Date	Week #	Topic	Homeworks
Aug. 30	1	Course Overview; meet and greet; ML intro	
Sep. 3	2	no class (Labor Day)	
Sep. 6		intro to linear algebra and probability	
Sep. 10	3	Supervised learning overview; linear regression	HW1
Sep. 13		Logistic regression	
Sep. 17	4	Decision Trees	HW2
Sep. 20		RL intro, modeling	
Sep. 24	5	Intro to Multi-armed Bandits	
Sep. 27		Bayesian Bandits	
Oct. 1	6	Finite state automata, Markov chains	HW3
Oct. 4		Markov reward processes	
Oct. 8	7	Markov decision processes	HW4
Oct. 11		No class	
Oct. 15	8	Dynamic Programming, Policy/Value iteration	HW5
Oct. 18		Monte Carlo methods	
Oct. 22	9	Temporal difference learning	
Oct. 25		Q-learning	
Oct. 29	10	Fully connected neural networks	HW6
Nov. 1		Optimization	
Nov. 5	11	Convolutional neural networks, Image Classification	HW7
Nov. 8		Regularization	
Nov. 12	12	Generalization	HW8
Nov. 15		Q-learning with function approximation	
Nov. 19	13	Policy gradient theorem, REINFORCE algorithm	
Nov. 22		Policy gradients with function approximation, actor-critic	
Nov. 26	14	Deterministic policy gradients	HW9
Nov. 29		no class (Thanksgiving)	
Dec. 3	15	Soft Actor-Critic, Rainbow	
Dec. 6		Offline RL	
Dec. 10	16	Imitation Learning	HW10