Contribution ID: **125** Type: **Lecture**

Multi-objective RL

Tuesday 8 April 2025 11:15 (1 hour)

Multi-objective reinforcement learning (MORL) extends traditional reinforcement learning (RL) by addressing environments where multiple conflicting objectives must be optimized simultaneously. In real-world applications, such as autonomous systems, particle accelerator optimization and control, agents often face trade-offs between competing goals. This lecture provides an overview of the key concepts, techniques, and challenges in MORL. We will explore different approaches to handling multiple objectives, including scalarization methods, Pareto dominance, and reward decomposition. Additionally, we will discuss the concept of the Pareto front and how it can guide decision-making in complex environments. The lecture will also highlight the role of exploration-exploitation trade-offs in multi-objective settings and the development of algorithms that can balance these objectives effectively. By the end of the session, attendees will gain a deeper understanding of how to design and implement multi-objective RL systems capable of navigating the complexities of real-world decision-making tasks.

Presenter: RAJPUT, Kishan (Jefferson Lab)