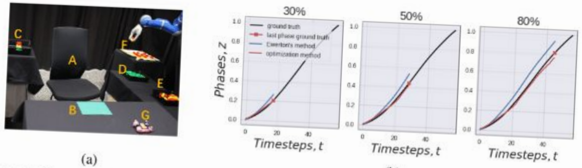

Figure 2: (a) Illustration of the experimental setup where a human and the robot perform a pick-and-place task in a shared workspace. Adapted from [9]. (b) Comparison of estimated phases of a trajectory from goal 2 to goal 1 after observing 30, 50 and 80 percent of trajectory points using Ewerton's method (blue) 3.3.1 and optimization method (red) 3.3.2.

In the experiment environment illustrated in Figure 2a, a human and the robot perform a pick-and-place task while sharing their workspace. The task of a human seated at A is to assemble blocks, which need to be collected from D and E in a freely chosen order, at the work area B and deliver them to C. During the task, the human is allowed to use only one hand to collect the blocks. On this hand, the human wears a glove G with visual markers that are used to capture the hand movement via an optical-tracking system. The task of the robot is deliver blocks from F to E. During the task, the robot should adapt its movement according to the predicted movement of the human. However, since the focus of the project lies in predicting the movement of the human, we use the recorded data from the experiments where the robot is not involved in the task so that it does not influence the movements of the human. The task is completed when there are no more building blocks in D. Note that the hand of a human is moving between areas B, C, D, and E, so these areas represent motion goals.

