Safexpress Centre for Data, Learning and Decision Sciences

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

Day 1: January 13, 2026 — Venue: AC-02-LR-007

Time	Speaker	Title
9:45 to 10:00	Somak Raychaudhury, Vice Chancellor; Sandeep Juneja, SCDLDS Director	Opening remarks
10:00 to 11:00	†Peter Glynn (Stanford University)	Online Linear Programming with Batching
11:00 to 11:30	Break	
11:30 to 13:00	Arun Sai Sugalla (Google DeepMind)	Understanding State-of-the-Art LLMs: From Architecture to Reasoning - I
13:00 to 14:30	Lunch (venue: First floor, Dining Hall)	
14:30 to 16:00	Dheeraj Nagaraj (Google DeepMind)	Diffusion and Flow models for Generative Modeling - I
16:00 to 16:30	Break	
16:30 to 18:00	Arun Sai Sugalla (Google DeepMind)	Understanding State-of-the-Art LLMs: From Architecture to Reasoning - II

† Keynote talk

Safexpress Centre for Data, Learning and Decision Sciences

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

Day 2: January 14, 2026 — Venue: AC-02-LR-007

Time	Speaker	Title
9:30 to 11:00	Aditya Gopalan (IISc)	Multi-armed Bandits & Reinforcement Learning - I
11:00 to 11:30	Break	
11:30 to 13:00	Dheeraj Nagaraj (Google DeepMind)	Diffusion and Flow models for Generative Modeling - II
13:00 to 14:30	Lunch (venue: First floor, Dining Hall)	
14:30 to 16:00	Aditya Gopalan (IISc)	Multi-armed Bandits & Reinforcement Learning - II
16:00 to 16:30	Break	
Small talk		
16:30 to 16:45	Bismark Singh (University of Southampton)	Small talk: Balancing user accessibility and facility load fairness in undesirable facility location problems
16:45 to 17:00	Shubhada Agrawal (IISc)	Small talk: TBA
17:00 to 18:30	†Nisheeth Vishnoi (Yale University)	Creativity, Discovery, and Learning in the Age of AI

† Keynote talk

Safexpress Centre for Data, Learning and Decision Sciences

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

Day 3: January 15, 2026 — Venue: AC-02-LR-007

Time	Speaker	Title
9:30 to 10:30	†Venkat Padmanabhan (Microsoft Research)	From FLOPS to Hits: An Overview of (AI) Explorations at Microsoft Research India
10:30 to 11:00	Break	
11:00 to 12:15	Panel Discussion: Peter Glynn (Stanford University), Venkat Padmanabhan (Microsoft Research), Alexandre Proutiere (KTH), and Nisheeth Vishnoi (Yale University). Moderated by: Gautam Menon (Ashoka University).	Why I became a researcher?
12:15 to 12:30	Group photo session	
12:30 to 15:00	Lunch and industry stalls (venue: Lawn near the atrium)	
15:00 to 16:15	Panel discussion: Arko Chaterjee (Bains and Company), Peter Glynn (Stanford University), Ashish Goel (Stanford University), Gautam Shroff (IIIT Delhi) Moderated by: Debayan Gupta (Ashoka University)	Collaborations between academia and industry - What works and what doesn't

† Keynote talk

Safexpress Centre for Data, Learning and Decision Sciences

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

Day 4: January 16, 2026 — Venue: AC-02-LR-007

Time	Speaker	Title
9:30 to 10:30	†Ashish Goel (Stanford University)	Decentralized Finance: Theory, Practice, and Architectures
10:30 to 11:00	Break	
11:00 to 11:50	Sarah Cen (Carnegie Mellon University)	Bridging the Gap Between Research and Policy in AI Safety and Accountability
11:50 to 12:20	Break	
12:20 to 13:10	Andrew Ilyas (Carnegie Mellon University)	TBD
13:10 15:20	Poster presentations and lunch (venue: Lawn near the atrium)	
15:20 to 16:10	Karthyek Murthy (USC Los Angeles)	The Scaling Behaviors in Achieving High Reliability via Chance-constrained Optimization
16:10 to 17:00	Alexandre Proutiere (KTH Royal Institute of Technology)	Emergence of low-rank structures in Reinforcement Learning
17:00 to 17:15	Break	
Small talk		
17:15 to 17:30	Sumit Kunnumkal (ISB)	Strawberry Or Vanilla This Week? How To Optimize Tailored Assortments.For

