LEHIGH UNIVERSITY | Institute for Data, Intelligent Systems, and Computation (I-DISC)

Robot Learning Workshop

October 14-15, 2019

Sumeet Singh

Google Brain Robotics, NY

Learning Stabilizable Nonlinear Dynamics with Contraction-Based Regularization

When it works, model-based Reinforcement Learning (RL) typically offers major improvements in sample efficiency in comparison to model-free techniques such as policy gradients that do not explicitly estimate the underlying dynamical system. Yet, all too often, when standard supervised learning is applied to model complex dynamics, the resulting controllers do not perform at par with model-free RL methods in the limit of increasing sample size, due to compounding errors across long time horizons. In this talk, I will present novel algorithmic tools leveraging Lyapunov-based analysis and semi-infinite convex programming to derive a control-theoretic regularizer for dynamics fitting, rooted in the notion of trajectory stabilizability. I will demonstrate how to embed these control-theoretic conditions as constraints within a semi-supervised algorithm for learning dynamical systems from user demonstrations. The constraints act as a form of context-driven hypothesis pruning to yield learned models that jointly balance regression performance and stabilizability, ultimately resulting in generated trajectories for the robot that are conditioned for feedback control. Experimental results on a quadrotor testbed will illustrate the efficacy of the proposed algorithms and clear connections between theory and hardware.