Haskayne School of Business | University of Calgary

Land Acknowledgment

The University of Calgary, located in the heart of Southern Alberta, both acknowledges and pays tribute to the traditional territories of the peoples of Treaty 7, which include the Blackfoot Confederacy (comprised of the Siksika, the Piikani, and the Kainai First Nations), the Tsuut'ina First Nation, and the Stoney Nakoda (including Chiniki, Bearspaw, and Goodstoney First Nations). The city of Calgary is also home to the Métis Nation within Alberta (including Nose Hill Métis District 5 and Elbow Métis District 6).

The townsite of Banff is also located on traditional Treaty 7 territory. These sacred lands are a gathering place for the Niitsitapi from the Blackfoot Confederacy, of whom the Siksika, Kainai, and Piikani First Nations are part; the Îyârhe Nakoda of the Chiniki, Bearspaw, and Goodstoney First Nations; the Tsuut'ina First Nation; and the Métis Nation of Alberta.

Welcome Message

Dear Colleagues,

It is our great pleasure to welcome you to the second workshop on Artificial Intelligence and Machine Learning in Operations Management Research (AIMOR), taking place May 14–15 at the Royal Canadian Lodge, in the heart of the breathtaking Banff National Park.

This gathering brings together a vibrant community of researchers dedicated to advancing the role of AI/ML in the evolving field of Operations Management. As organizations face increasing complexity and uncertainty, the insights shared here will help shape the future of how we analyze, optimize, and lead operational systems through intelligent technologies.

We would like to extend a special thank-you to the Haskayne School of Business for their generous support of this event. Their substantial funding has enabled us to offer significantly discounted registration fees, making it more accessible for a broad and diverse group of participants. Their commitment to fostering innovation and academic exchange has been instrumental in bringing this workshop to life.

Over the next two days, we invite you to engage fully with a dynamic program of keynote presentations, research sessions, and collaborative discussions. We also hope you'll take the time to enjoy the natural beauty that surrounds us here in Banff—an inspiring setting for both fresh thinking and bold ideas.

Thank you for being part of this workshop. We are thrilled to have you with us and look forward to a memorable and thought-provoking workshop experience.

Warm regards,
Osman Alp, Marco Bijvank, Yiwen Jin and Hossein Piri

Workshop Schedule

May 14, 2026

Start	End	Topic
8:00 am	8:50 am	Registration + Breakfast
8:50 am	9:00 am	Welcome Message
9:00 am	10:30 am	Keynote: *Reinforcement Learning as a Sequential Decision Problem using the Universal Modeling Framework*, Warren Powell (Princeton University) – Part 1
10:30 am	11:00 am	Break
11:00 am	12:30 pm	Keynote: *Reinforcement Learning as a Sequential Decision Problem using the Universal Modeling Framework*, Warren Powell (Princeton University) – Part 2
12:30 pm	1:30 pm	Lunch (provided)
1:30 pm	3:00 pm	Technical Session 1
3:00 pm	4:00 pm	Break + Poster Session 1
4:00 pm	5:30 pm	Technical Session 2
6:30 pm	8:00 pm	Dinner (provided)

May 15, 2026

Start	End	Topic
8:00 am	9:00 am	Breakfast
9:00 am	10:30 am	Keynote: *Learning in Stochastic Models: An Asymptotic Approach*, Amy Ward (University of Chicago Booth School of Business) – Part 1
10:30 am	11:00 am	Break
11:00 am	12:30 pm	Keynote: *Learning in Stochastic Models: An Asymptotic Approach*, Amy Ward (University of Chicago Booth School of Business) – Part 2
12:30 pm	1:30 pm	Lunch (provided)
1:30 pm	3:00 pm	Technical Session 3
3:00 pm	4:00 pm	Break + Poster Session 2
4:00 pm	5:30 pm	Technical Session 4

Keynote Speakers

Warren B. Powell (Princeton University)

Reinforcement Learning as a Sequential Decision Problem using the Universal Modeling Framework

Reinforcement learning emerged in the 1980s in computer science for approximating Bellman's equation, building on the framework of Markov decision processes introduced in the 1950s by Bellman. This work paralleled a line of research from the optimal control community starting in 1974 for deterministic control problems, and then the work on approximate dynamic programming in operations research in the 1990s. Ultimately, all three lines of research came to the same conclusion: approximating Bellman's equation is really hard and generally does not work.

