

Course Code		Course Title	Online learning Algorithms			
Dept. /Faculty proposing the course	Dr R Sri Prakash	Structure (LTPC)	L	T	P	C
			3	1	0	4
To be offered for	UG/PG	Type	Core ☐		Elective ☒	
		Status	New ☒		Modification ☐	
Pre-requisite	Probability theory	Submitted for approval				
Learning Objectives	● Understand challenges in adversarial and stochastic online learning with sequential data. Grasp the multi-armed bandit problem as a sequential decision-making framework under uncertainty. ● Comprehend Markov Decision Process (MDP) components and foundational concepts like value functions and Bellman optimality equations for solving MDPs.					
Learning Outcomes	● Students will analyse online learning algorithms in adversarial and stochastic settings and will formulate, solve multi-armed bandit problems, balancing exploration and exploitation. ● Students will also model problems as Markov Decision Processes (MDPs), compute value functions, and apply Bellman equations to determine optimal policies.					
Contents of the course (With approximate break-up of hours for L/T/P)	Basics of probability: review of conditional probability, random variables, expectations, conditional expectations, concentration inequalities, and random processes (L6+T2) Online learning in adversarial settings, Learning from experts, Bandit setting (L12+T4) Stochastic bandits: Regret minimisation, Information theoretic bounds, Bayesian setting (L12+T4) Pure exploration: Best arm Identification with fixed confidence, Best arm identification under fixed budget (L6+T2) Introduction to Markov Decision Process, Value functions, Bellman optimality equation, Reinforcement Learning algorithms, Markovian bandits (L6+T2)					
Text Books	1. Tor Lattimore, and Csaba Szepesvári, Bandit Algorithms, Cambridge University Press, ISBN: 1108486827, 2020. 2. Aleksandrs Slivkins, Introduction to Multi-armed Bandits, Now publishers, ISBN: 9781680836202, 2019.					
Reference Books	1. S. Shalev-Schwartz, Online Learning and Online Convex Optimization, NOW Publishers, ISBN: 1601985460, 2011. 2. Sutton, Richard S., and Andrew G. Barto, Reinforcement learning: An introduction, 2nd edition, MIT Press, ISBN : 9781461536192, 2018.					