

# Distributed computing for RL

Asynchronous lock-free parallel Deep  
Reinforcement Learning

# RL Intro

- An agent learns to *interact* with an unknown stochastic and dynamic environment such as to *maximize some notion of reward*.
- In many practical applications, *agents evolve in a large state space* and must therefore *extract meaningful representations from rich sensory inputs*, which is typically achieved with neural networks.

# Learning representations in the RL setting

- *Explore* the state-action space
- Data is *non stationary*
- Data is *correlated*
  
- Experience replay = sample batch of observed transitions
- *asynchronously executing multiple agents* that share the same parameters in different instances of the environment.

# Asynchronous lock-free parallel training

- *Asynchronous lock-free parallel optimization* with *shared RMSProp*: a single shared vector is updated asynchronously and without locking to update the parameters of a neural network
- *Multiple threads* on multi-core CPU = *No network communication costs* :)
- *Linear speedup* (computational gains + improved exploration)