The lines of research emerging from these communities (computer science using "reinforcement learning", optimal control using "adaptive dynamic programming," and operations research using "approximate dynamic programming") would all grow to embrace a wide range of solution methods to meet the needs of the vast range of problems being addressed. Over a series of articles from 2014 to 2019, I formalized the idea that all RL problems are sequential decision problems: decision, information, decision, information, …, which can be modeled using what I call the "universal modeling framework," a derivative of what is used in optimal control.

I then realized that every method for making decisions could be organized into four classes of policies, only one of which uses Bellman's equation. Policies in each of the four classes can be found throughout the optimal control literature, as well as in the 2018 edition of Sutton and Barto's book Reinforcement Learning. This laid the foundation for my 2022 volume, Reinforcement Learning and Stochastic Optimization: A unified framework for sequential decisions which is written entirely around the universal modeling framework and all four classes of policies.

I will demonstrate how this framework identifies approaches for any sequential decision problem, spanning the traditional problems with discrete actions up to the high-dimensional problems faced in operations research. All four classes of policies can be applied to stochastic search problems known as multi-armed bandit problems (I like the name "intelligent trial-and-error"), which are arguably the most common decision problem in practice.

Amy R. Ward (The University of Chicago Booth School of Business)
Learning in Stochastic Models: An Asymptotic Approach

We consider queueing models with reneging in which the primitive distributions are unknown and must be learned. The question of interest is how to prioritize arriving individuals for service. When the reneging distribution is non-exponential, the state space is very complex, because we must track the amount of time each customer in queue has been waiting in order to have a Markovian system. As a result, learning an exact optimal policy is very challenging.

The approach we take is to first ignore the discrete and stochastic nature of arrivals, and to instead assume that the number of individuals to arrive in a time interval of length $t>0$ is proportional to t. Then, we use historical (offline) data to formulate a data-driven (fluid) optimization problem (that may incorporate machine learning predictions) whose objective is to maximize long-run average reward. The solution to the data-driven optimization problem motivates a data-driven policy for deciding how to prioritize arriving individuals for service. With an infinite amount of historical data available, the data-driven policy is asymptotically optimal as the arrival rate and service capacity become large. With a finite amount of historical data available, the challenge is to establish regret bounds and/ or finite sample statistical guarantees.

Technical Session 1

May 14 1:30 pm – 3:00 pm

Solving Assortment Optimization with First-Order Methods and Neural Networks: A Computational Framework and Public Benchmark

Ningyuan Chen (University of Toronto) ,Qing Guo (Renmin University), Saman Lagzi (UNC Chapel Hill), Chenhao Wang (Tongji University), Guillermo Gallego (CUHK-SZ), Sumit Kunnumkal (Indian School of Business), Yao Wang (Xi'an Jiaotong University), Li Yu (Renmin University)

Assortment optimization is a central challenge in revenue management, due to non-linear objective function and discrete decision variables. Meanwhile, first-order methods like gradient descent have seen widespread adoption for continuous optimization in large-scale AI systems. We propose a computational framework that combines first-order methods and neural networks to efficiently solve assortment optimization. Our framework features straight-through estimators, which enable gradients to flow through discrete variables, and utilizes neural networks to perturb the gradient updates. We theoretically ground our framework by proving that our method is guaranteed to converge to the globally optimal solution for the unconstrained problem under the Multinomial Logit model (MNL). Furthermore, recognizing the need for standardized evaluation in this domain, we develop and release a public benchmark dataset, available at https://github.com/wch444/Assortment-Benchmark. This dataset, comprising several challenging assortment optimization problems, serves both to empirically test our proposed framework and to provide a robust testbed for the research community.

Learning for Assortment Decisions: A Decision-Focused Approach under MNL Model

Yu Sun (University of Toronto), Ming Hu (University of Toronto), Hansheng Jiang (University of Toronto), Sheng Liu (University of Toronto), Daniel Zhuoyu Long (The Chinese University of Hong Kong), Ruxian Wang (Johns Hopkins University)

We study assortment optimization under multinomial logit (MNL) choice models from a decision-focused perspective. Classical approaches estimate model parameters by maximizing likelihood and then optimize assortments under the estimated model, which can lead to poor decisions when estimation errors are amplified by the optimization step. We propose a decision-focused learning framework that directly trains choice models to optimize downstream assortment performance. By exploiting the revenue-ordered structure of the MNL model, we derive a tractable reformulation of the resulting problem. Using real data, we compare our approach with both the classical MNL plug-in method and distributionally robust optimization (DRO). Our results reveal a clear tradeoff: decision-focused learning performs strongly when the MNL model is correctly specified or close to the ground truth, while DRO provides superior robustness under greater model misspecification.

