ECE 598: Interplay between Control and Machine Learning (Fall 2020)

Course Syllabus

Catalog Description: Advanced graduate course focuses on interplay between control and machine learning. The first half of the course focuses on tailoring control tools to study algorithms in large-scale machine learning. In the second half of the course, students will study how to combine reinforcement learning and model-based control methods for control design problems.

The following topics will be covered: empirical risk minimization; first-order methods for large-scale machine learning; stochastic optimization; dissipation inequality; jump system theory; Lur'e-Postnikov type Lyapunov functions; integral quadratic constraints; KYP Lemma; graphical interpretations for optimization methods; implicit bias; neural tangent kernel and adaptive control; control-oriented analysis tools for temporal difference learning and Q-learning; reinforcement learning for linear quadratic regulator (LQR) problems; learning model predictive control for iterative tasks; zeroth-order optimization and evolutionary strategies; policy gradient for robust control (global convergence and implicit bias); adversarial reinforcement learning; data-driven control of large-scale switching systems; iterative learning control; imitation learning for control; regularization of model-free control via prior model-based design; constrained policy optimization.

Course Objectives: Upon successful completion of the course, students will be able to explain two basic ideas: (i) applying control methods for learning; (ii) applying learning methods for control. Students will be able to understand the mathematics underlying such interdisciplinary research. Students will be prepared to tailor their own expertise (in either control or machine learning) for research at the intersection of control and machine learning.

Outline of Topics:

- A Control Perspective on Optimization Methods in Large-Scale Machine Learning (12 hours):
    - Brief introduction to empirical risk minimization and related optimization Methods
    - Dynamical system perspectives for optimization methods
    - A unified analysis of large-scale optimization methods via dissipativity
    - Jump system theory for stochastic optimization
    - Lure-Postnikov type Lyapunov functions
    - Weighted off-by-one integral quadratic constraint
    - Design of accelerated stochastic optimization methods
    - Multiplier theory and KYP lemma
    - A control perspective on implicit bias
    - Connections between neural tangent kernel and adaptive control

- Control-Oriented Analysis Tools for Reinforcement Learning Algorithms (6 hours):
    - Analysis of temporal difference learning with linear function approximators
    - Analysis of Q-learning with linear function approximators
    - Analysis of emphatic TD learning
    - Variance reduction techniques in TD learning

- Reinforcement learning for linear quadratic control (12 hours):
    - Brief review of linear quadratic regulator (LQR)
    - Policy gradient for LQR: REINFORCE and Evolutionary Strategies (ES)
    - Convergence guarantees of policy gradient on LQR problem
    - Q-learning and approximate policy iteration for LQR
    - Iterative feedback tuning for gradient estimation
    - Policy gradient for linear quadratic control of large-scale switching systems
    - Policy gradient for linear quadratic Gaussian (LQG) with partial observation
    - Iterative learning control
    - Sample-based model predictive control

- Robustness and safety in reinforcement learning/control (12 hours):
    - Model-based robust control: H-infinity control, DK iteration, and IQC-synthesis

- Risk-sensitive control: linear exponential quadratic Gaussian (LEQG)
- Policy gradient for mixed H2/H-infinity control and LEQG (global convergence and implicit bias)
- Adversarial robust reinforcement learning and linear quadratic games
- Zeroth-order optimization for robust control
- Model-based robust reinforcement learning and system level synthesis
- Robustness in imitation learning
- Regularization of model-free deep reinforcement learning via prior model-based control design
- Safe reinforcement learning via constrained policy optimization (CPO)

Text: none required; instructor's lecture notes will be used. Some papers will be used as references:

- L. Lessard, B. Recht, and A. Packard, Analysis and design of optimization algorithms via integral quadratic constraints, SIAMopt 2016.

- B. Hu, P. Seiler, and A. Rantzer, A unified analysis of stochastic optimization methods using jump system theory and quadratic constraints, COLT 2017.

- B. Hu and L. Lessard, Dissipativity theory for Nesterov's accelerated method, ICML 2017.

- B. Hu, S. Wright, and L. Lessard, Dissipativity theory for accelerating stochastic variance reduction: A unified analysis of SVRG and Katyusha using semidefinite programs, ICML 2018.

- A. Rantzer, On the Kalman-Yakubovich-Popov Lemma, SCL 1996.

- S. Du, X. Zhai, B. Poczos, and A. Singh, Gradient descent provably optimizes over-parameterized neural networks, ICLR 2019.

- M.A. Belabbas, On implicit regularization: Morse functions and applications to matrix factorization, 2020.

- R. Srikant and L. Ying, Finite-time error bounds for linear stochastic approximation and TD learning, COLT 2019.

- B. Hu and U. Syed, Characterizing the exact behaviors of temporal difference learning algorithms using Markov jump linear system theory, NeurIPS 2019.

- D. Soudry, E. Hoffer, M. Nacson, and N. Srebro, The implicit bias of gradient descent on separable data, ICLR 2018

- K. Krauth, S. Tu, and B. Recht, Finite-time analysis of approximate policy iteration for the linear quadratic regulator, NeurIPS 2019.

- M. Fazel, R. Ge, S. Kakade, and M. Mesbani, Global convergence of policy gradient methods for the linear quadratic regulator, ICML 2018.

- J.P. Jansch-Porto, B. Hu, and G. Dullerud, Convergence guarantees of policy optimization methods for Markovian jump linear systems, ACC 2020

- U. Rosolia and F. Borrelli, Learning model predictive control for iterative tasks: A data-driven control framework, TAC 2017.

- H.K. Venkataraman, and P. Seiler, Recovering robustness in model-free reinforcement learning, ACC 2019.

- S. Chen, K. Saulnier, N. Atanasov, D.D. Lee, V. Kumar, G.J. Pappas, M. Morari, Approximating explicit model predictive control using constrained neural networks, ACC 2018.

- J. Veenman and C. Scherer, IQC-synthesis with general dynamic multipliers, IJRNC 2013.

- S. Wang, H. Pfifer, and P. Seiler, Robust synthesis for linear parameter varying systems using integral quadratic constraints, Automatica2016.
- L. Pinto, J. Davidson, R. Sukthankar, and A. Gupta, Robust adversarial reinforcement learning, ICML 2017.
- K. Zhang, B. Hu, and Tamer Başar, Policy Optimization for H2 Linear Control with H-Infinity Robustness Guarantee: Implicit Regularization and Global Convergence, 2019.
- R. Cheng, A. Verma, G. Orosz, S. Chaudhuri, Y. Yue, and J. Burdick, Control Regularization for Reduced Variance Reinforcement Learning, ICML 2019.
- J. Achiam, D. Held, A. Tamar, and P. Abbeel, Constrained policy optimization, ICML 2017.

Coursework: 50% regular homework sets; 50% written research report

Prerequisites: ECE 515. ECE 534 and ECE 490 are recommended, but not required.