Safexpress Centre for Data, Learning and Decision Sciences

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

		Variety-Seeking/Avoiding Consumers
17:30 to 17:45	Sunny Manchanda (DRDO)	TBD
18:30 to 19:30	Cultural event	
20:00 onwards	Workshop dinner (buses will leave at 19:45 from Gate #2)	

† Keynote talk

Safexpress Centre for Data, Learning and Decision Sciences

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

Day 5: January 17, 2026 — Venue: AC-02-LR-007

Time	Speaker	Title
9:30 to 10:30	†Assaf Zeevi (Columbia University)	Thompson Sampling through an Optimization Lens
10:30 to 11:00	Break	
11:00 to 11:50	Dootika Vats (IIT Kanpur)	MCMC Importance Sampling via Moreau-Yosida Envelopes
11:50 to 12:10	Break	
12:10 to 13:00	Vineet Goyal (Columbia University)	TBD
13:00 to 14:30	Lunch (venue: First floor, Dining Hall)	
Presentations by exSPLOre 26 Best Young Researcher Paper Competition finalists		
14:30 to 14:40	Achal Bassamboo (Northwestern University)	Award overview by a chair
14:40 to 15:05	Agniv Bandyopadhyay (TIFR)	Optimal Top-Two Method for Best Arm Identification and Fluid Analysis
15:05 to 15:30	Santanu Das (TIFR)	A Direct Proof of a Unified Law of Robustness: Bregman Divergence Losses
15:30 to 15:55	Piyushi Manupriya (IISc)	Multi-agent Multi-armed Bandits with Minimum Reward Guarantee Fairness
15:55 to 16:20	Soumen Pachal (TCS Research,	Generalized Simultaneous

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

	Chennai, India)	Perturbation-Based Gradient Search With Reduced Estimator Bias
16:20 to 16:45	Tamojit Sadhukhan (ISI Kolkata)	Changepoint Estimation in Sparse Dynamic Stochastic Block Models under Near-Optimal Signal Strength
16:45 to 16:55	Achal Bassamboo (Northwestern University)	Announcement of winners

† Keynote talk

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

Speaker : Sarah Cen

Affiliation : USC Los Angeles

Title : Bridging the Gap Between Research and Policy in AI Safety and Accountability

Abstract : As AI becomes increasingly integrated into both the private and public sectors, challenges around AI safety and accountability have arisen. There is a growing, compelling body of work around the legal and societal challenges that come with AI, but there is a gap in our rigorous understanding of these problems. In this talk, I dive deep into a few topics in AI safety and accountability. We will discuss AI supply chains (the increasingly complex ecosystem of AI actors and components that contribute to AI products) and study how AI supply chains complicate machine learning objectives. We'll then shift our discussion to AI audits and evidentiary burdens in cases involving AI. Using Pareto frontiers as a tool for assessing performance-fairness tradeoffs, we will show how a closed-form expression for performance-fairness Pareto frontiers can help plaintiffs (or auditors) overcome evidentiary burdens or a lack of access in AI contexts. I'll conclude with a longitudinal study of LLMs during the 2024 US election season. If time permits, we may touch on formal notions of trustworthiness.

Speaker : Peter W. Glynn

Affiliation : Stanford University

Title : Online Linear Programming with Batching

Abstract: The online linear programming problem is a resource allocation problem in which resources need to be allocated to orders sequentially and irrevocably, with the goal of maximizing total reward. This problem generalizes the Multi-Secretary Problem, and arises naturally in many applied settings. We will briefly survey what is known about this decision problem when the sequential stream of orders is modeled as a sequence of random vectors. When the resources scale in proportion to the number n of orders, then a natural measure of regret scales as O(log n) when appropriate algorithms are applied. We will then describe our new results related to the situation when decisions can be deferred, so that the acceptance decisions for the orders within a “batch” are made only when all the orders within that batch have been received. When the batches are of equal size, we will provide an algorithm that achieves a regret that scales as O(log K) in the number of batches K, uniformly in n, and that matches the regret lower bound. In addition, we introduce the concept of a "delay budget" and

discuss how to select the optimized batch size when delay in order fulfillment is important. This work is joint with Haoran Xu and Yinyu Ye.