Return Concerned Adaptive Control of Product Offerings

Recep Bekci (University of Waterloo), Mehmet Gumus (McGill University), Sentao Miao (University of Colorado Boulder)

In collaboration with a major European fast-fashion retailer, we address the challenge of optimizing product assortments under high return rates and delayed feedback. We formulate a dynamic optimization problem where customer demand and return probabilities are initially unknown. To solve this, we introduce the COBRA (COnfidence Bounds for Return-aware Assortment) algorithm, which employs an optimistic learning strategy to balance exploration and exploitation. We prove that COBRA achieves a regret bound of $\tilde{O}(\sqrt{NT})$, establishing its near-optimality against a matching lower bound. Extensive experiments using synthetic simulations and real-world data demonstrate that COBRA significantly outperforms standard benchmarks. By effectively integrating censored return data, our approach not only enhances profitability but also substantially reduces return volumes, offering a scalable solution for sustainable retail operations.

Reinforcement Learning for Cost Estimation in Two-Level Combinatorial Decision Problems

X.Y. Han (Chicago Booth), Yuan Zhong (Chicago Booth)

In large-scale AI training, Sparse Mixture-of-Experts (s-MoE) layers enable scaling by activating only a small subset of experts per token. An operational challenge in this design is token-to-expert load-balancing, which is important for the efficient utilization of (costly) GPUs. We provide a theoretical framework for analyzing the Auxiliary-Loss-Free Load Balancing (ALF-LB) procedure -- proposed by DeepSeek's Wang et al. (2024) -- by casting it as a computationally-efficient primal-dual method. First, in a stylized deterministic setting, our framework yields several insightful structural properties like monotonic improvement of certain metrics, a preference rule, and an approximate-balancing guarantee. Then, we incorporate the stochastic nature of AI training using online optimization. In the online setting, we derive a strong convexity property that leads to a logarithmic expected regret bound under certain step-sizes. Additionally, we present experiments on 1B-parameter DeepSeekMoE models. These results build a principled framework for analyzing the ALF-LB procedure in AI training.

Poster Session 1

May 14 3:00 pm – 4:00 pm

A Theoretical Framework for Auxiliary-Loss-Free Load Balancing of Sparse Mixture-of-Experts in Large-Scale AI Models

Fausto Errico (École de technologie supérieure - CIRRELT and GERAD), Mohsen Dastpak (École de technologie supérieure - CIRRELT and GERAD), Ola Jabali (Politecnico di Milano)

Many combinatorial decision problems exhibit a two-level decision structure, where high-level strategic choices influence lower-level operational decisions. Examples include Facility Location, Location Routing, and Outsourcing problems, where second-level decisions determine costs of first-level choices. A challenge in solving such problems is accurately and rapidly estimating second-level costs to guide first-level decision-making. In this presentation, we propose leveraging Reinforcement Learning (RL) to obtain fast and accurate approximations of second-level costs, which can then be integrated into traditional (meta)-heuristic frameworks for exploring first-level decisions. We demonstrate this approach on a daily distribution problem, where a fleet must serve a set of customers whose demands are stochastic. Each day, the set of customers and their demand distributions are revealed. The decision-maker can outsource a subset of customers at a given cost while committing to serve the remaining ones. The objective is to determine the optimal outsourcing decision and an efficient routing policy.

Climate Risk Exposure and Supply Chain Restructuring: A Large Language Model Approach

Jiong Sun (Purdue University), Leyao Tan (University of British Columbia), Yi Qian (University of British Columbia)

This study examines how to quantify firm-level climate risks and how supplier-level climate risks prompt firms to restructure their supply networks. We develop a novel approach to measuring climate risk by combining Large Language Model sentence embeddings with supervised machine learning. Our approach improves classification precision by 20% compared to dictionary-based approaches. Using this climate risk measure, we find that supplier climate risk exposure increases the likelihood of terminating existing suppliers by 5.3% and adding new suppliers by 6.3%. Our analysis further suggests that restructured supply networks exhibit greater geographic diversity, indicating firms' efforts to mitigate future climate-related disruptions by diversifying supplier locations. In addition, firms prioritize adding new suppliers that produce more unique products or occupy more central positions in the supply chain network. Finally, we find that firms with higher operational efficiency and greater product differentiation respond more aggressively to supplier climate risk.

