Hierarchical Deep Reinforcement Learning for Robotics

Postdoc description

In recent years, Deep Reinforcement Learning (DRL) has lead to impressive results such as computers outperforming humans in video games or classic board games (Mnih et al., 2015, Silver et al., 2016). In contrast, the application of DRL to robotics tasks has turned out to be much more challenging (Lilicrap et al. 2015). Current algorithms still lack a human-like causal understanding and suffer from poor sample efficiency. One key to scaling up current approaches to human-level competence appears to lie in humans' abilitiy to discover generalizable hierarchical decompositions of control problems and the learning of corresponding hierarchically structured world models. This learning appears to be driven by intrinsic motivations to explore and gain control over the body and the environment. In this project, we will try to mimic these human abilities by blending ideas from hierarchical (Nachum et al., 2018) and model-based (Hafner et al., 2019 ; Moerland et al., 2020) DRL with intrinsic motivations. Through this, we aim to move current DRL approaches for robotics a decisive step closer to human-like skill acquisition.

Context

This postdoc is funded through the FrenchTech/I-site Cap2025 *chaire d'Excellence* program. The candidate will be joining the *Image, Perception Systems and Robotics* group of Institute Pascal which has long experience in computer vision and mobile robots. The research will be conducted in the context of an ongoing collaboration between Institut Pascal and Prof. Jochen Triesch from the Frankfurt Institute for Advanced Studies (FIAS) in Frankfurt, Germany.

Advisors:
- Prof. Jochen Triesch (Frankfurt Institute for Advanced Studies)
- Dr. Céline Teulière (Institut Pascal, UCA)

Duration: One year contract
Starting date: At your earliest convenience
Research Group: Institut Pascal
University: Université Clermont Auvergne (UCA) – Clermont Ferrand - France
Contact: celine.teuliere@uca.fr , triesch@fias.uni-frankfurt.de

Position Requirements

- PhD in Machine Learning, Robotics, Computational Neuroscience, or other relevant subjects
- Experience with Deep Reinforcement Learning ; knowledge of Computational Neuroscience and/or Computer Vision are a plus.
- Strong programming skills (C/C++, Python)
- Solid analytical ability
- Good spoken and written English
- High motivation

References

Hafner, D., Lillicrap, T., Ba, J., & Norouzi, M. (2019). Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*.

Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., ... & Wierstra, D. (2015). Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Petersen, S. (2015). Human-level control through deep reinforcement learning. Nature, 518(7540), 529.

Moerland, T. M., Broekens, J., & Jonker, C. M. (2020). Model-based reinforcement learning: A survey. *arXiv preprint arXiv:2006.16712*.

Nachum, O., Gu, S., Lee, H., & Levine, S. (2018). Data-efficient hierarchical reinforcement learning. *arXiv preprint arXiv:1805.08296*.

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., ... & Dieleman, S. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, *529*(7587)