Write a Scientific Paper

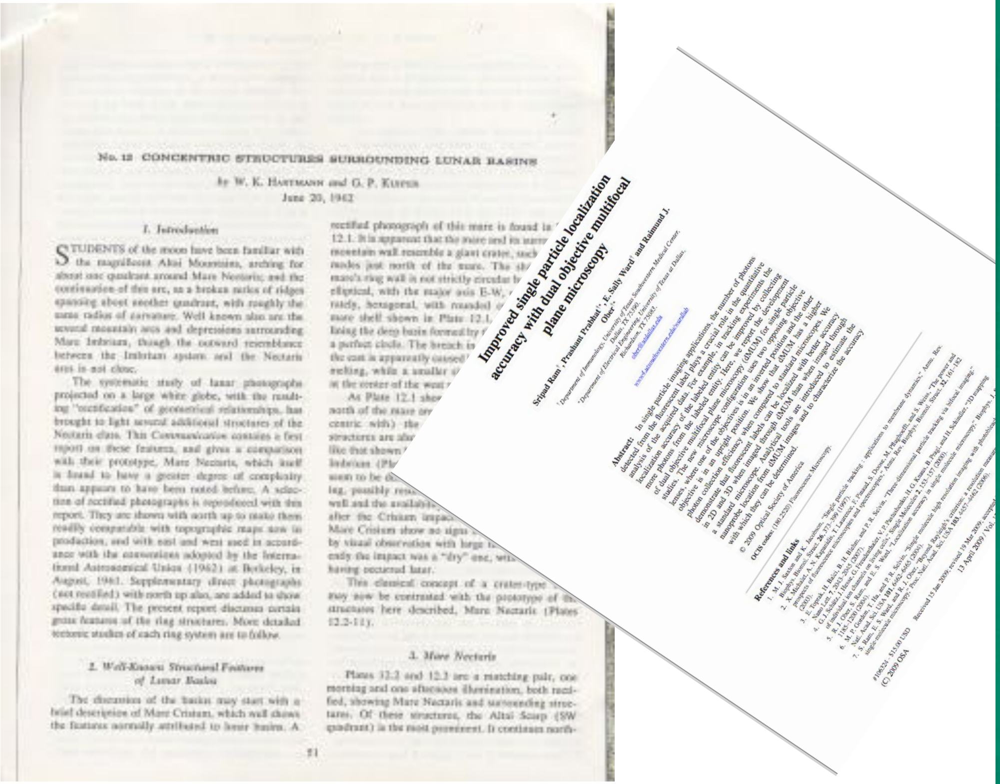

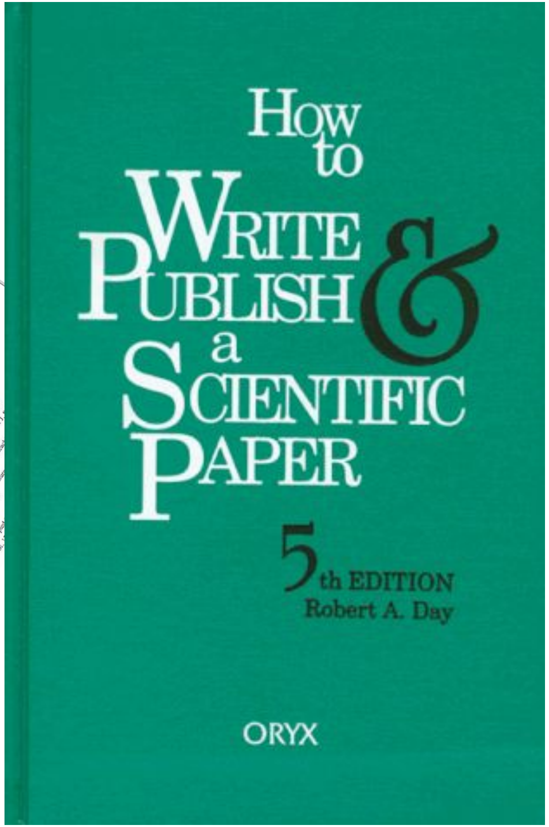

Do a Mini-Conference!

What We Offer....

We offer you a glimpse how life as a researcher in robot learning is like

- Use the knowledge from the robot learning lecture right away
- Decide what problem you are interested in and implement it in our simulator
- Write a scientific paper with a team
- Have a mini-conference at the semester's end
- Good projects can be continued as a Masters or Bachelors theses
    - [Coming soon] Outstanding projects are invited to attend the Seminar on Advanced Robot Learning (SARoL)
- **You are trying out how research life is like!**

Background (technical) Knowledge

- Is very project-dependent
- But might help if you have:
  - Mathematics from the first semesters (calculus, statistics)
  - Programming (project dependent, usually C/C++, Python)
  - Computer science fundamentals (algorithms)
- Simultaneous or previous attendance of the Statistical Machine Learning and Robot Learning lectures is very helpful!
- **Most important is that you have a wish to learn new topics!**

The Timeline

The Timeline

1. Choose a project, email the IP coordinator and supervisors (Until **25.04.2024**)

2. Topic Assignment (02.05.2024)

3. Work on the project…

4. First round submission of IP report (05.09.2024)

5. Peer review (12.09.2024)

6. Final submission (19.09.2024)

7. Presentation to the group (mini-conference) (26.09.2024)

Regulation regarding Report Submission

You could fail the course because of

- Fail to submit the reports and review in time. Including the first and second rounds submission of the report, and the review.

The Projects

Hands-on Control:
Tactile Feedback for Remote Robot Assembly

Supervisor: Tim Schneider, Kay Hansel

- Motivation:
  - Highlighting the crucial role of haptic feedback in assembly tasks.
  - Addressing the unresolved challenges in providing haptic feedback for Teleoperation.
- Goals :
  - Integrate haptic feedback using affordable Manus Gloves through vibration signals.
  - Successfully tackle complex assembly tasks via remote control
- Prerequisites :
  - Programming skills in Python, and optionally C++
  - Prior experience with Linux, and optionally ROS

Source: [3]

[1] Weber, Paul, et al. "A low-cost sensor glove with vibrotactile feedback and multiple finger joint and hand motion sensing for human-robot interaction." *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2016.
[2] Fritsche, Lars, et al. "First-person tele-operation of a humanoid robot." *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015.
[3] https://www.nist.gov/el/intelligent-systems-division-73500/robotic-grasping-and-manipulation-assembly/assembly

Control Barrier Functions for Assistive Teleoperation

Supervisor: Kay Hansel, Berk Guler

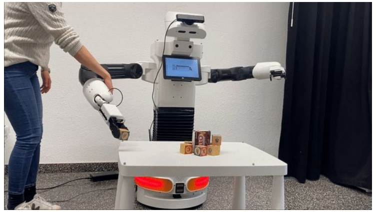

- Motivation:
  - Teleoperation faces challenges environmental changes, partial observability and time delays
  - How can we ensure safety, stability and robustness in the presence of such uncertainties and disturbances?
- Goals :
  - Implement control barrier functions (CBF)[1] for assistive teleoperation and shared control[2]
  - Compare against state-of-the-art (SOTA) methods, such as RelaxedIK[3]
  - Successfully handle complex teleoperation tasks.
- Prerequisites :
  - Programming skills in Python, optionally C++ or C#
  - Prior experience with Linux, optionally ROS
  - Optional: Prior Experience with Unity