Speaker : Ashish Goel

Affiliation : Stanford University

Title : Decentralized Finance: Theory, Practice, and Architectures

Abstract : The emergence of blockchains and stablecoins has resulted in new kinds of financial markets and applications, commonly labeled DeFi (decentralized finance). In the first part of this talk, we will describe recent theoretical work from our group on algorithms for pricing, matching, and clearing in these new markets. We will show how multi-currency exchanges can be efficiently cleared using simple convex programs, and how to design optimum pricing curves for automated market makers. The second part of this talk will focus on regulatory risks associated with stablecoins and blockchains. We will present an architecture for DeFi that we believe is optimized for the Indian payments ecosystem, satisfies regulatory concerns, and still enables innovative DeFi applications.

Speaker : Sumit M Kunnumkal

Affiliation : ISB

Title : Strawberry Or Vanilla This Week? How To Optimize Tailored Assortments For Variety-Seeking/Avoiding Consumers

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

Abstract: We consider the problem of a retailer personalizing an assortment to a consumer who is variety-seeking or variety-avoiding, that is, less or more likely to buy the same product as in the previous purchase occasion. To capture the consumer's attitude towards variety, we construct a model based on that of Givon (1984) and later modified by Feinberg et al. (1994). The problem can be formulated as a dynamic program. Although it is intractable in general, we identify special cases where the optimal assortment can be determined efficiently. Moreover, we characterize the structure of the optimal assortment and compare it to the case where the consumer is variety neutral and makes choices according to the multinomial logit model. Our computational experiments indicate that the retailer can lose substantial revenues if the consumer's attitude towards variety is not considered while making the assortment decision. This is joint work with Dorothee Hohon (UTD), Ismail Kirci (UIUC) and Sridhar Seshadri (UIUC).

Speaker : Aditya Gopalan

Affiliation : Indian Institute of Science

Title : Multi-armed Bandits & Reinforcement Learning

Abstract : These lectures will discuss models and algorithms for sequential decision-making in the face of uncertainty, often when there is an information constraint leading to an exploration-exploitation tradeoff for the learning agent. The quintessential example of such a problem is the multi-armed bandit, which is a special case of more general reinforcement learning. Sequential learning problems arise in the realm of many modern data-driven intelligent systems (think Internet recommendation engines, hyperparameter tuning for machine learning algorithms, game playing, financial portfolio allocation, resource allocation in communication systems, etc.), in which one desires the ability to continuously adapt to changing conditions.

The tentative plan for the lectures is as follows:

1. Introduction to multi-armed bandit problems, the stochastic k-armed bandit model, performance metrics: regret, best-arm identification
2. Algorithms: Explore-Then-Commit (Epsilon-Greedy), UCB, Thompson Sampling, Median elimination, Successive Rejects
3. Lower bounds: Lai-Robbins lower bound on asymptotic regret, minimax lower bounds
4. More general Reinforcement learning

Safexpress Centre for Data, Learning and Decision Sciences

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

Speaker : Karthyek Murthy

Affiliation : USC Los Angeles

Title : The Scaling Behaviors in Achieving High Reliability via Chance-Constrained Optimization

Abstract : We consider the problem of resource provisioning under stringent reliability or service-level requirements, which arise in applications such as power distribution, emergency response, cloud server allocation, and regulatory risk management. With chance-constrained optimization serving as a natural starting point for modeling this class of problems, our main contribution is to characterize how the optimal costs and decisions scale for a generic joint chance-constrained model as the target probability of satisfying the service/reliability constraints is raised to one. Besides revealing the behavior of optimal solutions in the high reliability regime, our scaling result has two key consequences for learning to optimize under high-reliability constraints:

