

Title and Abstracts

Dimitris Bertsimas (MIT)

Algorithms for Life: AI and Operations Research Transforming Healthcare, Education, and Agriculture.

For over four decades, I have pursued the vision of using mathematics and algorithms to address problems that improve the human condition—saving lives, improving learning, and feeding the world. This talk traces the convergence from operations research and optimization under uncertainty, through machine learning and AI, to their impact in healthcare, education, and agriculture. I share how breakthrough research becomes real-world impact, with concrete examples from medical diagnosis and precision medicine to precision agriculture and personalized education. And I explore what is next: how emerging AI developments will fundamentally reshape these fields in ways we are only beginning to imagine.

Rakesh Vohra (University of Pennsylvania)

Arrow's Theorem and Linear Programming

The recent interest in AI alignment has resurrected interest in the old questions about preference aggregation inspired by Arrow's celebrated (im)possibility theorem. About a quarter of a century ago, CP, Jay Sethuraman and myself took Arrow's axioms and turned them directly into a 0-1 integer program which satisfied total unimodularity, hence a linear program. Not only did this yield a new proof of Arrow's theorem, but new and unified proofs of a whole host of results about preference aggregation.

Jay Sethuraman (Columbia University)

Core Allocations in a Stable Sharing Model

I'll describe a model in which agents are nodes on a graph and two agents are potential partners if they are connected by an edge in the graph. Agents have to be matched in pairs, and each pair must complete a task that requires one unit of effort. Each agent has symmetric preferences around an ideal contribution level. I'll discuss strong and weak core allocations in this simple model, drawing connections with natural matching and covering problems on the graph. (Joint work with Sonal Yadav)

Zuo-Jun (Max) Shen

LLMs in Operations Management: Exploration and Evidence

Large language models (LLMs) have demonstrated strong capabilities across a wide range of tasks, creating new opportunities for operations management. However, making LLMs genuinely useful for operational decision-making requires more than simply applying general-purpose models to existing problems. Their outputs may be inaccurate or infeasible, while repeated reasoning and verification can introduce substantial computational and latency costs. This talk explores how LLMs can be used effectively and efficiently in operations management through two research projects. The first studies LLM-based optimization modeling, and the second studies scheduling and online operational decisions. Experimental results show that carefully designing the role, placement, and inference process of LLMs can substantially improve their effectiveness. Together, we highlight a two-way relationship between LLMs and operations management: LLMs can help solve operational problems, while OM principles can help design more effective ways to use LLMs.

Melvyn Sim (NUS)

Streamlining Robustness for Sequential Decision Making

Sequential managerial decisions require early commitments under incomplete information and adaptation as uncertainty unfolds. We develop an end-to-end framework connecting uncertainty learning, robust satisficing, and sequential optimization. The framework uses tolerated robust constraints, a managerially interpretable variant of globalized robust constraints, to control where a cost target must be maintained, how far protection extends, and how rapidly cost may deteriorate. We construct tractable formulations using period-wise segregated affine recourse. Under unrestricted support, convexifying the lifted uncertainty set introduces no additional conservatism; with an asymmetric L_1 deviation gauge and zero strict robustness level, the formulation is lossless relative to fully adaptive, non-anticipative recourse. We also derive an exact single-period reformulation for general strict robustness levels. The deviation-gauge parameters are learned from historical prediction errors through minimum-volume estimation, which admits a closed-form solution under componentwise normalization. A multiperiod ambulance relocation study, using only historical demand observations, demonstrates strong out-of-sample performance and favorable computational scaling relative to a sample-average stochastic optimization benchmark.

Guodong Lyu (HKUST)

Service-Level-Constrained Capacity Planning: A Little Flexibility Goes a Long Way

Enterprise planning systems often translate aggregate service-level targets into static safety-stock or capacity rules, making it difficult to adapt the underlying decision logic to dynamic operating environments. Motivated by the Gaze Heuristic, we develop a service-level-constrained approach to joint capacity planning and resource allocation under uncertainty. Across three settings, we show that limited but well-designed flexibility can deliver substantial operational value. In flexible production networks with stationary stochastic demand, sparse long-chain flexibility achieves fill-rate performance close to that of full flexibility while requiring much less capacity than dedicated systems. In crowdsourced delivery with nonstationary demand, modest volume flexibility in delivery walker deployment helps maintain service targets without continual capacity reconfiguration. In high-speed railway ticket allocation with decision-dependent demand, selectively offering covering-ticket alternatives can recover otherwise lost sales while preserving origin-destination service requirements. By contrast, excessive flexibility may inefficiently consume scarce network capacity. Together, these results show how service guarantees, capacity configuration, resource allocation, and flexibility design can be integrated to build planning systems that are transparent, adaptive, and operationally effective.

