Karlsruhe Institute of Technology

Autonome Lernende Roboter (ALR)

Prof. Gerhard Neumann

Project Type

- ☑ Master Thesis
- ☐ Bachelor Thesis
- ☑ Research Project

Supervisors

Niklas Freymuth

niklas.freymuth@kit.edu

Difficulty

Algorithmic

Math

Application

Representation-Invariant Latent Spaces for Robot Video Learning

Description

The prediction quality of Deep Neural Networks (DNNs) varies greatly depending on the representation of the data they are trained on. Additionally, it can be very difficult to translate between different representations, even if they contain similar information. As an example, the top part of Figure 1 shows both a human and a robot arm pushing an object to a certain target position. While both contain roughly the same semantic information ('the object is pushed towards the target position'), their representations are quite different.

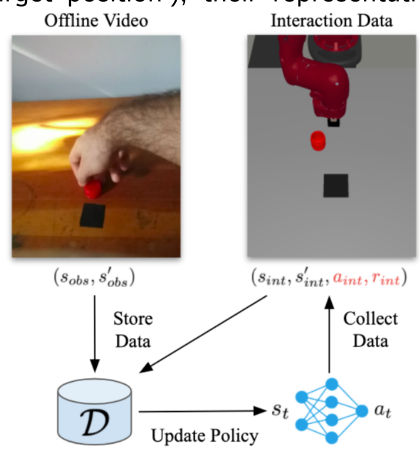

Figure 1: How can we integrate human demonstrations into a Reinforcement Learning framework [3]?

Motivated by the vast amount of videos available online, this project aims to use DNNs for a Reinforcement Learning agent that is able to efficiently learn from human video demonstrations. We will use recent advancements in Transfer Learning (see e.g., [4]) and Representation Learning [1] to develop a *representation-invariant latent space*, i.e., some hidden layer of the network that contains useful information about the data regardless of its input representation. For example, a first step could be to train a Variational Auto Encoder [2] jointly on different representations and shape its latent space by requiring that provided data pairings map to the same point in this space (c.f. [3]). This latent space will then be used to learn from demonstrations in both simulated and physical robot tasks.

Tasks

- Literature Review: Get familiar with existing transfer learning approaches and common extensions of VAEs.
- Architecture Design: Develop a novel DNN architecture or training regime for latent spaces suitable for teaching robots from videos.
- Evaluation: Evaluate the new approach on simulated and real robot tasks.

References

[1] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

[2] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *stat*, 1050:10, 2014.

[3] Karl Schmeckpeper, Oleh Rybkin, Kostas Daniilidis, Sergey Levine, and Chelsea Finn. Reinforcement learning with videos: Combining offline observations with interaction. *Conference on Robot Learning (CoRL)*, 2020.

[4] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, 2020.