[1] Ames, Aaron D., et al. "Control barrier functions: Theory and applications." *IEE ECC (2019);*
[2] Selvaggio, Mario, et al. "Autonomy in physical human-robot interaction: A brief survey." *IEEE RA-L (2021);*
[3] Rakita, D., et al. "RelaxedIK: Real-time Synthesis of Accurate and Feasible Robot Arm Motion." *RSS (2018);*

Learning Torque Control for Quadrupeds

Supervisor: Nico Bohlinger

- Motivation
  - Torque control allows for more agile behaviour
  - Learning high frequency torque control is hard because exploration is far more difficult compared to PD control
- Goals
  - Implement a NN architecture that combines PD and torque control
  - Show benefits in challenging tasks, e.g. high speed running, jumping, climbing, etc.
- Prerequisites
  - Great programming skills in Python
  - Experience with MuJoCo
  - Experience with JAX or PyTorch

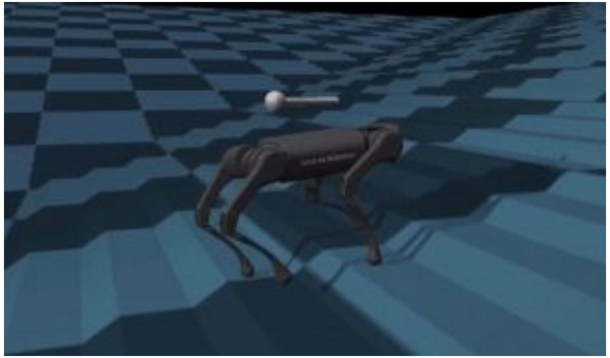

[1] Lee, Joonho, et al. "Learning quadrupedal locomotion over challenging terrain."
[2] Chen, Shuxiao, et al. "Learning Torque Control for Quadrupedal Locomotion."

Learning Locomotion with Human-Based Simulation

Supervisor: Michael Drolet, Firas Al Hafez

- Motivation:
    - Simulate locomotion / gait behaviors using SconePy [1], an open source software for performing predictive simulations of biological motion.
- Goals :
    - Gain familiarity with bio-inspired models and algorithms [2,3], using Reinforcement Learning and Imitation Learning
    - Create a controller for an exoskeleton
- Prerequisites :
    - Proficient in python and knowledgeable in robot learning methods

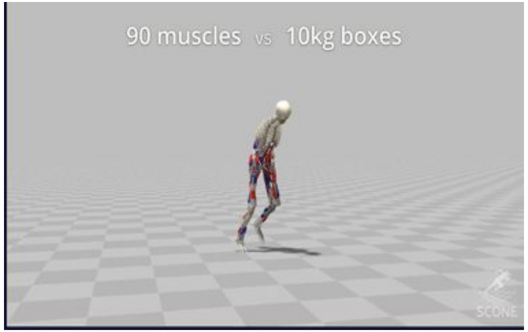

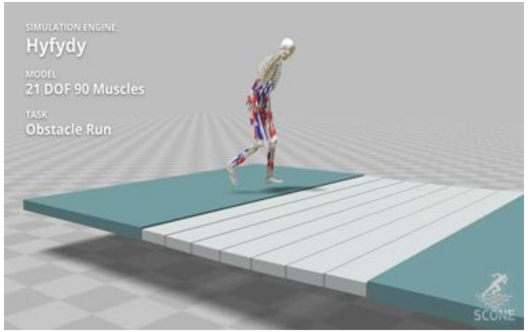

Source: https://github.com/martius-lab/depRL

[1] https://scone.software
[2] Schumacher, Pierre, et al. "Dep-rl: Embodied exploration for reinforcement learning in overactuated and musculoskeletal systems." arXiv preprint arXiv:2206.00484 (2022).
[3] Al-Hafez, Firas, et al. "LocoMuJoCo: A Comprehensive Imitation Learning Benchmark for Locomotion." arXiv preprint arXiv:2311.02496 (2023).

