Safexpress Centre for Data, Learning and Decision Sciences

exSPLORe 2025: Schedule

14th-18th January, 2025

Day 1: 14th January, 2025

Time	Speaker	Title
9:30 AM - 10:00 AM	Pramath R. Sinha, Chairperson, Board of Trustees; Somak Raychaudhury, Vice Chancellor; Sandeep Juneja, SCDLDS Director	Opening Remarks
10:00 AM - 11:30 AM	Jayakrishnan Nair (IIT Bombay)	Constructing Confidence Intervals - 1
11:30 AM - 11:50 AM	Break	
11:50 AM - 1:20 PM	Ketan Rajawat (IIT Kanpur)	Introduction to Optimization: Algorithms & Implementation - 1
1:20 PM - 2:30 PM	Lunch	
2:30 PM - 4:00 PM	Jayakrishnan Nair	Constructing Confidence Intervals - 2
4:00 PM - 4:30 PM	Break	
4:30 PM - 5:30 PM	* Ankur Puri (McKinsey)	AI from an Executive's Lens
9:00 PM - 10:00 PM	Balaji Prabhakar - Online (Stanford University)	What Is A Good Network And How Can We Build One?

Day 2: 15th January, 2025

Time	Speaker	Title
9:30 AM - 11:00 AM	Praneeth Netrapalli (Google DeepMind)	Introduction to Reverse Diffusion Sampling - 1
11:00 AM - 11:20 AM	Break	
11:20 AM - 12:50 PM	Ketan Rajawat	Introduction to Optimization: Algorithms & Implementation-2
12:50 PM - 2:00 PM	Lunch	
2:00 PM - 3:00 PM	Poster Presentation	
3:00 PM - 4:30 PM	Praneeth Netrapalli	Introduction to Reverse Diffusion Sampling - 1
4:30 PM - 5:00 PM	Break	
5:00 PM - 6:00 PM	* Manish Gupta (Google DeepMind)	The Transformative Power of AI and Open Challenges
6:00 PM - 7:00 PM	Poster Presentation contd.	
9:00 PM - 10:00 PM	Peter Glynn - Online (Stanford University)	Incorporating Forecasts Into MDPs

* Industry Talk

Day 3: 16th January, 2025

Time	Speaker	Title
9:30 AM - 10:30 AM	† Vijay Vazirani (University of California, Irvine)	The Assignment Game: Equitable Core Imputations
10:30 AM - 10:50 AM	Break	
10:50 AM - 11:50 AM	† Ronnie Sircar (Princeton University)	ORFEUS - Operational Risk Financialization of Electricity under Stochasticity
11:50 AM - 12:40 PM	Ruta Mehta (University of Illinois at Urbana-Champaign)	Computability of Competitive Equilibrium with Chores
12:40 PM - 2:00 PM	Lunch	
2:00 PM - 2:50 PM	Siddhartha Banerjee (Cornell University)	Artificial Replay: How to get the most out of historical data in decision-making
2:50 PM - 3:40 PM	Achal Bassamboo (Northwestern University)	Asymptotically Optimal Adaptive A/B Tests for Average Treatment Effect
3:40 PM - 4:00 PM	Break	
4:00 PM - 4:50 PM	Sunita Sarawagi - Online (IIT Bombay)	Understanding the Emergence of In-Context Learning in LLMs

† Keynote Talk

Day 4: 17th January, 2025

Time	Speaker	Title
9:30 AM - 10:30 AM	† Garud Iyengar (Columbia University)	Online Matching with Heterogeneous Supply and Minimum Allocation Guarantees
10:30 AM - 10:50 AM	Break	
10:50 AM - 11:50 AM	† Rama Cont (Oxford University)	Asymptotic analysis of Deep Residual Networks
11:50 AM - 12:40 PM	Srikanth Iyer (IISc)	Static Hedging of European Options
12:40 PM - 2:00 PM	Lunch	
2:00 PM - 2:50 PM	Varun Gupta (Northwestern University)	Greedy Algorithm for Multiway Matching with Bounded Regret
2:50 PM - 3:40 PM	Siddharth Barman (IISc)	Collective Welfare as a Metric in Algorithmic Decision Making
3:40 PM - 4:00 PM	Break	
4:00 PM - 4:50 PM	Rahul Vaze (TIFR)	Optimal Bounds for Constrained Online Convex Optimization
6:00 PM - 7:30 PM	Cultural Event	
8:00 PM onwards	Workshop Dinner	

