

APS1080 Reinforcement Learning Course Description

Artificial Intelligence (AI) is a disruptive technology that represents the pinnacle of computer science and engineering. Reinforcement Learning (RL) is the dominant toolset of modern AI.

In this course, we will provide an introduction to Reinforcement Learning, with the goal of raising students from a level where they have no knowledge of RL to one where they can productively employ advanced RL methods in their research and professional work. RL addresses the design of AI machines (“agents”) capable of learning from experience about how to behave in a goal-directed manner within unknown environments and problem spaces. Although RL is positioned as one of three facets of Machine Learning, RL in fact has much broader scope than the narrower tools of supervised and unsupervised learning; RL is the only ML toolset concerned directly with AI.

We will cover the foundations, including dynamic programming, Monte Carlo, Temporal-Difference Learning methods. From here, we will cover more advanced schemes that exploit function approximation methods (neural networks) with RL: DQNs, Policy Gradient methods, Planning (including MuZero), Reinforcement Learning with Human Feedback (e.g., for LLMs), and how RL-only LLMs (e.g., Deepseek-R1-zero) are trained.

Through course work students will develop a theoretical and practical understanding of the subject, with the goal of being able to practically apply the toolset to novel problems.