Universal Humanoid Motion Representations for Physics-Based Control

Supervisor: Boris Belousov

- Motivation:
  - Solve HumanoidBench (https://humanoid-bench.github.io/) using PULSE (https://github.com/ZhengyiLuo/PULSE)
- Goals :
  - Run PULSE (code provided) - if necessary, record additional human demonstrations using Optitrack at the IAS lab
  - Figure out how PULSE controllers can be used to attack some tasks from HumanoidBench
- Prerequisites :
  - Proficient in python
  - Attended lectures in robot learning, RL, SML, optimization, etc.

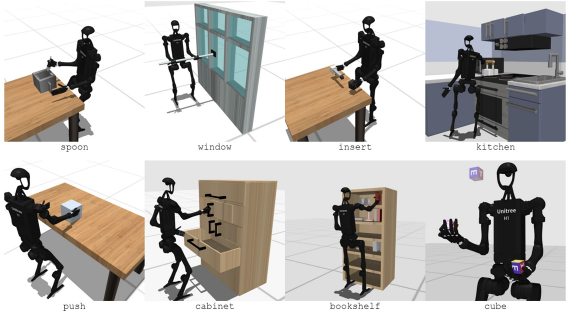

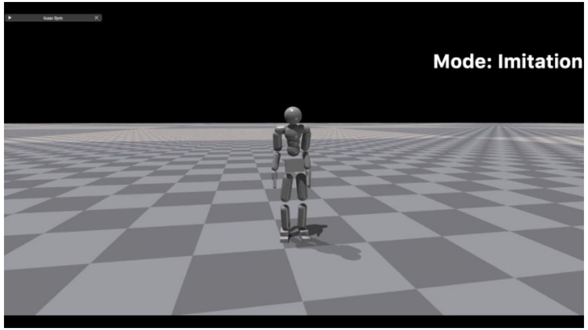

Real-Time Stream Transformation for Robot Control

Supervisor: Boris Belousov, Manisha Luthra-Agnihotri

- Motivation:
  - How to do real-time streaming transformation of events on robots? E.g., multimodal sensor fusion, vision-based feature extraction, feedback control, etc.
  - Currently ROS/ROS-2 are mainly used, but slow/bulky/unreliable/old in certain cases
  - Luckily, new **distributed dataflow systems** offer much lower latency, e.g., Dora (https://github.com/dora-rs/dora)
- Goals:
  - Design and implement real-time stream transformation of robot events using Dora-rs
  - Create & evaluate at least 2 dataflow applications, e.g., https://github.com/dora-rs/dora/tree/main/examples/python-operator-dataflow with vision input and robot action output
- Prerequisites:
  - Programming skills in Python, optionally C++ or Rust
  - Preferred: knowledge in distributed streaming systems

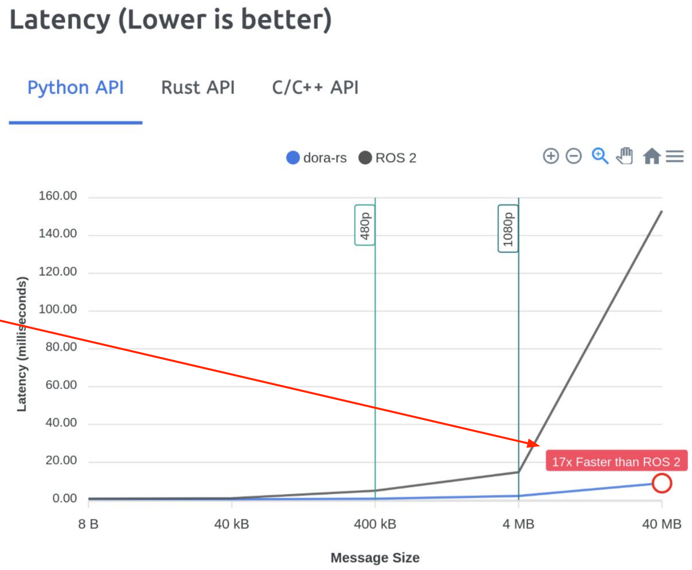

https://twitter.com/i/status/1777025381020688470

Robot Learning for Dynamic Motor Skills : A Case Study with Paper Planes

Supervisor: Kai Ploeger, Alap Kshirsagar

- Motivation
  - Teaching robots to perform dynamic motor skills remains a challenge
  - Throwing a paper plane involves both gross and fine motor skills, making it an interesting problem for robot learning