† Keynote Talk

Day 5: 18th January, 2025

Time	Speaker	Title
9:30 AM - 10:30 AM	Dheeraj Nagaraj (Google DeepMind)	Training and Sampling Algorithms for Diffusion Models
10:30 AM - 10:50 AM	Break	
10:50 AM - 11:50 AM	Vivek Borkar (IIT Bombay)	Reinforcement Learning: From Markov to Non-Markov
11:50 AM - 12:40 PM	Sanjay Bhat (TCS Research)	ExSPLORations in bandit optimization on continuous decision spaces
12:40 PM - 2:00 PM	Lunch	
2:00 PM - 2:50 PM	Rohit Vaish (IIT Delhi)	Fair Scheduling of Indivisible Chores
2:50 PM - 3:40 PM	Alexandre Proutiere (KTH Royal Institute of Technology)	Low-rank Reinforcement Learning
3:40 PM - 4:00 PM	Break	
4:00 PM - 5:00 PM	Devavrat Shah (Massachusetts Institute of Technology)	Various Incarnations of Matrix Completion

Emergency contacts

- Infirmary +918199977075
- Accommodations +917082000568

Abstracts

Siddhartha Banerjee

Title: Artificial Replay: How to get the most out of historical data in decision-making
Abstract: How best to incorporate historical data for designing decision-making policies is an important open question for using online control and reinforcement learning in practice: more data should mean better performance, but naive use of historical samples has high computation and storage cost, and may still lead to bad policies.

To get around this, I will propose a simple paradigm called Artificial Replay, that is based on the following core insight: using historical samples, we can infer counterfactual consequences, that then guide policy decisions. I will first illustrate this for simpler 'bandit' settings, showing how our approach uses a fraction of the data while achieving identical regret compared to a full warm-start policy. Next I will extend this paradigm to a much more general class of problems online decision-making settings with state, where the uncertainty in the system can be represented as being exogenous to the system state. Here, I will discuss how we can show formal regret guarantees for such systems using a technique called the compensated coupling, and also demonstrate its use in a variety of settings, where it outperforms state-of-the-art RL algorithms.

Siddharth Barman

Title: Collective Welfare as a Metric in Algorithmic Decision Making
Abstract: Regret minimization is a pre-eminent objective in the study of decision making under uncertainty. Indeed, regret is a central notion in multi-armed bandits, reinforcement learning, game theory, decision theory, and causal inference. In this talk, I will present our recent work that extends the formulation of regret with a welfarist perspective.

In particular, we quantify the performance of a decision maker by applying a fundamental welfare function–namely the Nash social welfare (NSW)–and study Nash regret, defined as the difference between the (a priori unknown) optimum and the decision maker's performance. Since NSW is known to satisfy fairness axioms, our approach complements the utilitarian considerations of average (cumulative) regret, wherein the algorithm is evaluated via the arithmetic mean of its expected rewards. I will present our recent works that obtains essentially tight Nash regret guarantees for stochastic multi-armed bandits (MAB) as well as linear bandits.

Achal Bassamboo

Title: Asymptotically Optimal Adaptive A/B Tests for Average Treatment Effect
Abstract: Typically, in A/B testing, an experiment designer sequentially assigns treatment A or B to arriving individuals to identify the better treatment (known as the best treatment identification (BTI) problem), that is, the one with better mean performance, with the probability of error restricted to a pre-specified δ. We focus on a related, equally important, but more informative problem of estimating the difference between the two means or the average treatment effect (ATE). For computational efficiency, we restrict accuracy to a confidence interval (CI) of width at most ϵ, where again the probability of CI not containing ATE is restricted to at most δ. The objective of the experiment

