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 of Alberta (Districts 5 and 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 first 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 and Hossein Piri

Workshop Schedule

May 14, 2025

Start	End	Topic
8:00 am	8:50 am	Registration
8:50 am	9:00 am	Welcome Message
9:00 am	10:30 am	Keynote: *Bridging AI and OM: Enhancing Decision-Making Through Machine Learning*, Pengyi Shi (Purdue University) – Part 1
10:30 am	11:00 am	Break
11:00 am	12:30 pm	Keynote: *Bridging AI and OM: Enhancing Decision-Making Through Machine Learning*, Pengyi Shi (Purdue University) – Part 2
12:30 pm	1:30 pm	Lunch (provided)
1:30 pm	3:00 pm	Keynote: *AI-Inspired Operations: How AI Is Shaping the Field of OM and How OM Can Shape AI*, Tinglong Dai (Johns Hopkins University) – Part 1
3:00 pm	3:30 pm	Break
3:30 pm	5:00 pm	Keynote: *AI-Inspired Operations: How AI Is Shaping the Field of OM and How OM Can Shape AI*, Tinglong Dai (Johns Hopkins University) – Part 2
6:00 pm	8:00 pm	Dinner (provided)

May 15, 2025

Start	End	Topic
9:00 am	10:30 am	Technical Session 1
10:30 am	11:00 am	Break
11:00 am	12:00 pm	Lightning Session 1
12:00 pm	1:00 pm	Lunch (provided)
1:00 pm	2:30 pm	Technical Session 2
2:30 pm	3:30 pm	Poster Session
3:30 pm	4:00 pm	Break
4:00 pm	5:00 pm	Lightning Session 2

Keynote Speakers

Pengyi Shi (Purdue University, Daniels School of Business)
Bridging AI and OM: Enhancing Decision-Making Through Machine Learning

This tutorial explores the integration of AI and ML into OM research. We will look at how state-of-the-art tools such as generative AI, reinforcement learning, and online learning algorithms can help to address operational challenges, such as workforce scheduling, resource allocation, and personalized intervention in dynamic and high-stakes environments. By bridging the gap between theory and practice, this tutorial aims to equip participants with actionable insights and tools to enhance the impact of OM research in industries such as healthcare, public sector, and service operations.

Tinglong Dai (Johns Hopkins University, Carey Business School)
AI-Inspired Operations: How AI Is Shaping the Field of OM and How OM Can Shape AI

The rapid advancement of artificial intelligence (AI) is driving profound change in the field of operations management (OM). While much attention has been paid to how AI can shape OM practices, the reverse perspective—how OM principles can enhance AI development, deployment, and scaling—remains little explored. This talk highlights the unique opportunity for the OM community to address this gap and produce high-quality research with broader societal impact.

OM, with its systemic perspective and analytical rigor, is uniquely positioned to address critical challenges across the AI lifecycle. OM methods can improve AI training processes, generate real-world evidence of AI's impact on productivity, efficiency, and sustainability, and balance trade-offs such as model accuracy versus pattern recognition (including those beyond human capabilities). In addition, OM can play a critical role in scaling the AI marketplace, fostering continuous learning, and ensuring alignment between AI developers and stakeholders to enhance safety and accountability.

Using examples primarily from healthcare, as well as insights from other industries, this talk will illustrate how OM can bring order to the world of AI. Time permitting, breakout sessions will encourage brainstorming of potential research topics and collaboration

Technical Session 1

May 15 9:00 am – 10:30 am

Contextual Linear Optimization with Bandit Feedback

Yichun Hu (Cornell University), Nathan Kallus (Cornell University), Xiaojie Mao (Tsinghua University), Yanchen Wu (Tsinghua University)

Contextual linear optimization (CLO) uses predictive contextual features to reduce uncertainty in random cost coefficients and improve average-cost performance. An example is the stochastic shortest path problem with random edge costs and contextual features. Existing work on CLO assumes the data has fully observed cost coefficient vectors, but in many applications, we can only see the realized cost of a historical decision (i.e., bandit feedback). We study a class of offline learning algorithms for CLO with bandit feedback, which we term induced empirical risk minimization (IERM), where we fit a predictive model to directly optimize the downstream performance of the policy it induces. We show a fast-rate regret bound for IERM that allows for misspecified model classes and flexible choices of the optimization estimate, and develop computationally tractable surrogate losses. We compare the performance of different modeling choices numerically using a stochastic shortest path example and provide practical insights.

When Less is More: Optimizing Prescription Alerts under Fatigue

Michael Lingzhi Li (Harvard University), Hossein Piri (University of Calgary)

As healthcare systems grow more complex, pharmacists play a crucial role in preventing medication errors. Computerized Provider Order Entry (CPOE) systems assist by identifying drug interactions and dosage issues, but excessive alerts lead to fatigue, reducing effectiveness. This study develops a fluid model to optimize CPOE alert strategies under pharmacist fatigue, balancing risk avoidance and alert management. We find that the optimal strategy follows a dynamic threshold structure, adjusting alert frequency as fatigue accumulates. Risk distributions with heavier tails require more conservative policies, while a constant threshold policy remains near-optimal in the long run, simplifying implementation. Numerical results using real-world data show that the optimized strategy significantly reduces patient risk while minimizing unnecessary alerts.