Methodological Advances in Human-LLM Decision-Making Research: Interactive AI Experimental Design Protocol and Conversational Analysis

Tracy Jenkin (Queen's University), Yang Chen (Western University), Samuel N. Kirshner (University of New South Wales,), Anton Ovchinnikov (Queen's University, INSEAD), Meena Andiappan (McMaster University)

With their rapid ascension, large language models (LLMs) are increasingly used as advisors. The way individuals interact with LLMs differs from traditional predictive AI, representing a paradigm shift in human-AI interaction, where interactions unfold dynamically through multi-turn natural language exchanges. Thus, to study human-LLM decision making, we need experimental designs that capture how individuals interact with LLMs in real time as well as approaches to analyze these conversations. Accordingly, we develop a novel interactive AI experimental design protocol, embedded into a Qualtrics survey, which captures study participants' free-form conversations with an LLM, going beyond the typical experimental interfaces in which participants cannot seek information not provided by the experimenters. We then use GPT to assist in the qualitative inductive coding of the resulting conversation data, supplementing the GPT coding with human review and oversight. To demonstrate effectiveness, we apply these approaches to study decision biases, comparing static to interactive conditions.

Generative Artificial Intelligence for Decision Making in Supply Chain

Mohammad Moshref-Javadi (University of Illinois at Urbana Champaign), Amlan Baruah (University of Illinois)

This paper presents a structured framework for uses of Generative Artificial Intelligence in Supply Chain Management. It examines how different GenAI models—including Generative Adversarial Networks (GANs), Transformers, Variational Autoencoders (VAEs), and flow-based models—are applied across six levels of supply chain processes. Using bibliometric analysis, density analysis, temporal analysis, and topic modeling based on more than 692 articles, the study summarizes key practical applications, major focuses of research in the literature, identifies major challenges, and outlines directions for future research.

Technical Session 2

May 14 4:00 pm – 5:30 pm

Balancing Optimism and Pessimism in Offline-to-Online Learning

Ilbin Lee (University of Alberta), Flore Sentenac (HEC Paris), Csaba Szepesvari (University of Alberta, Deepmind London)

We consider the offline-to-online learning setting, focusing on stochastic finite-armed bandit problems. In offline-to-online learning, a learner starts with offline data passively collected from interactions with an unknown environment, and then begins interacting with the environment to maximize total reward. The learner faces a fundamental dilemma: if the horizon is short, a suitable strategy (in a number of senses) is the Lower Confidence Bound (LCB) algorithm, based on pessimism. LCB effectively competes with any policy "covered" by the offline data. However, for longer time horizons, a preferred strategy is the Upper Confidence Bound (UCB) algorithm, based on optimism. We develop a new algorithm that gradually transitions from LCB to UCB as more rounds are played. Both our theoretical and empirical results show that our new algorithm performs nearly as well as the better of LCB and UCB at any point in time.

Finite-Sample Analysis of Decentralized Q-Learning for Stochastic Games

Zuguang Gao (University of California, Irvine), Qianqian Ma (Coupang, Inc), Tamer Başar, University of Illinois Urbana-Champaign), John R. Birge (The University of Chicago Booth School of Business)

Learning in stochastic games is arguably the most standard and fundamental setting in multi-agent reinforcement learning. We establish the finite-sample complexity of fully decentralized Q-learning algorithms in a significant class of general-sum stochastic games – weakly acyclic games, which includes all potential games. We focus on the practical while challenging setting of fully decentralized MARL, where neither the rewards nor the actions of other agents can be observed by each agent. Both tabular and linear function approximation cases are considered. In the tabular setting, we analyze the sample complexity for the decentralized Q-learning algorithm to converge to a Markov perfect equilibrium. With linear function approximation, the results are for convergence to a linear approximated equilibrium, a new notion of equilibrium that we propose, which describes that each agent's policy is a best reply (to other agents) within a linear space. Numerical studies are also provided to demonstrate the results.

Reinforcement Learning for Reasoning in Time Series Analysis

Kimia Ghobadi (Johns Hopkins University), Felix Parker (Johns Hopkins University)

Many operations decisions rely on time-series forecasts informed by unstructured context, yet conventional pipelines separate perception from prescription and offer little auditability. We introduce a two-stage framework that equips a pretrained language model to reason over numerical sequences. A residual vector-quantized variational autoencoder tokenizes time-series patches into a discrete vocabulary that interleaves with text, enabling mixed-modality prompts. Supervised tuning teaches parsing; reinforcement learning with verifiable rewards elicits compact chain-of-thought before answers. Across electrocardiogram question answering, context-aware forecasting, and few-shot UCR classification, the model improves accuracy over baselines while producing machine-readable rationales suited to downstream optimization and healthcare operations.