designer is to estimate a CI of ATE with the minimum sample size, i.e., minimize the total number of individuals getting treatment A or B. We first establish a lower bound on the expected sample size of the A/B test (experiment) needed for any adaptive experimental policy, which constructs a CI of ATE with desired properties as the solution to a max-min optimization problem for small δ. Using the insights provided by the max-min optimization problem, we construct an adaptive policy that is asymptotically optimal, i.e., matches the lower bound on the expected sample size for small δ. To reduce the computational burden of our policy, we propose another adaptive policy that is asymptotically optimal for small ϵ and δ. We find that, for small ϵ and δ, the asymptotically optimal fraction of treatment assignment for A and B is inversely proportional to the square root of Fisher's information of the outcome distributions of treatments A and B, respectively. Further, we compare the popular randomized controlled policy with any asymptotically optimal adaptive policies and show that there are meaningful gains from any asymptotically optimal adaptive policy in terms of the length of the experiment. Finally, we present a comparative analysis between our ATE problem and the BTI problem revealing marked differences in the asymptotically optimal assignment of treatments in both ATE and BTI problems. (Joint work with Vikas Deep and Sandeep Juneja.)

Sanjay Bhat

Title: ExSPLORations in bandit optimization on continuous decision spaces
Abstract: To be Announced

Vivek Borkar

Title: Reinforcement Learning: From Markov to Non-Markov
Abstract: This talk will give a rapid overview of stochastic approximation algorithms and reinforcement learning as an application thereof, followed by some recent results about reinforcement learning in non-markovian environments.

Rama Cont

Title: Asymptotic analysis of Deep Residual Networks
Abstract: Residual Networks (ResNets) are neural network architectures which have been successfully used for image and speech recognition tasks. They have also been the focus of much theoretical interest due to their (assumed) link with "neural differential equations". We rigorously investigate the asymptotic properties of deep ResNets as the number of layers increases. We first show the existence of scaling regimes for trained weights markedly different from those implicitly assumed in the "neural ODE" literature. We study the convergence of the hidden state dynamics in these scaling regimes, showing that one may obtain an ODE, a stochastic differential equation (SDE) or neither of these. In particular, our findings point to the existence of a diffusive regime in which the deep network limit is described by a class of forward-backward stochastic differential equations (FBSDEs).

Peter Glynn

Title: Incorporating Forecasts Into MDPs
Abstract: In an increasing number of applications contexts, one has access to daily, monthly, or quarterly forecasts of some exogenous quantity that strongly affects the system to be controlled. For example, in the setting of the district level heating and cooling systems that are used to control temperatures of buildings within the "district", hourly and daily temperature forecasts localized to the district may be available. The optimal control of the amount of heated and chilled water to be sent to the buildings is strongly influenced by these forecasts. In this setting, it is highly desirable that the dynamics of the system, when conditioned on the forecasts, remain low dimensional. We also discuss theoretical implications of our martingale-based forecast modeling approach for the theory underlying model predictive control. This is joint work with Jacques de Chalendar and Haoran Xu. Our joint work is motivated by our efforts to help support more efficient operation of Stanford University's district heating and cooling system.

Manish Gupta

Title: The Transformative Power of AI and Open Challenges
Abstract: We begin by presenting the recent advances in the area of artificial intelligence, and the high level ideas underlying the progressively narrower domains of machine learning, deep learning, and foundation models, which have emerged over time as dominant paradigms for artificial intelligence. We describe the tremendous progress of these models on problems ranging from understanding, prediction and creativity on one hand, and open technical challenges like safety, fairness and transparency on the other hand. These challenges are further amplified as we seek to advance Inclusive AI to tackle problems for billions of human beings in the context of the Global South. We present our work on multilingual models to democratize information access in a diverse set of Indian languages, on healthcare in environments where we lack data in digital form to begin with, and on analysis of satellite imagery to help transform agriculture and improve the lives of farmers. Through these examples, we hope to convey the excitement of the potential of AI to make a difference to the world, and also a fascinating set of open problems to tackle.