- Goals
  - Build a simulation environment for robotic paper plane throwing scenario
  - Train a barrett WAM robot arm to throw paper planes with high speed and accuracy
  - Evaluate imitation learning and reinforcement learning algorithms

- Prerequisites
  - Good python programming skills
  - Prior experience with simulation frameworks (MuJoCo/Gym)
  - Optional: Good paper plane throwing skills

AutoPlan: Autonomous Gearbox Assembly

Supervisor: Aiswarya Menon, Arjun Datta

- Motivation
  - Gearbox assembly is extremely complex and difficult to program
- Goals
  - Develop an advanced learning-based method that can autonomously generate the assembly sequence of a gearbox
  - Demonstrate solution for gear box assembly using the Panda robot
  - If successful at scale, present to Schaeffler and publish
- Pre-Requisites
  - Good programming skills in python and C++
  - Prior experience with robotics

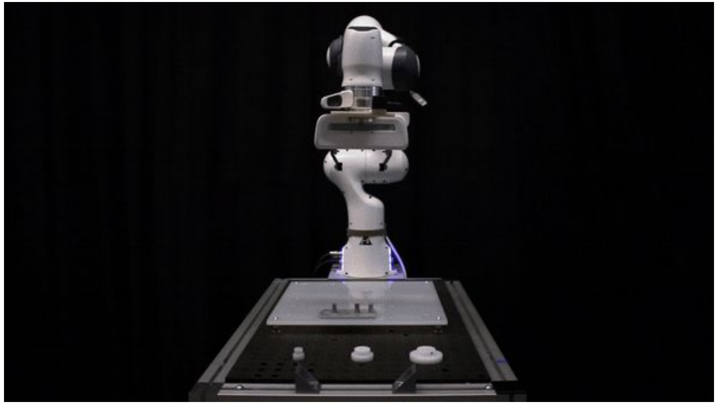

In cooperation with

TELEKINESIS

[1] Tian et al. ”Assemble Them All: Physics-Based Planning for Generalizable Assembly by Disassembly”
[2] Tian et al. “ASAP: Automated Sequence Planning for Complex Robotic Assembly with Physical Feasibility”
[3] telekinesis.ai

AutoScrew: Learning to Autonomously Screw

Supervisor: Aiswarya Menon, Arjun Vir Datta, Suman Pal

Description:

- Enhancing robot capabilities through advanced screwing skills acquisition

Goals:

- Develop a versatile learning framework to enhance robot adaptability across diverse screwing scenarios, validated through obstacle avoidance and tightening experiments.
- Demonstrate on multiple robots - Franka Emika, UR5
- If successful at scale, present to Schaeffler and publish

Pre-Requisites:

- Programming skills in C++, Python
- Basic knowledge of Robotics

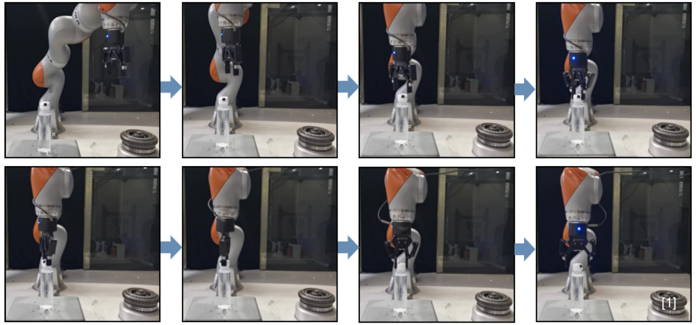

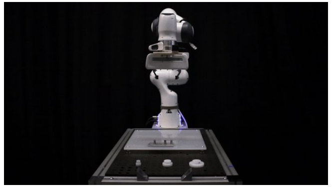

In cooperation with

TELEKINESIS

[1] Li F, Bai Y, Zhao M, Fu T, Men Y, Song R. Research on Robot Screwing Skill Method Based on Demonstration Learning. Sensors (Basel). 2023 Dec 19;24(1):21. doi: 10.3390/s24010021. PMID: 38202883; PMCID: PMC10780978.
[2]telekinesis.ai

AutoWeld: Learning to Autonomously Weld

Supervisor: Aiswarya Menon, Arjun Vir Datta, Suman Pal

Description:

- Automated offline programming framework, utilizing CAD models and advanced robotics technology to streamline programming and enhance welding efficiency.