1. Distributionally robust optimization (DRO) modeling of chance constraints. In DRO modeling of chance constraints, widely used ambiguity sets based on KL divergences, Wasserstein distances, or moment constraints substantially distort the scaling properties of optimal decisions, often inflating costs exponentially. In contrast, DRO formulations of the P-model variant are less sensitive to the ambiguity set choice and offer decision-makers a greater control over the robustness–conservativeness tradeoff.
2. Data-driven estimation. Given N data samples, the scaling framework enables estimation of nearly Pareto-optimal decisions whose reliability levels are significantly larger than the $1-\Omega(1/N)$ -estimation barrier a decision-maker encounters in the absence of parametric distribution assumptions.

Speaker : Dheeraj Nagaraj

Affiliation : Google DeepMind

Title : Diffusion and Flow models for Generative Modeling

---

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

Speaker : Venkat Padmanabhan

Affiliation : Microsoft Research

Title : From FLOPS to Hits: An Overview of (AI) Explorations at Microsoft Research India

Abstract : In this talk, I will give an overview of work at Microsoft Research India, which completed two decades in 2025. Given the times we live in, much of our current work unsurprisingly related to AI but nevertheless spans a wide range. I will walk up the "AI stack", going from the systems infrastructure to effectively utilize all the FLOPS on offer from the hardware to the intelligence infrastructure comprising building blocks such as retrieval that are key to AI models being effective to applications in such diverse areas such as education, health care, and road safety that have arguably been "hits" with users. I will also touch on work on cross-cutting challenges such as security and responsible AI, and also that on the "science of AI", which is not to be confused with "AI for science"! Woven through my talk will be thoughts on and learnings from how, as a corporate research lab, we have long partnered with the ecosystem --- academia, government, startups, and NGOs --- in myriad but mutually beneficial ways.

Speaker : Alexandre Proutiere

Affiliation : KTH Royal Institute of Technology

Title : Emergence of low-rank structures in Reinforcement Learning (Joint work with Bastien Dubail and Stefan Stojanovic)

Abstract : Low-rank structure is a common implicit assumption in many modern reinforcement learning (RL) algorithms. For instance, reward-free and goal-conditioned RL methods often presume that the successor measure admits a low-rank representation. In this work, we challenge this assumption by first remarking that the successor measure itself is not approximately low-rank. Instead, we demonstrate that a low-rank structure naturally emerges in the shifted successor measure, which captures the system dynamics after bypassing a few initial transitions. We provide finite-sample performance guarantees for the entry-wise estimation of a low-rank approximation of the shifted successor measure from sampled entries. Our analysis reveals that both the approximation and estimation errors are primarily governed by a newly introduced quantitity: the spectral recoverability of the corresponding matrix. To bound this parameter, we derive a new class of functional inequalities for Markov chains that we call Type II Poincaré inequalities and from which we can quantify the amount of shift needed for effective low-rank approximation and estimation. This analysis shows in particular that the required shift depends on decay of the high-order singular values of the shifted successor measure and is

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

hence typically small in practice. Additionally, we establish a connection between the necessary shift and the local mixing properties of the underlying dynamical system, which provides a natural way of selecting the shift. Finally, we validate our theoretical findings with experiments, and demonstrate that shifting the successor measure indeed leads to improved performance in goal-conditioned RL.

Speaker : Arun Sai Suggala

Affiliation : Google DeepMind

Title : Understanding State-of-the-Art LLMs: From Architecture to Reasoning

Abstract : Large Language Models (LLMs) have recently demonstrated remarkable capabilities, achieving near-human performance in domains ranging from complex mathematics at the International Mathematical Olympiad level to accelerating scientific discovery. This tutorial provides an overview of some of the key concepts powering these advances. We will examine the various stages of the modern LLM development stack, detailing architectural innovations in pre-training and crucial post-training techniques, such as RLHF and reasoning. The tutorial will conclude by outlining open research questions and some exciting new directions.