Varun Gupta

Title: Greedy Algorithm for Multiway Matching with Bounded Regret
Abstract: We consider a finite horizon online resource allocation/matching problem where the goal of the decision maker is to combine resources (from a finite set of resource types) into feasible configurations. Each configuration is specified by the number of resources consumed of each type and a reward. The resources are further subdivided into three types - offline, online-queueable (which arrive online and can be stored in a buffer), and online-nonqueueable (which arrive online and must be matched on arrival or lost). We prove the efficacy of a simple greedy algorithm when the corresponding static planning linear program (SPP) exhibits a non-degeneracy condition called the general position gap (GPG). In particular we prove that, (i) our greedy algorithm gets bounded any-time regret when no configuration contains both an online-queueable and an online-nonqueueable resource, and (ii) O(log t) expected any-time regret otherwise

(we also prove a matching lower bound). By considering the three types of resources, our matching framework encompasses several well-studied problems such as dynamic multi-sided matching, network revenue management, online stochastic packing, and multiclass queueing systems.

Srikanth Iyer

Title: Static Hedging of European Options
Abstract: We consider the problem of static hedging of a target European option using options of shorter maturities. In the first part, we assume that the underlying asset follows a single-factor Markovian model. We extend a result of Carr and Wu (2014) to include options of various short maturities. In the second part of the talk, we consider this problem in the model-free or robust framework. The problem is postulated as a min-max problem where the inner maximization problem is a modified martingale optimal transport problem. Here we present some preliminary results in this direction.

Garud Iyengar

Title: Online Matching with Heterogeneous Supply and Minimum Allocation Guarantees
Abstract: We propose a novel modeling framework for two-sided matching platforms that accounts for worker churn. Existing literature suggests that workers churn because they are not allocated enough work. Our model allows the workers to be heterogeneous in their minimum requirements (in addition to their heterogeneous quality and capacity) and provides the decision maker the ability to control churn by guaranteeing an appropriate amount of work to each worker. In each period, a random number of jobs sequentially arrive, and have to be matched to an available worker in an online manner. The matching policy maximizes the steady state average match quality. Churn plays a critical role here because, in its absence, the Greedy policy that matches a job to the highest quality worker with residual capacity is optimal. We show that policies that are blind to the quality of the workers (e.g., first-come-first-serve (FCFS)) or do not account for the heterogeneity in the workers' preferences (e.g., Greedy) can perform arbitrarily badly in the worst case. We propose a policy (Tracker) that tracks a target allocation level for each worker by allocating an incoming job to the most under-allocated worker with residual capacity. The target levels are computed by solving a carefully constructed static deterministic relaxation that results in a non-linear binary program with exponentially many constraints. In addition to establishing the optimality of Tracker via a careful Lyapunov-style analysis, we develop an efficient algorithm to compute a solution of the non-linear binary program with provable performance guarantees. We perform a simulation study calibrated to the data from a specific labor platform. The proposed policy (Tracker) dominates the two benchmark policies (FCFS and Greedy) in every scenario we tested, with gains of over 10% in many scenarios. Our work influenced the labor platform to change their matching policy, and they observed gains of a similar magnitude.

Ruta Mehta

Title: Computability of Competitive Equilibrium with Chores
Abstract: Competitive equilibrium, also known as market equilibrium, is arguably one of the most fundamental solution concepts within Economics. The existence and computability of CE have been extensively studied when all items are disposable goods. However, the problem is less explored when some or all items are non-disposable chores (bads), despite being equally relevant, for example, the various labor markets like Ola/Uber/Taskrabbit. In this talk, I will discuss recent algorithmic advances on the computation of CE when the item set includes chores.

Surprisingly, this problem stands in sharp contrast to the goods-only case, even under linear (additive) utility functions: for the case of goods, the CE set is known to be a convex set, while with chores, it may be non-convex and disconnected. I will discuss how to handle this non-convexity through new (continuous) optimization methods, and via a novel non-convex formulation, leading to fast algorithms for computing approximate CE. These methods may be of independent interest for finding local optima of certain classes of non-convex programs.

Jayakrishnan Nair

Title: Constructing Confidence Intervals
Abstract: Starting with sub-Gaussian distributions, moving to sub-Exponential/Light-tailed distributions, and finally moving to the heavy-tailed setting, where estimators other standard empirical ones (like median-of-means) are preferable. Along the way, some bandit-style hypothesis testing applications will be introduced.