Reinforcement Learning for Optimizing Physician Effort Allocation in Emergency Departments

Parvin Malekzadeh (University of Toronto), Opher Baron (University of Toronto), Simai He, (Shanghai Jiao Tong University), Dmitry Krass (University of Toronto), Hongsong Yuan, (Shanghai University of Finance and Economics)

In emergency departments (EDs), physicians dynamically split effort between new patients awaiting initial assessment and in-system patients needing reassessment or discharge. Physicians often prioritize throughput, while administrators seek to reduce waiting—creating competing objectives. We model effort allocation as a finite-horizon control problem in a two-station queueing system with stochastic arrivals, service times, and abandonment. We develop a reinforcement learning (RL) framework that leverages network structure to separate deterministic and stochastic state components, enabling augmented sample generation and improving sample efficiency. Experiments show the method outperforms RL baselines and produces near-optimal policies for physician-, ED-, and balanced (Pareto-efficient) objectives. The throughput-maximizing policy has a time-threshold structure, while the waiting-time–minimizing policy follows a reversed-priority (Erlang-type) rule.

Technical Session 3

May 15 1:30 pm – 3:00 pm

Data-Driven Optimal and Myopic Policies for Inventory Systems with Demand Covariates

Kevin Shang (Duke University), Jingkai Huang (Zhejiang University), Yi Yang (Zhejiang University), Weihua Zhou (Zhejiang University)

Companies increasingly collect auxiliary covariates (e.g., economic indicators, search trends) to improve inventory decisions. We study data-driven replenishment policies for a multi-period backorder system using historical demand and feature data. Applying Sample Average Approximation to the dynamic program yields a Data-Driven Optimal (DDO) policy that is statistically consistent but computationally expensive. To address this, we propose two efficient heuristics: a Data-Driven Myopic (DDM) policy and a Data-Driven Semi-Myopic (DDSM) policy. We establish finite-sample performance guarantees for all three policies, showing that the sample size required for ϵ-optimality scales polynomially with the planning horizon and feature dimension. A key determinant of performance is the demand signal-to-noise ratio, under which the DDM policy is near-optimal in high-signal regimes. Moreover, in small-sample settings, we prove that the DDO policy degenerates to the DDM policy. These results provide clear guidance on choosing policies based on data availability and computational resources.

Dynamic Matching for Refugee Resettlement: A Case Study

Rad Niazadeh (The University of Chicago Booth School of Business), Kirk Bansak (UC Berkeley), Soonbong Lee (Yale School of Management), Vahideh Manshadi (Yale School of Management), Rad Niazadeh (The University of Chicago Booth School of Business), Elisabeth Paulson (Harvard Business School)

Refugee resettlement is an international effort that aims to provide a durable solution for the current global refugee crisis. The goal is to help refugee families to find a new home in a host country and eventually find a new job to get “resettled”. In this talk, I will talk about a recent paper in partnership with a major national agency working on refugee resettlement in the United States. In this work, we re-design the core dynamic matching algorithm used by our partner, for sequential yearly assignment of refugee cases to our partner's affiliate locations. I discuss various operational intricacies in this dynamic matching problem---such as lack of reliable arrival prior data, predicting employment outcomes of each match, and controlling backlogs in affiliate locations--and explain the details of the design and the analysis of our primal-dual algorithm.

Exploring the Frontier of AI in Optimization Research

Michael Wagner (University of Washington)

This research presents three projects at the intersection of modern AI and optimization: (i) an Amazon deployment exploring how AI-guided modeling choices and lightweight search procedures can substantially improve the tractability of a challenging nonconvex integer program, enabling large-scale pilots; (ii) controlled "AI-as-doctoral-student" evaluations showing that top models can outperform typical PhD students on advanced optimization exam performance; and (iii) hierarchical multi-agent teamwork that materially improves success on difficult optimization problems that defeat single agents. Together, these results suggest AI can (a) unlock new algorithmic ideas, (b) accelerate technical problem solving, and (c) make complex optimization methods more usable via structured multi-agent workflows.

