Tutorial Sessions

The 58th Annual Allerton Conference on Communication, Control, and Computing Tutorials were given Tuesday, September 27, 2022, by **Professor Mihailo Jovanovic** from the Ming Hsieh Department of Electrical and Computer Engineering at the University of Southern California and **Professor Sean Meyn** from the Department of Electrical & Computer Engineering at the University of Florida.

Mihailo Jovanovic presented his tutorial from 9:30 a.m. to 1:30 p.m., which included a one-hour lunch break.

Title: ***Robustness of Gradient Methods for Data-driven Decision Making***

Abstract:
Gradient descent and its accelerated variants are increasingly used for learning and data-driven control design for uncertain dynamical systems in which an approximation of the gradient is sought through noisy measurements. In the first part of the talk, we utilize techniques from control theory to quantify robustness of accelerated first-order algorithms to stochastic uncertainties in gradient evaluation. In particular, for unconstrained, smooth, strongly convex problems, we demonstrate how tight upper bounds on the mean-square error in the optimization variable can be established when the iterates are perturbed by additive white noise. Our analysis reveals fundamental trade-offs between noise amplification and convergence rates for any acceleration scheme similar to Nesterov's or

heavy-ball methods. To gain additional analytical insight, for strongly convex quadratic problems we explicitly evaluate noise amplification in terms of the spectrum of the Hessian matrix. We specialize this result to the problem of distributed averaging over undirected networks and examine roles of network size and topology on robustness of noisy accelerated algorithms.

In the second part of the talk, we focus on model-free reinforcement learning which attempts to find an optimal control action for an unknown dynamical system by directly searching over the parameter space of controllers. The convergence behavior and statistical properties of these approaches are often poorly understood because of the non-convex nature of optimization problems and the lack of exact gradient computation. We examine performance and efficiency of such methods by focusing on the standard linear quadratic regulator problem for systems with unknown state-space parameters. We establish exponential stability of the gradient descent method over the set of stabilizing feedback gains and provide theoretical bounds on the convergence rate and sample complexity of the random search method with two-point gradient estimates. We also prove that the required simulation time for achieving ε-accuracy in the model-free setup and the total number of function evaluations both scale as $\log(1/\varepsilon)$.

Biography:
Mihailo Jovanovic is a professor in the Ming Hsieh Department of Electrical and Computer Engineering and the founding director of the Center for Systems and Control at the University of Southern California. He was a faculty member in the Department of Electrical and Computer Engineering at the University of Minnesota, Minneapolis, from 2004 until 2017, and has held visiting positions with Stanford University, the Institute for Mathematics and its Applications, and the Simons Institute for the Theory of Computing. His current research focuses on large-scale and distributed optimization; design of controller architectures; fundamental limitations in the control of networks of dynamical systems; and modeling, dynamics, and control of fluid flows. Prof. Jovanovic received a CAREER Award from the National Science Foundation in 2007, the George S. Axelby Outstanding Paper Award from the IEEE Control Systems Society in 2013, and the Distinguished Alumnus Award from the University of California at Santa Barbara in 2014. He is a Fellow of the American Physical Society (APS) and the Institute of Electrical and Electronics Engineers (IEEE).

Sean Meyn presented his tutorial from 2:00 p.m. to 5:30 p.m., which included a 30-minute break.

Title: ***Control Systems and Reinforcement Learning***

Abstract:
By the end of the three hours, you will create decision and control algorithms that are truly automatic: without any knowledge of physics or biology or medicine, your RL algorithm tunes itself to become a super controller: the smoothest ride into space, and the most expert micro-surgeon! (Humor like this will be sprinkled throughout the tutorial.)

The tutorial is designed for newcomers to control systems, and consists of two parts:

1. The ODE Method. The basics of algorithm design grounded in recent refinements and interpretations of stochastic approximation.
2. What is Q? A survey of approaches to Q-learning for approximate optimal control. The magic of the methodology is emphasized, along with warnings that the gap between theory and practice is enormous. Why Q-learning is so often so successful remains a mystery.

Prerequisites: My dream audience would consist of students who mastered a first graduate-level course on linear systems (such as ECE 515 at UIUC) and also have an appreciation for the basic ideas from classical feedback control (such as ECE 486 at UIUC). However, the course can be followed

with only background in ODEs and linear algebra. No probability theory is required. The tutorial will be based on Part I of my recent book *Control Systems and Reinforcement Learning* (https://meyn.ece.ufl.edu/control-systems-and-reinforcement-learning/).

Biography:
Sean Meyn was raised by the beach in Southern California. Following his BA in mathematics at UCLA, he moved on to pursue a PhD with Peter Caines at McGill University. After about 20 years as a professor of ECE at the University of Illinois, in 2012 he moved to beautiful Gainesville. He is now Professor and Robert C. Pittman Eminent Scholar Chair in the Department of Electrical and Computer Engineering at the University of Florida, and director of the Laboratory for Cognition and Control. He also holds an Inria International Chair to support research with colleagues in France. His interests span many aspects of stochastic control, stochastic processes, information theory, and optimization. For the past decade, his applied research has focused on engineering, markets, and policy in energy systems.