ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Representation Learning for Reinforcement Learning (RL4RL)

Reinforcement learning (RL) agents often struggle with complex environments and are unable to generalise to out-of-distribution variations. Recent work by Dunion et al. [1] tries to address this issue using temporal disentanglement (TED). TED works as a self-supervised auxiliary task that encourages disentangled representations of the encoder based on the sequential nature or RL observations. Dunion et al. [1] claim that this leads to faster adaption to changes in environment variables and better generalization to unseen values of task-irrelevant variables.

In the field of disentanglement, Klindt et al. [2] developed a novel VAE loss, called SlowVAE, that achieves SOTA performance for disentanglement on data exhibiting temporal sparse coding.

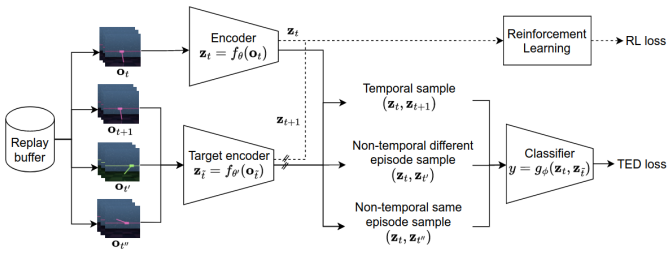

We are specifically interested in the benefit that SlowVAE or similar methods could bring over TED as auxiliary task in a reinforcement learning setting.

In this thesis, four separate agents are to be implemented and trained on a variety of pixel-based DeepMind Control Suite environments:

1. Baseline agent: AE loss, no auxiliary classifier
2. TED agent: AE loss, TED auxiliary classifier
3. SlowVAE agent: SlowVAE loss, no auxiliary classifier
4. SlowVAE+TED agent: SlowVAE loss, TED auxiliary classifier

All agents are to be trained with multiple off-policy algorithms under identical hyperparameters wherever possible. Subsequently, each agent's sample efficiency and performance shall be evaluated and compared, to assess whether the combination of SlowVAE and TED leads to an overall improvement in the reinforcement learning context.

The Student's Duties

- One meeting per week with the advisers to discuss current matters.
- A final report in English, presenting work and results.
- A final presentation (15 min) of the work and results obtained in the project.

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