ITSC 2018 Tutorial on Deep Reinforcement Learning and Transportation

Aim and Scope

Deep reinforcement learning is an emerging area in machine learning, which has demonstrated promise in widely varied and challenging control problems, such as discrete tasks (Atari, board games), continuous tasks (hand manipulation, humanoid walking), and even seemingly combinatorial problems (Go). Control problems in transportation are highly varied and complex, consisting of cascaded dynamical systems of heterogeneous vehicles, humans, and infrastructure. Disparate techniques have been employed over the decades to address challenges in modeling and control, such as partial differential equations, queuing models, model predictive control, evolutionary algorithms, game theory, and many others. Each of these techniques provides theoretical or empirical results for different pieces and resolutions of the transportation system, which are crucial to the development of transportation control. However, these techniques are limited in the context of studying multiple pieces of or the holistic transportation problem. For instance, multi-lane freeways and intersections are typically modeled and studied separately using different mathematical techniques.

Deep reinforcement learning provides a data-driven scenario-agnostic framework for control, and recent empirical and theoretical results have demonstrate promise of these techniques for studying disparate transportation control problems in a unifying mathematical and algorithmic framework. This tutorial aims to provide an introduction to deep reinforcement learning to interested participants, providing a formal and algorithmic foundation, state-of-the-art practices, case studies in transportation, and a hands-on tutorial for getting started.

Outline of topics and respective expected duration (in minutes)

Expected duration (including ample time for Q&A)	Topic
60 min	Reinforcement learning and approximate dynamic programming • Reinforcement learning, stochastic control, Markov decision processes • Value iteration, policy iteration • Deep learning (briefly) • (Deep) Q-learning, critics
60 min	Policy optimization • Cross-entropy methods • Policy gradient methods • Actor-critic • Self-play
60 min	Case studies of deep reinforcement learning in transportation • Mixed-autonomy traffic (mixed automated and human-driven vehicles)

	• Infrastructure control
60 min	Tools of the trade • Flow (https://github.com/cathywu/flow) • SUMO (http://sumo.dlr.de/) • Ray/RLlib (https://github.com/cathywu/ray)
90 min	Hands-on tutorial on Flow, SUMO, and RLlib • Stabilize a ring road with mixed automated and human-driven vehicles • Design a custom traffic control task • Jointly control traffic lights and automated vehicles
60 min	Advanced topics in deep reinforcement learning • Hierarchical deep reinforcement learning • Deep multi-agent reinforcement learning • Safe reinforcement learning • Model-based reinforcement learning • Intrinsic rewards

Intended length

- Full day

Intended audience and assumed background knowledge

- Familiarity with basic control (feedback control, linear system theory) or machine learning
- Familiarity with Python

If any, personal computer and software requirements for attendees

- Flow (https://github.com/cathywu/flow)
- SUMO (http://sumo.dlr.de/)
- Ray (https://github.com/cathywu/ray)

Oganizers

- Alexandre Bayen, Liao-Cho Professor of Engineering; Director, Institute of Transportation Studies; Professor, EECS and CEE; UC Berkeley; bayen@berkeley.edu
- Cathy Wu, Ph.D. candidate, EECS, UC Berkeley; cathywu@eecs.berkeley.edu