DOINA PRECUP, McGill University / Mila / DeepMind
On Hierarchical Reinforcement Learning

Reinforcement Learning allows intelligent agents to learn by interacting with their environment over time and receiving rewards. Hierarchical Reinforcement Learning (HRL) approaches promise to provide more efficient solutions to sequential decision making problems, both in terms of statistical as well as computational efficiency. While this has been demonstrated empirically over time in a variety of tasks, theoretical results quantifying the benefits of such methods are still few and far between. In this talk, I will discuss the theoretical underpinnings of HRL in the framework of Markov and semi-Markov Decision Processes. I will describe the kind of structure in a Markov decision process which gives rise to efficient HRL methods. Specifically, we formalize the intuition that HRL can exploit repeating sub-structures. We show that, under reasonable assumptions, such algorithms can achieve statistical efficiency, as established through a finite-time regret bound, as well as near-optimal and computationally efficient planning, using hierarchical models.