Goals:

- Develop a precise framework for automatic generation of robot programs from CAD models, employing advanced robot motion planning algorithms, specifically tailored for welding.
- Demonstrate on multiple robots - Franka Emika, UR5
- If successful at scale, present to VW, Porsche etc. and publish

Pre-Requisites:

- Programming skills in C++, Python
- Basic knowledge of Robotics

In cooperation with

TELEKINESIS

[1] Li F, Bai Y, Zhao M, Fu T, Men Y, Song R. Research on Robot Screwing Skill Method Based on Demonstration Learning. Sensors (Basel). 2023 Dec 19;24(1):21. doi: 10.3390/s24010021. PMID: 38202883; PMCID: PMC10780978.
[2] Sarivan, IM., Madsen, O. & Wæhrens, B.V. Automatic welding-robot programming based on product-process-resource models. *Int J Adv Manuf Technol* (2024). https://doi.org/10.1007/s00170-024-13409-x
[3] https://www.methodtool.com/ar-1440-stem-weld-cell
[4] telekinesis.ai

MultiPlan: Multi-Robot Assembly

Supervisor: Aiswarya Menon, Arjun Vir Datta, Suman Pal

Motivation:

- Assembly automation via CAD-driven sequence generation and multi-robot motion coordination.

Goals:

- Develop an AI that, given the design files (CAD) of assembled model, auto-generates the assembly sequence plan and the multi-robot motion plan for successful assembly
- Demonstrate on multiple robots - Franka Emika, UR5
- If successful at scale, present to Schaeffler and publish

Pre-Requisites

- Programming skills in C++, Python
- Basic knowledge of Robotics

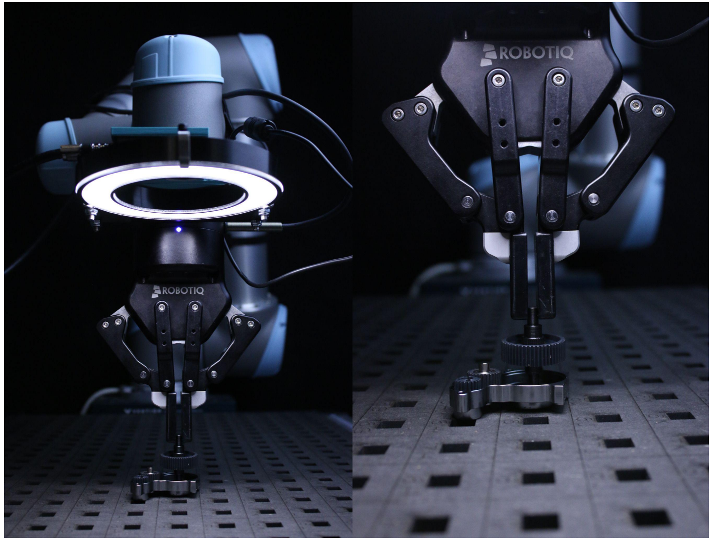

In cooperation with

TELEKINESIS

[1] Michniewicz et al. "CAD-based automated assembly planning for variable products in modular production systems"
[2] Tian et al. "ASAP: Automated Sequence Planning for Complex Robotic Assembly with Physical Feasibility"
[3]telekinesis.ai

DocPlan: Learning to Assemble from Instruction Manual

Supervisor: Aiswarya Menon, Vignesh Prasad, Arjun Datta

- Motivation
  - SMEs receive batch orders typically in the form of documents containing text, descriptions and industrial drawings
  - Humans interpret the documents to devise assembly plan
- Goals
  - Create vision language model that analyzes the instruction documents and extract essential features about the action type
  - Create action primitive library for gear box assembly
  - Demonstrate gearbox assembly on real Panda robot
  - If successful at scale, present to Schaeffler and publish
- Pre-Requisites
  - Good programming skills in python, C++ is a plus
  - Prior experience with computer vision and robotics

In cooperation with

TELEKINESIS

[1]Radford, Alec, et al. "Learning Transferable Visual Models From Natural Language Supervision."
[2]Michniewicz et al. “CAD-based automated assembly planning for variable products in modular production systems”
[3]telekinesis.ai

