Curriculum-Based Progressive Removal of Privileged Information in Reinforcement Learning

Technical University of Munich

School of Computation, Information, and Technology

Chair of Robotics, Artificial Intelligence and Real-time Systems

Background

Reinforcement learning agents trained in simulation often have access to privileged information (e.g., ground-truth dynamics, terrain parameters, object states) that is unavailable during real-world deployment. While existing methods either use fixed stochastic masking or ad-hoc gradual dropping of this information, no systematic study exists on how to optimally schedule the transition from privileged to deployment-realistic observations. This project investigates principled curriculum-based approaches for progressively removing privileged information during training, with the goal of improving sim-to-real transfer and policy robustness.

Description

Deep RL has achieved success in control tasks when trained in simulation with access to privileged state information. However, deploying these policies in the real world, where such information is unavailable, remains challenging due to the sim-to-real gap. Current approaches to this problem include: training separate teacher-student policies [1], using fixed dropout rates on privileged features [2], or learning compressed representations via information bottleneck methods [3]. While recent work has shown that gradually dropping privileged information can aid transfer, these methods lack theoretical grounding and systematic design principles.

This project aims to develop and analyze curriculum-based strategies for progressively removing privileged information during RL training. Rather than using fixed masking probabilities or ad-hoc schedules, we will investigate trigger mechanisms for transitioning between curriculum stages (performance-based vs. information-theoretic criteria).

Tasks

- Literature review on reinforcement learning with privileged information
- Design of curriculum strategies for progressive removal of privileged information
- Implementation of the proposed methods in a suitable RL framework
- Evaluation of the methodology on benchmark control tasks
- Documentation and thesis writing

References

[1] D. Chen, B. Zhou, V. Koltun, and P. Krähenbühl, “Learning by cheating,” in *Conference on Robot Learning*, 2020.

[2] P.-A. Kamienny, K. Arulkumaran, F. Behbahani, W. Boehmer, and S. Whiteson, “Privileged information dropout in reinforcement learning,” *arXiv:2005.09220*, 2020.

[3] H. He, P. Wu, C. Bai, H. Lai, L. Wang, L. Pan, X. Hu, and W. Zhang, “Bridging the sim-to-real gap from the information bottleneck perspective,” in *Conference on Robot Learning*, 2024.

Supervisor:
Prof. Dr.-Ing. Matthias Althoff

Advisor:
Olivia Garland, M.Sc.
Marlon Müller, M.Sc.

Research project:
—

Type:
Master’s thesis
(Bachelor’s thesis)

Research area:
Machine Learning

Required skills:
Reinforcement learning
Python, Git

Beneficial skills:
Robotic simulation
Curriculum learning

Language:
English

Date of submission:
27. Januar 2026

For more information please contact us:

Phone: —

E-Mail: olivia.garland@tum.de
marlon.mueller@tum.de

Website: www.ce.cit.tum.de/cps