Speaker : Bismark Singh

Affiliation : University of Southampton

Title : Balancing user accessibility and facility load fairness in undesirable facility location problems

Abstract : Typically, within facility location problems, fairness is defined in terms of accessibility of users. However, for facilities perceived as undesirable by communities hosting them, fairness between the usage of facilities becomes especially important. Limited research exists on this notion of fairness. In a series of works, we have attempted to close this gap by developing new classes of combinatorial optimization models for the allocation of populations of users to facilities such that access for users is balanced with a fair utilization of facilities. The optimality conditions of the underlying nonconvex quadratic models state the precise balance between accessibility and fairness. Further, we define new axioms of fairness and a metric to quantify the extent to which fairness is achieved in both optimal and suboptimal allocations. We

show that a continuous relaxation of our central model is sufficient to achieve a perfect extent of fairness, while a special case reduces to the classical notion of proportional fairness. We present computational results using actual data from the state of Bavaria in Germany. The main part of this work is based on two articles published with my students both in the INFORMS Journal on Computing. In ongoing work, we are studying supermodularity properties of this new class of objective functions.

Speaker : Dootika Vats

Affiliation : IIT Kanpur

Title : MCMC Importance Sampling via Moreau-Yosida Envelopes

Abstract : Markov chain Monte Carlo (MCMC) is the workhorse computational algorithm employed for inference in Bayesian statistics. Gradient-based MCMC algorithms are known to yield faster converging Markov chains. In modern parsimonious models, the use of non-differentiable priors is fairly standard, yielding non-differentiable posteriors. Without differentiability, gradient-based MCMC algorithms cannot be employed effectively. Recently proposed proximal MCMC approaches, however, can partially remedy this limitation. These approaches employ the Moreau-Yosida (MY) envelope to smooth the nondifferentiable prior enabling sampling from an approximation to the target posterior. In this work, we leverage properties of the MY envelope to construct an importance sampling paradigm to correct for this approximation error. We establish asymptotic normality of the importance sampling estimators with an explicit expression for the asymptotic variance which we use to derive a practical metric of sampling efficiency. Numerical studies show that the proposed scheme can yield lower variance estimators compared to existing proximal MCMC alternatives.

Speaker : Nisheeth Vishnoi

Affiliation : Yale University

Title : Creativity, Discovery, and Learning in the Age of AI

Abstract : When you learn or create something new, much of what matters actually happens beneath the surface. Real understanding is more than just a mental exercise; it is a lived experience. It grows through the physical and emotional work of trying, failing, and repeating. It

exSPLOre 26: Schedule

Explorations in Statistics, Probability, Learning and Optimization Research 2026

January 13-17, 2026

is the slow process of building a gut feeling for what works. This is how you develop a coherent understanding of the world and the confidence to judge your own ideas.

Modern AI tools change this experience. Because these systems can produce fluent, confident answers instantly, they shift your attention toward the finished result and away from the messier process of learning. This contrast makes a vital distinction clearer than ever before: producing an answer is not the same as understanding it, and getting an output is not the same as knowing why it matters.

This talk uses AI as a mirror to reflect on how we learn. We will explore how your intuition and judgment grow out of your own experiences, why the hard work of checking and verifying ideas is essential, and how a growing reliance on AI might reshape your independence as a thinker. In a world where machines can generate almost anything, we will ask: what does it truly mean to discover, to create, and to know for yourself?

Speaker : Assaf Zeevi

Affiliation : Columbia Business School

Title : Thompson Sampling through an Optimization Lens

Abstract : Multi-armed bandits are widely studied abstractions of sequential decision making problems that allow, among other things, a distilled focus on the so-called exploration-exploitation tradeoff in online learning. Various families of learning algorithms have been developed over the years for this class of problems, which continues to be the subject of very active research in several academic communities. In addition, many of said algorithms are now deployed on scale at various large technology companies. Perhaps the most pervasive implementations are variations on a Bayesian posterior sampling approach known as Thompson Sampling. The core idea underlying this algorithm is exceedingly simple, intuitive, and surprisingly effective, but despite Thompson's paper soon celebrating its centennial, the reasons underlying its remarkable success remain somewhat shrouded in mystery. In this talk we will focus on a simple question: what exactly is Thompson Sampling optimizing?