

Internship offer: Benchmarking Hierarchical Reinforcement Learning algorithms in robotic simulation

Laboratory: U2IS, ENSTA Paris & LIX, Ecole Polytechnique

Duration: 6 months

Context: Fully autonomous robots have the potential to impact real-life applications, like assisting elderly people. Autonomous robots must deal with uncertain and continuously changing environments, where it is not possible to program the robot tasks. Instead, the robot must continuously learn new tasks and how to perform more complex tasks combining simpler ones (i.e., a task hierarchy). This problem is called lifelong learning of hierarchical tasks. Hierarchical Reinforcement Learning (HRL) is a recent approach for learning to solve long and complex tasks by decomposing them into simpler subtasks. HRL could be regarded as an extension of the standard Reinforcement Learning (RL) setting as it features high-level agents selecting subtasks to perform and low-level agents learning actions or policies to achieve them.

Summary: This internship studies the applications of Hierarchical Reinforcement Learning methods in robotics: Deploying autonomous robots in real world environments typically introduces multiple difficulties among which is the size of the observable space and the length of the required tasks.

Reinforcement Learning typically helps agents solve decision making problems by autonomously discovering successful behaviours and learning them. But these methods are known to struggle with long and complex tasks. Hierarchical Reinforcement Learning extend this paradigm to decompose these problems into easier subproblems with High-level agents determining which subtasks need to be accomplished, and Low-level agent learning to achieve them.

During this internship, the intern will :

- Get acquainted with the state of art in Hierarchical Reinforcement Learning including the most notable algorithms [1, 2, 3], the challenges they solve and their limitations.
- Reimplement some of these approaches and validate their results in robotics simulated environments such as iGibson [4].
- Establish an experimental comparison of these methods with respect to some research hypothesis.

The intern is expected to also collaborate with a PhD student whose work is closely related to this topic.

References:

[1] Nachum, O.; Gu, S.; Lee, H.; and Levine, S. 2018. Data- Efficient Hierarchical Reinforcement Learning. In Bengio, S.; Wallach, H. M.; Larochelle, H.; Grauman, K.; Cesa- Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montre´al, Canada*, 3307–3317.

- [2] Kulkarni, T. D.; Narasimhan, K.; Saeedi, A.; and Tenenbaum, J. 2016. Hierarchical Deep Reinforcement Learning: Integrating Temporal Abstraction and Intrinsic Motivation. In Lee, D.; Sugiyama, M.; Luxburg, U.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- [3] Vezhnevets, A. S.; Osindero, S.; Schaul, T.; Heess, N.; Jaderberg, M.; Silver, D.; and Kavukcuoglu, K. 2017. FeU- dal Networks for Hierarchical Reinforcement Learning. *CoRR*, abs/1703.01161.
- [4] Chengshu Li, Fei Xia, Roberto Martín-Martín, Michael Lingelbach, Sanjana Srivastava, Bokui Shen, Kent Vainio, Cem Gokmen, Gokul Dharan, Tanish Jain, Andrey Kurenkov, C. Karen Liu, Hyowon Gweon, Jiajun Wu, Li Fei-Fei, and Silvio Savarese. igibson 2.0: Object-centric simulation for robot learning of everyday household tasks, 2021. URL <https://arxiv.org/abs/2108.0327>