Dheeraj Nagaraj

Title: Training and Sampling Algorithms for Diffusion Models
Abstract: Diffusion models have been extremely successful in the generative modeling of images and videos. In this talk, I will first discuss the training and inference algorithms for Diffusion models. I will then introduce my recent work on improving these algorithms with mathematically principled approaches.

Praneeth Netrapalli

Title: Introduction to Reverse Diffusion Sampling
Abstract: In this tutorial, we will develop the theory behind reverse diffusion sampling. Specifically, we will show how to convert between particle flows and probability flows, use this to present a simple proof of convergence of Langevin dynamics for strongly log concave densities, and then motivate score matching and reverse diffusion sampling. Time permitting, we will also briefly extend this theory to flow matching approaches for sampling.

Balaji Prabhakar

Title: What Is A Good Network And How Can We Build One?
Abstract: The answer depends on when the question was asked. In the mid-90s a good network was one which provided connectivity, leading to the rapid proliferation of TCP/IP "best-effort networks" and access technologies like DSL and WiFi. In the mid 2000s, a good network was one which enabled the convergence of local area networks (LANs), storage area networks (SANs) and high-performance computing (HPC) networks. This led to the birth of data centers whose interconnect networks had to simultaneously provide high bandwidth and low latency. In the 2010s, a good network was one which supported server virtualization, leading to the Cloud Computing revolution. A good network in this case was one which supported multiple tenants running heterogeneous applications and was ubiquitously accessible. Today, a good network is one which interconnects large GPU clusters and provides the highest possible performance. Concretely, this means providing close to 95 percent network utilization at near zero packet drops or delays—a far cry from the best-effort service provided by TCP/Ethernet networks.

In this talk, we will describe the quest for the "good network" over time. We will contrast the network- and edge-centric approaches to building a good network—the former needing special network hardware and protocols while the latter doesn't. We will show how accurate clock synchronization provides a foundation for building good networks which provide a very high performance for both the CPU and GPU clouds without requiring modifications to commodity network hardware.

Alexandre Proutier

Title: Low-rank Reinforcement Learning
Abstract: To be announced

Ankur Puri

Title: AI from an Executive's Lens
Abstract: How do executives, especially senior managers in traditional organizations , view AI? What are the applications they find valuable? What has been their experience in using AI? What have we learnt about successful adoption?

Whether you develop new algorithms, or engineer new systems that deliver the power of AI to enterprises, or develop applications that solve the problems faced by businesses, the speaker hopes that these perspectives from the frontier of application in the real world inspire and inform you.

Ketan Rajawat

Title: Introduction to Optimization: Algorithms and Implementation
Abstract: In this hands-on tutorial, we cover both fundamental principles as well as practical aspects of optimization. The session will begin with a foundational overview, setting the stage with essential preliminaries and the importance of optimization across domains. We will delve into convexity, exploring its pivotal role in ensuring efficient and reliable solutions. The tutorial then covers gradient descent, elucidating its theoretical

underpinnings and practical applications. Finally, we examine stochastic gradient descent (SGD), a cornerstone algorithm in machine learning, with insights into its stochastic nature and convergence properties.

Participants will gain both theoretical insights and programming experience. We will provide interactive Google Colab notebooks, enabling attendees to experiment with code and observe optimization algorithms in action. These hands-on exercises are designed to bridge the gap between theory and practice. To make the most of this tutorial, participants are encouraged to bring a laptop/tablet where they can access Google Colab. This session is ideal for students, researchers, and professionals eager to grasp the fundamentals of optimization and apply these techniques effectively in programming contexts. Basic knowledge of python is necessary.

Sunita Sarawagi

Title: Understanding the Emergence of In-Context Learning in LLMs
Abstract: Large language models (LLMs), trained via next-token prediction on large-scale natural language corpora, demonstrate the surprising ability to perform in-context learning (ICL)—adapting to new tasks with just a few examples in context, despite not being explicitly trained for this capability. The mechanisms behind ICL remain poorly understood, leading to several competing hypotheses. In this talk, we will explore three explanations for ICL: unrolled gradient descent across transformer layers, Bayesian task selection, and associative recall through induction circuits. We will also discuss cases where in-context learning fails and suggest potential directions for future research.