Zhenzhen Yan (NTU)

Learning Consumer Choice under Framing: Joint Assortment and Pricing Experiments with Representative Consumer Models

Product presentation influences demand through both display position and assortment size. Building on data-driven representative-consumer pricing, we develop a separable representative consumer model (SRCM) that incorporates these framing effects while retaining nonparametric univariate perturbations. We identify the model through a two-stage design of pricing and ordered-offer experiments. With exact market shares, the first stage recovers effective perturbation derivatives on attainable intervals up to a common shift. The second stage separates normalized position and size effects from structural derivatives using as many ordered offers as there are products. Under reachability and uniform price-domain and curvature conditions, the required number of queries grows polynomially with product count and inverse coverage precision. The per-target query count grows logarithmically with inverse target accuracy. Finite-sample guarantees also hold when market size and no-purchases are observed and experiments use fresh conditionally independent batches. We further show that framing can eliminate every positive constant approximation guarantee for profit-nested assortments. For the multinomial logit specialization with a submodular transformed size effect, we obtain a convex relaxation for fixed-price assortment and display-order optimization. Synthetic experiments demonstrate framing-parameter recovery, more accurate estimates than MNL-based maximum likelihood under Uniform utility shocks, and accurate share prediction and profit calibration under correlated shocks. In an airline application, the full model achieves the lowest or near-lowest held-out prediction error among five specifications, with ablations supporting the predictive value of both framing and perturbation flexibility.

Huan Zheng (Shanghai Jiao Tong University)

On Fairness of Process Flexibility: a Max-min Fill Rate Perspective

While sparse process flexibility is known to match the efficiency of full flexibility, its performance regarding fairness—a critical objective in humanitarian and healthcare operations—remains under-explored. This study investigates the worst-case fairness of sparse structures using the max-min demand fill rate. We demonstrate that this metric is determined by the demand subset with the minimum ratio of neighboring capacity to demand, and we characterize the performance bounds for k -chain and Ψ -expander structures. Validated with real-world vaccine distribution data from Texas, the results conclude that integrating flexibility design can substantially balance resource allocation and enhance public health protection.

Hanzhang Qin (NUS)

Large Language Models (LLMs) for Automated Operations Modeling

This talk presents my recent work at the interface of large language models (LLMs) and operations research, especially regarding automated optimization formulations in collaboration with Prof. CP Teo.

Karthik Natarajan (SUTD)

Moments with CP: A Theorem, a Counterexample, and an Open Question

In this talk I will review some older and some newer results on moment problems. Some of the results use Completely Positive cones and some of the results are joint work with Chung Piaw and others.

Zheng Zhichao (SMU)

Robust Data Pooling under Data Poverty in Sequential Decision-Making

Data poverty poses a fundamental challenge for sequential decision-making: some groups may have too few observations to support reliable policy learning on their own. Data pooling offers a natural way to alleviate this scarcity by borrowing information from better-sampled groups. However, when data are unevenly distributed and the underlying rewards or transition dynamics differ across groups, pooled models can be disproportionately influenced by the better-sampled group and introduce systematic bias for the data-poor group. We study this trade-off in finite-horizon Markov decision processes. We first derive finite-sample regret bounds that separate the statistical benefit of pooling from the bias induced by cross-group heterogeneity, and characterize when empirical pooling provides a tighter regret guarantee for the data-poor group than learning from its own data alone. We then develop a robust data-pooling approach that mitigates the bias introduced by pooling while retaining its statistical benefits. Our analysis shows that robust data pooling can expand the range of settings in which borrowing data from the better-sampled group yields a tighter regret guarantee than separate learning, including cases where empirical pooling alone does not. Numerical experiments on an ICU extubation task corroborate our theoretical findings, showing that robust pooling improves overall performance while reducing performance disparities across groups.

Qingxia Kong (Erasmus University Rotterdam)

Costly Gambler's Fallacy

How do erroneous beliefs about random sequences - the Gambler's Fallacy- respond to economic incentives? We study this question in two online experiments in which 717 participants make 40 incentivized choices over independently drawn binary outcomes. We create identical expected-value differences between the two options through two distinct mechanisms: varying outcome probabilities while holding payoffs constant, and varying payoffs while holding outcome probabilities at 50-50. While participants respond to the expected-value advantage by choosing the favorable option more frequently, a minority consistently implement the payoff-maximizing strategy. Individual-level analysis reveals that incentives primarily shift participants from mixed choice toward near-maximizing behavior; they have little effect on the prevalence of pronounced gambler's-fallacy-like or hot-hand-like behavior. Approximately 18%-23% of participants exhibit strong negative recency across treatments, even when responding to recent outcomes is costly. These choice patterns extend to deeper outcome histories and predict the direction of updating in independently elicited probabilistic beliefs. Incentivized recommendations further indicate that many participants actively search for patterns in explicitly random sequences. Our findings show that economic incentives improve decision making on average but do not eliminate systematic misperceptions of randomness among a persistent minority.