Programme for **An Introduction to Stochastic Control and Reinforcement Learning**

Day 1:

Finite state Markov chain (discrete-time), Transition probabilities, Markov decision processes (MDP) (controlled Markov chain) and their applications, discrete-time stochastic control formulation with applications (1.5 hrs)

Finite horizon stochastic control problem, principle of optimality (Bellman equation), Dynamic Programming and its solutions, Closed form solution for the Linear Quadratic Gaussian (LQG) control problem (1.5 hrs)

Infinite horizon stochastic control problems (discounted and average cost with finite state and action space), Bellman optimality equation, existence of stationary control policy, solution methodologies for these problems – policy iteration and value iteration and related algorithms (3 hrs with a break)

Day 2:

Curse of dimensionality in solving Dynamic Programming algorithms, Approximate Dynamic Programming algorithms – approximation in policy space and value space, contraction properties and error bounds, simulation based implementation (3 hrs with a break)

Intro to reinforcement learning in the setting of MDP, Temporal difference methods (TD(0), TD(λ)), and their convergence properties, applications, On-policy TD control (SARSA), Off-policy TD control such as Q-learning and its convergence properties, Applications (3 hrs with a break)

Day 3:

Advanced reinforcement learning, Value function approximation with Linear methods and function approximation, deep reinforcement learning, policy gradient methods, actor-critic based reinforcement learning and their applications to continuous control (such as LQG) problems (3 hrs with a break)