Cloud Value Chains in the Age of AI: An Operations Management Perspective

Shi Chen (University of Washington), Vinayak Deshpande (University of North Carolina at Chapel Hill)

The rapid adoption of artificial intelligence, particularly large language models and generative AI, is transforming cloud computing from a scalable service platform into a tightly coupled, capital-intensive value chain spanning hardware, energy, data centers, platforms, and applications. This paper provides an end-to-end survey of AI-powered cloud value chains from an operations management perspective. We highlight how upstream bottlenecks in hardware and energy supply chains, midstream challenges in AI data center capacity planning and compute pipeline design, and downstream pricing, contracting, and ecosystem strategies jointly shape system performance and value capture. A defining feature of this setting is the fundamental heterogeneity between training and inference workloads, which introduces new tradeoffs among throughput, latency, and cost. Building on this synthesis, we identify key research directions related to capacity competition, dynamic capacity expansion, service guarantees, pricing strategies, and platform ecosystem design in AI-driven cloud markets.

Poster Session 2

May 15 3:00 pm – 4:00 pm

A Distributed Decision-Making Mechanism Integrating Optimization and Reinforcement Learning for Coordinated Restoration of Interdependent Complex Networks

M. Hadi Amini (Florida International University), Shabnam Rezapour (Florida International University), Namrata Saha (Coe College)

In this work, we introduced a hybrid decision-making framework that leverages optimization and reinforcement learning (RL), and physics-based modelling of large-scale critical infrastructures to enable coordinated decision making in interdependent networks. Our solution facilitates the recovery of essential services, such as power and transportation, after a disaster. As these systems are interconnected, failure in one can negatively affect others. To this end we combined RL with traditional optimization to solve restoration challenges slowed down by decentralized management, ignorance of interdependencies, and unpredictable post-disaster conditions. Our proposed solution has three major novel contributions: Decentralized Coordination to bridge the gap between independent and interdependent networks; Adaptiveness and Robustness to manage real-time uncertainties; and Real-world Evaluation in interdependent critical infrastructures. Testing the model on power and road networks after simulated tornado damage shows the coordinated restoration improved community service levels by up to 27.9%.

Human Effort in the AI Value Chain: Improving Crowdsourced Data Annotation

Setareh Farajollahzadeh (McGill University), Rob Glew (McGill University)

The expansion of artificial intelligence has increased demand for large volumes of high-quality annotated data, much of which is produced through crowdsourcing to non-experts. Managing both the quality and quantity of output at scale remains a central operational challenge for these platforms. We study how task design -- specifically, minimum quality requirements and their sequencing -- affects performance over time when annotation is gamified. Using data from a large-scale natural experiment in crowdsourced DNA annotation, where minimum requirements were randomly assigned while rewards were held fixed, we causally identify how requirements on past tasks affect subsequent performance. We find that exposure to higher minimum requirements increases contributors’ discretionary performance on later tasks, leading to higher-quality output beyond what is necessary to earn rewards.

Selective Decision-Making for Human–AI Systems

Michael Huang (CUNY Baruch College), Bradley Rava (University of Sydney)

Artificial intelligence (AI) is increasingly used in high-risk public decision-making, where policymakers face a fundamental tradeoff between coverage, the number of decisions made autonomously by AI, and accuracy, the proportion of correct decisions. Higher coverage improves scalability but risks errors, while prioritizing accuracy increases reliance on costly and slow human review.

We develop an end-to-end selective decision-making framework that optimizes this coverage-accuracy tradeoff in Human-AI systems. By integrating selective classification, end-to-end learning, and conformal inference, the proposed approach provides theoretically rigorous control over human-AI decision allocation without ad hoc thresholds.

We apply the method to New York City speed hump requests, showing how increased AI coverage can reduce processing times, particularly in underserved neighborhoods.

History-Dependent Fluid Models for Weakly Coupled MDPs: Improving Transition Feasibility

Seyedeh Parisa Moosavi (University of Toronto), Andre Augusto Cire (University of Toronto), Selvaprabu Nadarajah (University of Illinois at Chicago)

Many sequential decision-making problems can be modeled as weakly coupled Markov decision processes (W-MDPs), in which a large-scale system decomposes into component MDPs linked by coupling constraints. Standard fluid approximations for W-MDPs typically rely on fully marginalized (aggregated) transition kernels. In this work, we investigate the inconsistencies in transition dynamics introduced by such approximations and propose novel classes of approximations to address them. Building on new insights related to reachable states, we introduce a class of history-dependent fluid models that incorporate endogenous histories to better preserve transition dynamics. We study the effect of these models on both bound strength and the quality of the policies they induce, and identify problem classes in which they outperform fully marginalized fluid models along both dimensions. Experiments on a maintenance planning problem demonstrate that our approach tightens bounds and improves policy quality.