AutoPlan: Autonomous Body in White Assembly

Supervisor: Rukang Xu, Suman Pal

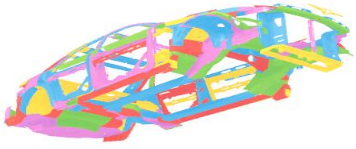

- Motivation
  - Body in white production is extremely complex to program
- Goals
  - Develop a novel AI that, given the industrial design files (CAD), auto-generates the robot code and the digital twin for Body in White production.
  - Demonstrate AI in professional simulation software such as Process Simulate, ABB RobotStudio, etc.
  - Present solution to companies like VW, Porsche, etc. if successful
- Pre-Requisites
  - Python
  - Prior experience with robotics and/or machine learning

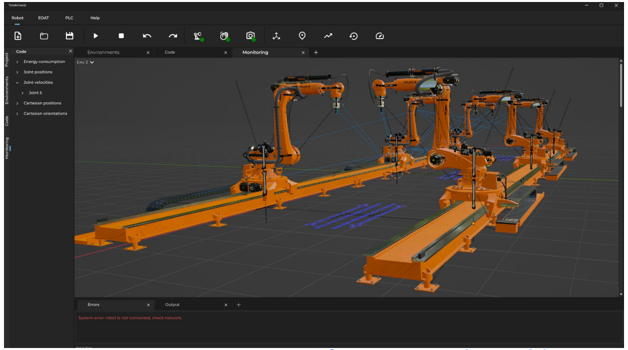

In cooperation with

TELEKINESIS

[1]Tian et al. ”Assemble Them All: Physics-Based Planning for Generalizable Assembly by Disassembly”
[2]Biesinger et al. “ A Case Study for a Digital Twin of Body-in-White Production Systems General Concept for Automated Updating of Planning Projects in the Digital Factory”
[3]telekinesis.ai

VisAutoPlan: Improved Autonomous Robotic Assembly with Visual Sensing

Supervisor: Rukang Xu, Aiswarya Menon

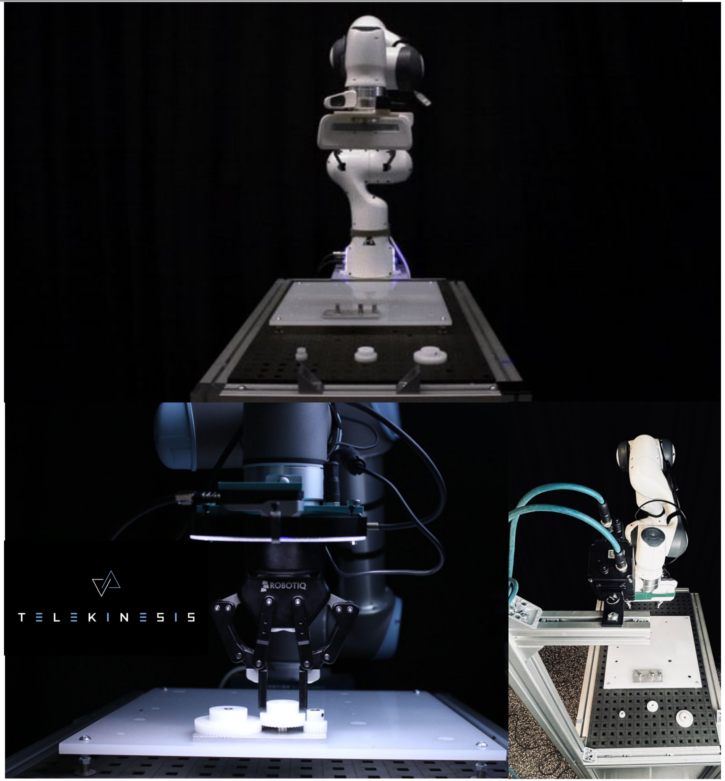

Motivation:

- Develop a visual-understanding-based AutoPlan that teaches robot autonomously for assembly task precisely and robustly.