Devavrat Shah

Title: Various Incarnations of Matrix Completion
Abstract: The objective of matrix completion is to estimate or complete an unknown matrix from its partial, noisy observations. Since its introduction as a model for recommendation systems in the early 1990s, it has been central to advances in machine learning, statistics, and applied probability. In this talk, I will discuss a few incarnations of it that arise in the context of time-series analysis, causal inference, reinforcement learning and empirical risk minimization.

Ronnie Sircar

Title: ORFEUS - Operational Risk Financialization of Electricity under Stochasticity
Abstract: We describe part of a multi-year project, sponsored by ARPA-E, to quantify, allocate and account for the risk introduced to electricity grids due to the unpredictable reliability of production from renewables. Incorporating this stochasticity into grid risk management is viewed by the industry (which has remained almost entirely tethered to a deterministic viewpoint, and in particular to weather forecasts) as increasingly crucial, as we aim for greater renewables penetration to reduce dependence on carbon-emitting fuels. Our methodology involves feeding Monte Carlo simulations of solar, wind and demand into a grid optimization software that emulates the performance and costs of

the Texas electricity grid. This outputs a distribution of running costs, from which we can numerically extract a measure of system (grid) risk. The more challenging part is to allocate this risk back (top down) to the individual renewable assets to assign them a reliability cost. This adapts existing approaches for the risk allocation problem related to Shapley values, but is computationally intense. We show results, project to potential future grids, and propose a way to incorporate the reliability costs back into the day ahead bid curve and thereby re-optimize unit commitment and economic dispatch of assets taking into (some) account the stochasticity.

Rohit Vaish

Title: Fair Scheduling of Indivisible Chores
Abstract: We will discuss the problem of fairly assigning a set of discrete tasks, or chores, among a set of agents. Each chore has a designated start and finish time, and each agent can perform at most one chore at any time. We will explore the existence and computation of "fair" (specifically, envy-free up to one chore) and "efficient" (specifically, maximal or Pareto optimal) schedules under various settings. The presentation will cover novel technical ideas, including a color-switching technique and an application of the "cycle-plus-triangles" theorem (originally conjectured by Erdős) for achieving approximate envy-freeness. We will also highlight several open problems and directions for future work.

Rahul Vaze

Title: Optimal Bounds for Constrained Online Convex Optimization
Abstract: A well-studied generalization of the standard online convex optimization (OCO) is the constrained online convex optimization (COCO). In COCO, on every round, a convex cost function and a convex constraint function are revealed to the learner after the action for that round is chosen. The objective is to design an online policy that simultaneously achieves a small regret while ensuring a small cumulative constraint violation (CCV). A long-standing open question in COCO is whether an online policy can simultaneously achieve $O(\sqrt{T})$ regret and $O(\sqrt{T})$ CCV unconditionally. For the first time, we answer this in the affirmative, where surprisingly, the analysis is short and elegant.

Vijay Vazirani

Title: The Assignment Game: Equitable Core Imputations
Abstract: The classic 1971 paper of Shapley and Shubik characterized the core of the assignment game. We observe that a sub-coalition consisting of one player (or a set of players from the same side of the bipartition) can make zero profit, and therefore its profit under a core imputation can be an arbitrary amount. Hence an arbitrary core imputation makes no fairness guarantee at the level of individual agents. Can this deficiency be addressed by picking a "good" core imputation?

To arrive at an appropriate solution concept, we give specific criteria for picking a special core imputation, and we undertake a detailed comparison of four solution concepts. Leximin and leximax core imputations come out as clear winners; we define these to be equitable core imputations. These imputations achieve "fairness" in different ways:

whereas leximin tries to make poor agents more rich, leximax tries to make rich agents less rich.

We give combinatorial strongly polynomial algorithms for computing these imputations via a novel adaptation of the classical primal-dual paradigm. The "engine" driving them involves insights into core imputations obtained via complementarity. It will not be surprising if our work leads to new uses of this powerful technique. Furthermore, we expect more work on computing the leximin and leximax core imputations of other natural games.