Adaptive Behaviour of Paramedics: The Heterogeneous Impact of Paramedics Workload on Scene Decisions

Maryam Zakeri (University of Alberta), Armann Ingolfsson (University of Alberta), Mohammad Delasay (Stony Brook University), Kenneth Schultz, Graham Vanderwater (Alberta Health Services), Gerald Lazarenko (Alberta Health Services), Jim Garland (Alberta Health Services)

Emergency medical services (EMS) play a critical role in patient care, yet the variability in demand by severity, timing, and location, makes capacity planning challenging. Paramedics must stabilize, treat, refer or transport patients while managing both physical and cognitive fatigue, which may influence their decision-making. This research explores how fatigue affects EMS crew decisions on scene. We develop a workload measure that represents fatigue and employ statistical methods to establish a causal link between workload and scene outcomes. Specifically, we design a complexity–discretion framework grounded in the operational characteristics of three call types—cardiac arrest, trauma, and opioid overdose—and analyze scene time and transport decisions across these call types through this operational lens. Our findings provide insights for optimizing EMS resource management and improving patient care efficiency.

Technical Session 4

May 15 4:00 pm – 5:30 pm

Pricing and Competition for Generative AI

Rafid Mahmood (University of Ottawa)

Generative AI models are typically priced at a fixed per prompt or token of user input. However, these models yield different value to users that apply them in different tasks (e.g., coding versus writing versus math). We develop a pricing-and-competition framework that accounts for multi-task demand and interactive prompt-based use. Users seek to minimize expected total cost of prompting, yielding a price–performance competitiveness metric for each task. We analyze a two-firm sequential game where a first mover sets an initial price, the entrant prices with full market information, and users select models task-by-task. Equilibrium pricing reduces to a piecewise optimization in which each firm targets a subset of tasks determined by ranking task-level competitive ratios. We characterize when late entry guarantees positive revenue, when the first mover earns zero regardless of price, and when differentiation induces task-specific monopolies with higher sustainable prices.

LLM Pricing

Hamid Arzani (University of Toronto), Ming Hu (University of Toronto)

Specialized large language models (LLMs) are increasingly used across domains such as healthcare, finance, legal services, and technical workflows. Providers offering specialized LLM services, therefore, face a trade-off in both pricing and LLM design decisions: encouraging more within-session interactions can improve output quality and increase user value, but it raises total computational cost. We develop a model of human–LLM interaction in which a user has a latent intent in a high-dimensional embedding space and generates prompts as noisy expressions of intent. The LLM starts from a Gaussian prior, updates its belief via Bayesian learning, and produces outputs at a per-interaction computational cost. The user chooses the session length as an optimal stopping problem, with recall based on the best-so-far quality score. We characterize the conditions under which subscription versus pay-per-use is optimal and show how computational cost, prior misalignment, prior uncertainty, and prompt variability shape pricing and LLM design choices.

Improving Behavioral Alignment in LLM Social Simulations via Context and Navigation

Qianran (Jenny) Jin (The Chinese University of Hong Kong), Letian Kong (The Chinese University of Hong Kong), Renyu Zhang (The Chinese University of Hong Kong)

Large language models are increasingly used to simulate human behavior in experimental settings, but they systematically diverge from human decisions in complex decision-making environments where participants must anticipate others' actions and form beliefs based on observed behavior. We propose a two-stage framework for improving behavioral alignment. The first stage, context formation, explicitly specifies the experimental design to establish an accurate representation of the decision task and its context. The second stage, context navigation, guides the reasoning process within that representation to make decisions. We validate this framework through three studies. Across four state-of-the-art models, we find that complex decision-making environments require both stages to achieve behavioral alignment with human benchmarks, whereas the simpler demand-estimation task requires only context formation. Our findings clarify when each stage is necessary and provide a systematic approach for designing and diagnosing LLM social simulations as complements to human subjects in behavioral research.

Why the Best Machine May Not Be the Best: Incentivizing Human-Machine Collaboration

Shouqiang Wang (University of Texas at Dallas), Xiaotong Guan (University of Texas at Dallas), Anyan Qi (University of Texas at Dallas)

Robust human-machine collaboration is paramount to modernizing crucial operational decisions. We study how the machine should be designed to incentivize human contributions. A principal designs a forecasting algorithm by specifying its precision and delegates the final forecasting task to an agent. After receiving the machine’s forecast, the agent decides whether to acquire an additional private signal before making the final forecast. We find that the principal may optimally reduce the machine’s precision below its technological limit—even when enhancing such precision within that limit is costless. Essentially, the machine’s precision serves not only a functional role but also motivates the agent to gather additional information. Interestingly, the agent may over-acquire more information, resulting in possibly higher forecasting accuracy, than what would be optimal if the forecasting task were not delegated. Humans and machines can be complementary: the better the information the agent can acquire, the more precise the machine should be.

