



CR15

Reinforcement Learning



Planning and Dynamic Programming:
Markov Decision Processes

Learning and Stochastic Approximation:
bandits, Q-learning, deep RL

Prerequisites: probability, linear algebra, basic programming

