1.	Title of the course	Stochastic Decision Processes
2.	Course number	CS521L
3.	Structure of credits	3-0-0-3
4.	Offered to	PG
5.	New course/modification to	Modification To CS5225/12
6.	To be offered by	Department of Computer Science and Engineering
7.	To take effect from	July 2022
8.	Prerequisite	CoT
9.	**Course Objective(s):** To enable the modeling of sequential decision making problems from different engineering applications. To provide the principles and algorithms for deriving optimal solutions for sequential problems.	
10.	**Course Content:** Basics: probability space, conditional probability, random variables, expectation and linearity of expectation, law of large numbers, central limit theorem, discrete and continuous time Markov chains, absorbing Markov chains; Finite horizon problems: Markov Decision Process (MDP) model formulation, optimality criteria, policy evaluation, Bellman optimality equations, optimality of deterministic Markov policies, backward induction; Infinite horizon MDP problems: expected total-reward criterion, expected discounted-reward criterion, Markov policies, policy evaluation, value iteration, policy iteration, linear programming; Optimal stopping problems; Introduction to partially observable MDP.	
11.	**Textbook(s):** 1. Bertsekas D P, *Dynamic Programming and Optimal Control, Volume-I*, 4th Edition, Athena Scientific (2017). 2. Puterman M L, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, 1st Edition, John Wiley and Sons (1994).	
12.	**Reference(s):** 1. Bertsekas D P, *Dynamic Programming and Optimal Control, Volume-II*, 4th Edition, Athena Scientific (2017). 2. Bertsekas D P, *Abstract Dynamic Programming*, 2nd Edition, Athena Scientific (2018). 3. Filar J and Vrieze K, *Competitive Markov Decision Processes*, Springer-Verlag (1996). 4. Sutton R S and Barto A G, *Reinforcement Learning - An Introduction* , 2nd Edition, MIT Press (2018).	