Reza Skandari, Imperial College Business School (UK)

Title. Balancing Complexity and Personalization in Sequential Decision Models for Chronic Disease Management

Abstract. Markov Decision Processes (MDPs) and their partially observable counterparts (POMDPs) are widely used in healthcare for decision-making across the continuum of care, from screening and diagnosis to treatment and long-term monitoring. These models provide a structured framework for integrating clinical data, patient heterogeneity, and uncertainty into medical decisions. However, healthcare systems face a critical trade-off: while fully personalized policies optimize patient-specific outcomes, their exponential complexity makes them impractical. Conversely, one-size-fits-all approaches, though easier to implement in real-world settings, often lead to inefficiencies and disparities in patient care. This talk explores how machine learning and optimization can help bridge this gap by determining the optimal size and configuration of population stratification. By clustering patients based on key characteristics, we can design equitable and scalable policies that improve health outcomes while remaining computationally and practically feasible. We will discuss recent advancements at the intersection of machine learning, optimization, and healthcare decision science, highlighting key methodological challenges and open questions. To illustrate these concepts, we present a case study on cancer surveillance, demonstrating how stratifying patients based on a small number of meaningful characteristics can significantly enhance equity and patient outcomes compared to a one-size-fits-all policy—without requiring fully personalized approaches. This work underscores the potential of integrating sequential decision-making frameworks with data-driven population stratification to improve healthcare delivery.