Participants

Name	Affiliation	E-mail
Osman Alp	University of Calgary	osman.alp@haskayne.ucalgary.ca
M. Hadi Amini	Florida International University	amini@cs.fiu.edu
Hamid Arzani	University of Toronto	hamid.arzani@rotman.utoronto.ca
Anant Balakrishnan	University of Texas at Austin	anantb@mail.utexas.edu
Recep Bekci	University of Waterloo	recep.bekci@uwaterloo.ca
Marco Bijvank	University of Calgary	marco.bijvank@haskayne.ucalgary.ca
Ghazal Bozorgmehry	University of Calgary	ghazal.bozorgmehrybo@ucalgary.ca
Ningyuan Chen	University of Toronto	ningyuan.chen@utoronto.ca
Shi Chen	University of Washington	shichen@uw.edu
Jianqiang Cheng	University of Arizona	jqcheng@arizona.edu
Andre Augusto Cire	University of Toronto	andre.cire@rotman.utoronto.ca
Mohammad Delasay	Stony Brook University	mohammad.delasay@stonybrook.edu
Serasu Duran	University of Calgary	serasu.duran@haskayne.ucalgary.ca
Fausto Errico	École de technologie supérieure	fausto.errico@etsmtl.ca
Setareh Farajollahzadeh	McGill University	setareh.farajollahzadeh@mcgill.ca
Zuguang Gao	University of California, Irvine	z.gao@uci.edu
Kimia Ghobadi	Johns Hopkins University	kimia@jhu.edu
X Y Han	University of Chicago	XY.Han@chicagobooth.edu
Michael Huang	CUNY Baruch College	michael.huang@baruch.cuny.edu
Raf Jans	HEC Montreal	raf.jans@hec.ca
Tracy Jenkin	Queen's University	tracy.jenkin@queensu.ca
Jenny Jin	Chinese University of Hong Kong	jennyjin@cuhk.edu.hk
Yiwen Jin	University of Calgary	yiwen.jin@ucalgary.ca
Ilbin Lee	University of Alberta	ilbin@ualberta.ca
Rafid Mahmood	University of Ottawa	mahmood@telfer.uottawa.ca
Houra Mahmoudzadeh	University of Waterloo	houra.mahmoudzadeh@uwaterloo.ca
Parvin Malekzadeh	University of Toronto	p.malekzadeh@rotman.utoronto.ca
Seyedeh Parisa Moosavi	University of Toronto	p.moosavi@mail.utoronto.ca
Mohammad Moshref-Javadi	University of Illinois at Urbana Champaign	moshref@illinois.edu
Sam Nazarinia	University of Calgary	sam.nazarinia@ucalgary.ca
Rad Niazadeh	University of Chicago	rad.niazadeh@chicagobooth.edu
Hossein Piri	University of Calgary	hossein.piri@ucalgary.ca

Name	Affiliation	E-mail
Warren Powell	Princeton University	wbpowell328@gmail.com
Danielle Ripsman	Johns Hopkins University	dripsma1@jh.edu
Feyza Sahinyazan	Simon Fraser University	fsahinya@sfu.ca
Amirhossein Shabani	University of Calgary	amirhossein.shabani@ucalgary.ca
Mohammadreza Shahsahebi	University of Calgary	seyedmohammadreza.sh@ucalgary.ca
Kevin Shang	Duke University	khshang@duke.edu
Jiong Sun	Purdue University	sun664@purdue.edu
Yu Sun	University of Toronto	yusun.sun@utoronto.ca
Sina Taheria	University of Calgary	sina.taheria@ucalgary.ca
Masoum Talebi	University of Alberta	mtalebi@ualberta.ca
Leyao Tan	University of British Columbia	leyao.tan@sauder.ubc.ca
Mustafa Tongarlak	Western University	mtongarlak@ivey.ca
Michael Wagner	University of Washington	mrwagner@uw.edu
Shouqiang Wang	University of Texas at Dallas	Shouqiang.Wang@utdallas.edu
Amy Ward	University of Chicago	amy.ward@chicagobooth.edu
Nasrin Yousefi	Queen's University	nasrin.yousefi@queensu.ca
Maryam Zakeri	University of Alberta	zakeriho@ualberta.ca

Organizing Committee

Dr. Osman Alp

Dr. Marco Bijvank

Dr. Yiwen Jin

Dr. Hossein Piri