

Date	Week #	Topic	Homeworks
Aug. 29	1	Course Overview; meet and greet; ML intro; RL intro	
Sep. 2	2	intro to linear algebra and probability	
Sep. 5		no class (Labor Day)	
Sep. 9	3	Supervised learning overview; linear regression	HW1
Sep. 12		Logistic regression	
Sep. 16	4	Decision Trees	HW2
Sep. 19		Intro to RL and control	
Sep. 23	5	Intro to Multi-armed Bandits	
Sep. 26		Bayesian Bandits	
Sep. 31	6	Markov chains	HW3
Oct. 3		Markov reward processes	
Oct. 7	7	Markov decision processes	HW4
Oct. 10		Bellman Optimality Equations	
Oct. 14	8	Dynamic Programming, Policy/Value iteration	HW5
Oct. 17		Monte Carlo methods	
Oct. 21	9	Temporal difference learning	
Oct. 24		Q-learning	
Oct. 28	10	Fully connected neural networks	HW6
Oct. 31		Optimization	
Nov. 4	11	Convolutional neural networks, Image Classification	HW7
Nov. 7		Regularization	
Nov. 11	12	Generalization	HW8
Nov. 14		Q-learning with function approximation	
Nov. 18	13	Policy gradient theorem, REINFORCE algorithm	
Nov. 21		Policy gradients with function approximation, actor-critic	
Nov. 25	14	no class (Thanksgiving)	
Nov. 28		no class (Thanksgiving)	
Dec. 2	15	Deterministic policy gradients	HW9
Dec. 5		Soft Actor-Critic, Rainbow	
Dec. 9	16	Offline RL	
Dec. 12		Imitation Learning	HW10