FOPAM 2026
Proceedings of the 3rd Foundations of Process/Product Analytics and Machine Learning
Decatur, USA, July 26-30 (2026)

Local-Global Learning of Interpretable Control Polices: The Interface between MPC and Reinforcement Learning

Ali Mesbah
University of California, Berkeley, Department of Chemical and Biomolecular Engineering

ABSTRACT

Optimal decision-making under uncertainty is a shared challenge across modern chemical, manufacturing, and energy systems that increasingly demand safe, data-driven autonomy. This talk revisits optimal control through the lens of the Bellman equation, emphasizing how optimal control theory and reinforcement learning have developed complementary, yet largely disconnected, perspectives on global optimality. In one view, central to reinforcement learning, the Bellman equation defines a global optimality condition that guides iterative policy learning from interacting with the system, but typically yields opaque control laws that are difficult to interpret, and deploy in safety-critical settings. In another view, widely adopted in model predictive control (MPC), the Bellman equation underpins tractable finite-horizon optimizations that deliver interpretable, constraint-aware, and modular local controllers, yet without explicit guarantees on alignment with global optimality. Building on these ideas, we introduce a local–global paradigm that treats MPC and related optimization-based controllers as structured function approximators designed to approximately satisfy the global Bellman optimality condition. We discuss algorithmic strategies for learning interpretable local decision makers whose adaptation is guided by Bellman residuals, along with the benefits and practical challenges that arise in terms of stability, constraint satisfaction, and sample efficiency. These concepts are illustrated through case studies that unify reinforcement learning and MPC for safe, high-performance control in complex, uncertain dynamical systems. The talk concludes by outlining open problems and research opportunities in learning interpretable control policies that achieve globally optimal performance while retaining the transparency and reliability required for real-world process control and optimization applications.