Universidade de Pernambuco
Programa de Pós-Graduação em Engenharia da Computação (PPGEC)

Proposta de Dissertação de Doutorado

Área: Computação Inteligente

Título: Deep Reinforcement Learning for Camera Control

Orientador – Bruno José Torres Fernandes (bjtf@ecomp.poli.br)
Co-orientador – Sergio Murilo Maciel Fernandes (smurilo@ecomp.poli.br)

Descrição – Reinforcement Learning (RL) [1] is a learning approach supported by behavioral psychology where an agent, e.g., a person or a robot, interacts with its environment trying to find an optimal policy to perform a particular task. In every time step, the agent performs an action reaching a new state and, sometimes, may obtain either a reward or a punishment. The agent tries to maximize the obtained reward by choosing the best action in a given state [2].

On the other hand, deep learning [3] is composed of many processing layers and has been successfully tested, among others, in image classification by representing different levels of abstraction [4]. Moreover, deep reinforcement learning [5] has combined the two aforementioned approaches to learning a motor policy mapping from a set of states to a set of actions. Deep reinforcement learning uses a neural network to learn the sum of direct rewards and expected future rewards for each action-state either in discrete or continuous domains [6].

In this project, the student will work with the deep reinforcement learning approach applied to a camera control scenario to track individuals in an environment. Moreover, the proposed technique could be extrapolated for other applications such as celestial bodies tracking through a telescope.

Referências Bibliográficas

1. R. S. Sutton and A. G. Barto. Reinforcement Learning: An Introduction. Cambridge, MA, USA: Bradford Book, 1998.
2. Francisco Cruz, Sven Magg, Yukie Nagai, and Stefan Wermter. Improving interactive reinforcement learning: What makes a good teacher? Connection Science, In Press, 2018.
3. I. Goodfellow, Y. Bengio, and A. Courville. Deep learning. Cambridge: MIT press, 2015.
4. Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. Nature, Vol. 521, Nr. 7553, pp. 436-444, 2015.
5. V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602, 2013.
6. M. Kerzel, H. Beik-Mohammadi, M. A. Zamani, S. Wermter. Accelerating Deep Continuous Reinforcement Learning through Task Simplification. In Proceedings of the International Joint Conference on Neural Networks (IJCNN), pp. 139-144, 2018.