Goals:

- Implement a scene perception feature on top of the existing approach[3] .
- Validate the full assembly pipeline with challenging real-world examples.
- If successful, present to Schaeffler and publish

Prerequisites :

- Programming skills in Python, optionally C++
- Prior experience with deep learning, computer vision and robotics

[1] Bowen Fu, et al. "6D Robotic Assembly Based on RGB-only Object Pose Estimation." *IROS (2022);*
[2] Andreas Wiedholz, et al. "Semantic 3D scene segmentation for robotic assembly process execution." *CASE (2023);*
[3] Yunsheng Tian, et al. "Assemble Them All: Physics-Based Planning for Generalizable Assembly by Disassembly." *SIGGRAPH Asia (2022);*

GearTrack: Tracking for Autonomous Gearbox Assembly

Supervisor: Andranik Aristakesyan, Arjun Datta, Vigment Prasad

Task Description:

- Apply 6D object pose estimation and tracking to gears for automated gearbox assembly

What you will gain:

- Deeper understanding of pose estimation and tracking for robotic tasks
- Hands-on experience with real robots and cameras
- If successful at scale, present to Schaeffler and publish

Requirements:

- Programming skills in Python
- Basic knowledge of Deep Learning, Computer Vision, and Transformers

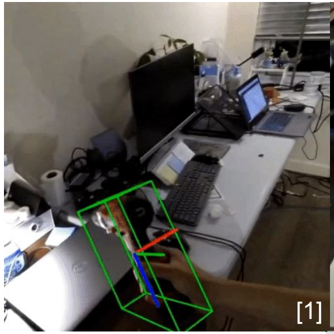

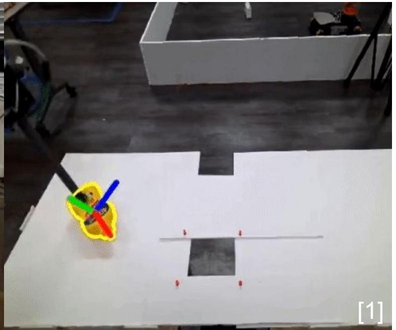

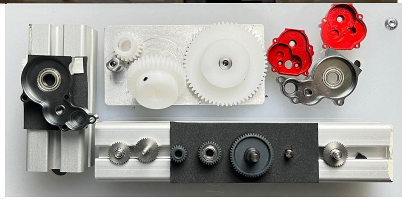

In cooperation with

TELEKINESIS

[1] Bowen Wen, Wei Yang, Jan Kautz, Stan Birchfield (2024, CVPR). FoundationPose: Unified 6D Pose Estimation and Tracking of Novel Objects
[2] Bowen Wen, Jonathan Tremblay, Valts Blukis, Stephen Tyree, Thomas Müller, Alex Evans, Dieter Fox, Jan Kautz, Stan Birchfield (2023, CVPR). BundleSDF: Neural 6-DoF Tracking and 3D Reconstruction of Unknown Objects

Visual Programming: Humanoids Learning from YouTube

Supervisor: Aiswarya Menon, Vignesh Prasad, Arjun Datta

- Motivation
  - Humanoids can learn from vast amount of YouTube videos
- Goals
  - Improve on the Visual Programming and imitation learning methods to ensure dynamically feasible humanoid motions.
  - Create a dataset of human videos and humanoid robot motions in MuJoco
  - Demonstrate multiple scalability to multiple humanoids such as Boston Dynamics, Unitree, Figure, etc.
  - If successful at scale, present to Boston Dynamics and publish
- Pre-Requisites
  - Good programming skills in python
  - Prior experience with computer vision and robotics

[1] LocoMujoco: https://github.com/robfiras/loco-mujoco
[3]telekinesis.ai

In cooperation with

TELEKINESIS

How to Apply

Homework (due end of day **Thursday, 25.04.2023**)

Write a short paragraph to answer the following questions:

1) Which project would **you** like to try and why?

2) Why do you think this project is **important**?

3) What **helpful background** do you have for the project and **what makes you special** for that project?

4) Your **academic aspirations**: *1 semester? 2 semesters? Future thesis?*

The participants can only send **two** such proposals to our PhD students.

Please specify your **priority** for the two projects.

If you already have a group, please send a joint email.

~Email **supervisors + kay.hansel@tu-darmstadt.de** (cc)

After a **meeting** with the potential supervisor(s), topics will be assigned to students.

Unfortunately, some students might not get topics!

Further information

Website of robot learning IP

https://www.ias.informatik.tu-darmstadt.de/Teaching/IP-RobotLearning#Part1 , you should be able to login with your TUCAN account and password.

Contact information of supervisors

https://www.ias.informatik.tu-darmstadt.de/Team/Members, you should be able to find the contact information of each supervisor here.

If you have any questions, feel free to email me:

kay.hansel@tu-darmstadt.de