Workshop on Learning and Control
IIT Mandi, 22 – 26 July, 2019

Organizer: Control Society

Schedule

Slot	Monday	Tuesday	Wednesday	Thursday	Friday
A1	MV-1	PPK-2	SPB-1	PPK-4	MV-3
A2	PPK-1	MV-2	PPK-3	PVR-2	PVR-3
P1	PVR-1	Demos	Free	SPB-2	Demos
P2	Demos	Demos	Free	Demos	Open Forum

Details of Speakers:

SPB = Sanjay P. Bhat, Tata Consultancy Services

PPK = Pramod P. Khargonekar, University of California at Irvine

PVR = Puduru Viswanadha Reddy, Indian Institute of Technology Madras

MV = M. Vidyasagar, Indian Institute of Technology Hyderabad

Titles of Lectures:

SPB-1: Multi-Arm Bandit Problems

SPB-2: Multi-Agent Reinforcement Learning

PPK-1: Deep Learning: From Foundations to Current Research

PPK-2: Advanced Topics in Deep Learning: Generative Adversarial Networks

PPK-3: Reinforcement Learning: Dynamic Programming to Q-Learning

PPK-4: Recent Advances in Reinforcement Learning

VR-1: Game Theory: Basic Notions and Solution Concepts

VR-2: Game Theory: Nash Equilibrium and Refinements of Nash Equilibrium

VR-3: Learning in Games: Fictitious Play

MV-1: A Historical Overview of AI and Machine Learning

MV-2: Compressed Sensing: Vector Recovery

MV-3: Compressed Sensing: Matrix Recovery

Brief Description of Lectures:

SPB-1: This lecture will begin with some examples of real world problems which fall into the stochastic multi-arm bandit (MAB) framework. These problems involve an agent who has to repeatedly choose between a set of alternatives that yield noisy rewards. The agent would like to sequentially select alternatives so as to maximize her total reward, given only the outcomes of past selections. In this lecture, we will understand the exploration versus exploitation dilemma that such an agent comes to face, introduce the notion of regret as a way of quantifying this dilemma, and look at a few well-known algorithms and their regret performance.

SPB-2: In this lecture, we will see how some of the ideas, tools and techniques of single-agent reinforcement learning carry over to multi-agent situations. Furthermore, we will look at a few well known algorithms.

PPK-1: In this lecture, I will give an overview of deep learning. I will begin with the fundamental result on the single hidden layer neural network. This will be followed by a presentation of the multi-layer neural networks training and validation. I will describe some salient applications of deep learning. The lecture will conclude with a discussion of important open research questions.

PPK-2: In this lecture, I will continue from where the first lecture ends. Generative Adversarial Networks (GANs) refers to the use of a pair of neural networks, generative and discriminiative, to represent models to generate training data. We will begin with motivations for studying GANs. We will then discuss methods for training the generative and discriminative models. The lecture will end with some key applications of GANs.

PPK-3-4: This lecture will give a succinct overview of fundamentals of reinforcement learning. We will cover dynamic programming, value and policy iterations, Q-learning, etc. After these fundamentals, we will discuss deep reinforcement learning and its recent applications. We will connect materials in Lectures 1-3 to future controls research.

VR-1: During the past decade, advances in information and communication technologies have resulted in the emergence of large-scale, distributed, networked and heterogeneous systems. Some applications include, demand-side management in smart grids, pricing and investment decisions in the internet, and security issues in cyber-physical systems. Game theory has emerged as a powerful tool to analyze these systems. In this lecture, we will start with basic non-cooperative game models such as strategic form games and extensive form games. We will then discuss solution concepts such as security strategies, dominant strategy equilibrium, and iterative elimination of dominated strategies.

VR-2: In this lecture, we will discuss the concept of Nash equilibrium, and refinements such as subgame perfect Nash equilibrium and Bayesian Nash equilibrium. We will then review some results on the existence and uniqueness of Nash equilibrium.

VR-3: In this lecture, we will discuss explanations as to why an equilibrium can arise through dynamic process where rational agents search for optimal strategies. We will discuss a learning rule called fictitious play and its convergence properties.

MV-1: In this lecture I will give an overview of the historical evolution of Artificial Intelligence (AI) as a separate discipline, starting in the late 1940s until now. Specifically, I will discuss how the aims of AI have shifted over the years, from trying to mimic human capabilities to replicating the functioning of the human brain and back again to mimicking human capabilities.

MV-2: Compressed sensing refers to the recovery of high-dimensional but low-complexity entities from a small number of linear measurements, via solving a convex optimization problem. Some examples of compressed sensing are: recovering a high-dimensional but sparse vector, and recovering a high-dimensional but low-rank matrix. Applications of compressed sensing include the reconstruction of bandwidth-limited signals from sampling at a few frequencies, sensor localization, recommendation engines, and the like. In this lecture I will discuss the problem of vector recovery.

MV-3: In this lecture I will discuss the problem of matrix recovery.