ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Representation Learning for Reinforcement Learning (RL4RL)

Reinforcement Learning has achieved significant success in various domains. However, agents often struggle with understanding the complex pixel-based inputs and a noticeable performance gap persists between state-based algorithms and their pixel-based counterparts.

Auxiliary representation learning tasks, such as the reconstruction objective used in SAC-AE [5], have been shown to increase the sample efficiency of RL agents. Similarly, the use of disentangled representations has been beneficial, not only in reinforcement learning but also in various downstream applications [2, 4]. To further explore this area, Dunion et al. recently introduced temporal disentanglement (TED) in their work [1]. TED, a self-supervised auxiliary task, focuses on promoting disentangled representations in the encoder, utilizing the sequential nature of RL observations. Dunion et al. argue that this method enhances adaptability to environmental changes and generalization across task-irrelevant variables [1].

In the field of disentanglement, Klindt et al. developed Slow-VAE, a novel VAE loss, achieving state-of-the-art (SOTA) performance in disentanglement on temporally sparse-coded data [3]. My previous projects, conducted under my supervision, have demonstrated that incorporating this auxiliary loss can enhance the performance of RL agents, although its effectiveness is limited by the VAE's capabilities.

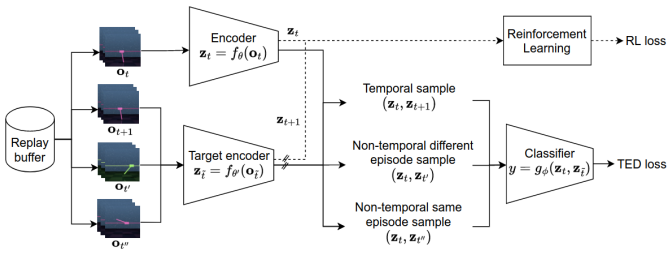

This thesis aims to develop novel auxiliary tasks for reinforcement learning, possibly inspired by previous work in disentanglement learning. The objective is to improve both the sample efficiency and generalization capabilities of RL agents.

Contact

- Benjamin Estermann: besterma@ethz.ch, ETZ G60.1

References

[1] Mhairi Dunion et al. "Temporal disentanglement of representations for improved generalisation in reinforcement learning". In: *arXiv preprint arXiv:2207.05480* (2022).

[2] Irina Higgins et al. "Darla: Improving zero-shot transfer in reinforcement learning". In: *International Conference on Machine Learning*. PMLR. 2017, pp. 1480–1490.

[3] David Klindt et al. “Towards nonlinear disentanglement in natural data with temporal sparse coding”. In: *arXiv preprint arXiv:2007.10930* (2020).

[4] Sjoerd Van Steenkiste et al. “Are disentangled representations helpful for abstract visual reasoning?” In: *Advances in neural information processing systems* 32 (2019).

[5] Denis Yarats et al. “Improving sample efficiency in model-free reinforcement learning from images”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35. 12. 2021, pp. 10674–10681.