Learning for Control Mini-Workshop

Organizers: Michael Zavlanos (Duke University) and George Pappas (University of Pennsylvania)

Purpose: Control of CPS is usually done in a model-based manner, where a desired control policy is computed from a high fidelity system model that has been derived at design-time, and potentially may be updated at runtime. However, this approach is not suitable for highly dynamical CPS, that potentially represent systems of systems whose spatial and temporal configurations may rapidly change. Consequently, it is critical to facilitate design of data-based controllers, with strong performance guarantees, in a way that allows for natural runtime control adaptation. The purpose of this workshop is to bring together experts in controls, machine learning, and optimization, to discuss recent developments, challenges, and new opportunities in data-driven control and optimization of CPS.

Attendance: 70+

Learning for Control Mini-Workshop

Speakers:

- Alexandre Bayen (UC Berkeley)
- Na Li (Harvard)
- Nikolai Matni (Penn)
- Aaron Ames (CalTech)

Topics Covered:

- Model-based learning for control, model-free learning for control, and distributed learning for control.
- Applications ranging from robotics to transportation.

Learning for Control Mini-Workshop

Future Directions on Learning for Control (Panel Discussion):

- End to end (pixel) learning
  - Model free deep-RL
  - Learning from videos (i.e. creating a video-input based model)
- Certifiable NN and explainable AI
  - Understanding I/O maps generated from learning
- Transfer learning
  - Learning in software transferred over multiple platforms
  - Demonstrations w. field operational tests (CAVs)
- Model-based RL
  - Learning on PDEs (hyperbolic conservation laws)
  - Apply to micro-models (and real world): atomic agents

Learning for Control Mini-Workshop

Future Directions on Learning for Control (Panel Discussion):

- Multi-agent/distributed learning
  - Decomposition methods, convergence rates
  - Cooperative vs. Non-cooperative methods
- Safe model-free learning
  - Definition of safety and robustness in a model-free setting
- Data-driven fundamental limits
  - If a data-driven controller is performing poorly, can we determine if this is due to poor exploration (learning), poor exploitation (control), or a bad system?
- System Co-design
  - Systematically build systems for which learning for control is easy
- Learning and control in autonomy
  - Combination of supervisor-level learning (tasks), certificate-level learning (invariants), and dynamics-level learning (plants)