Seminário de sistemas dinâmicos e estocásticos

IMECC - UNICAMP

Título: Reinforcement Learning algorithms and partially observable stochastic control. .

Alberto Ohashi
UNB

Resumo:

The course first introduces reinforcement learning in the classical Markov decision process framework, focusing on SARSA and Q-learning as stochastic approximation algorithms for the Bellman optimality equation. We then move beyond the Markov setting and consider partially observable and non-Markovian stochastic control problems, where the controlled state may be observable but the dynamics depend on latent factors or unknown parameters. In particular, we discuss a recent approach where the effective state is an information set (or path history) and the associated value function becomes a functional, approximated using deep neural networks.