PROF. LIAM PAULL, Université de Montréal
Training Robots in Simulators

One paradigm to building robotic agents is reinforcement learning. It is flexible and general. However, there are some particular challenges with respect to training RL agents on real physically embodied systems. For example: RL training tends to be quite innefficient and performing rollouts on a real robot system is expensive, real world environments don't automatically reset, and real world environments don't necessarily provide a reward signal to the agent explicitly. To overcome these challenges, training agents in simulators is appealing. However, the new problem becomes ensuring that an agent trained in a simulator generalizes to the real environment, the so-called sim2real problem. In this talk we will present two paradigms for tackling the sim2real, which we refer to as "Learn to Transfer" and "Learn to Generalize". We will also outline some future directions that we are pursuing in the Montreal Robotics and Embodied AI Lab (REAL) in this direction. Finally, I will also briefly describe our AI Driving Olympics project in connection to the problem of robotics benchmarking and "sim2real" transfer.