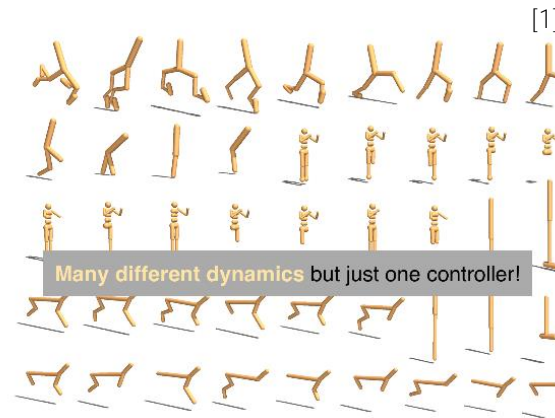


Bachelor/Student/Master Thesis (group work possible!)

Deep Reinforcement Learning with Foundation Models for Control



Description *Can one-and-the-same agent solve reference-tracking tasks in arbitrary systems?*

This question should be addressed within the thesis. In particular, the agent should, within a given interaction time, self-reliantly learn to steer the system to follow a given reference signal quickly and accurately. The control policy must be trained on a large amount of simulated dynamics to achieve generalization across a class of systems (domain randomization). This can be implemented using RL Gym Environments. The interaction time may be used for „fine-tuning“ a pre-trained foundation model.



Tasks

- Literature research for similar problem formulations
- Understanding of the existing RL Gym Environment
 - → <https://github.com/simon-bachhuber/L2L>
- Usage of RL algorithms from opensource RL libraries
- Scaling and hyperparameter optimization on LUIS cluster



Supervisors

Jan-Hendrik Ewering
 Daniel Weber

Requirements

- Self-dependent student with high intrinsic motivation
- Excellent Python Programming Skills
- Understanding of Machine Learning and Reinforcement Learning
- First experience with a deep-learning framework (such as TF, PyTorch, JAX)

Start
 As of now

[1] image source: <https://arxiv.org/abs/2007.04976>