Balancing Personalized Medicine and One-Size-Fits-All: Navigating Trade-Offs in Complexity, Efficiency, and Equity

Narges Mohammadi (Imperial College London), Niloofar Zamani Foroushani (Imperial College London), Reza Skandari (Imperial College London)

Personalized policies in medical care improve outcomes by tailoring guidelines to individual characteristics but are often resource-intensive and challenging to implement. To address this, we propose an interpretable decision-tree-based clustering method to aggregate subpopulations, aiming to minimize health inequalities while ensuring the guidelines remain practical and implementable.

Boosted Generalized Normal Distributions for Modelling Patient Wait and Service Time Distributions in EDs

Donald Lee (Emory University), Ragip Gurlek (Emory University), Francis de Vericourt (ESMT)

Patient time-to-treatment is a critical benchmark for evaluating emergency department (ED) performance. Over the past decade, machine learning (ML) models have been developed to improve wait time forecasting, but they do not incorporate well-established insights from the healthcare operations literature. Specifically, ED wait times are known to follow an approximately exponential distribution, while service times are approximately log-normal. We introduce bGND, a rigorous gradient-boosted distributional ML model that integrates these distributional knowledge into forecasting. Applied to data from a large U.S. academic ED, bGND improves distributional forecasting accuracy by between 6.1% to 8.8% compared to a leading nonparametric ML benchmark. These gains translate into a 9.4% reduction in patient dissatisfaction, and an estimated increase of $120,000 in hospital earnings per 10,000 visits.

Lightning Session 1

May 15 11:00 am – 12:00 pm

Supply Chain Coordination under Unknown Demand Distribution: Online Learning and Contracting Mechanisms

Shi Chen (University of Washington), Mengxiao Zhang (University of Iowa), Haipeng Luo (University of Southern California), Yingfei Wang (University of Washington)

Multi-echelon stochastic inventory models with known demand distributions have been fundamental in supply chain management (SCM), enabling optimal policies for centralized supply chains and contract designs to coordinate decentralized systems. This work extends the classic two-echelon inventory problem to an online setting with unknown demand distributions, introducing novel challenges. First, upstream firms only observe downstream orders, not true demand, further complicated by downstream experimentation with policies under uncertainty. Second, the multi-echelon nature creates complex inventory interactions, where upstream shortages impact downstream inventory levels. For the centralized system, we propose algorithms based on Online Newton Step with a low-switching property to converge to the first-best policy. For the decentralized system, we develop algorithms and a protocol for learning optimal contract parameters, achieving individual regret guarantees and ensuring convergence to first-best decisions, thereby coordinating the supply chain under an online, unknown-demand setting.

(Machine) Learning Preferences from Complex Choice Sets: An Application to Service Networks

Ken Moon (University of Pennsylvania)

Hidden complexity pervades commonplace choices. Customers at venues such as supermarkets, shopping centers, and amusement parks routinely choose from astronomically many available paths through the venues' stations (e.g., sections, stores, rides). Such choice complexity is captured by rich data and poses a challenge for empirical research. Even state-of-the-art choice estimation methods suffer serious issues in estimating consumer preferences from such data. Under common conditions, a minute fraction of choice utility comparisons provably captures all customer preference information. This hidden low-dimensional structure is automatically learned by machine learning. We develop a neural network-based estimator of customer preferences and demonstrate it on data tracking shoppers at an urban hypermarket serving 1M+ customers annually. Our ML-based method uncovers the hypermarket's demand structure and counterfactually raises its throughput by 5-25% through modest capacity reallocations. We address a long-standing challenge in the literature by using congestion shocks' cross-station demand effects to identify truly complementary consumption.

Multiclass Queue Scheduling Under Slowdown: An Approximate Dynamic Programming Approach

Berk Gorgulu (McMaster University), Jing Dong (Columbia University), Vahid Sarhangian (University of Toronto)

In many service systems, customer waiting times can result in increased service requirements. Such service slowdowns can significantly impact system performance. Scheduling under wait-dependent service times is challenging, especially when multiple customer classes are heterogeneously affected by waiting. In this work, we study scheduling policies in multiclass, multiserver queues with wait-dependent service slowdowns. We propose a simulation-based Approximate Dynamic Programming (ADP) algorithm to find close-to-optimal scheduling policies. The ADP algorithm (i) represents the policy using classifiers based on the index policy structure, (ii) leverages a coupling method to estimate the differences of the relative value functions directly, and (iii) uses adaptive sampling for efficient state-space exploration. Through extensive numerical experiments, we illustrate that the ADP algorithm generates close-to-optimal policies that outperform well-known benchmarks. We also provide insights into the structure of the optimal policy and conduct a case study on scheduling admissions into

Training Neural Networks for the *GI/G/K* Queue

Qi-Ming He (University of Waterloo), Haoran Wu (Sun Yat-sen University), Zhenggao Wu (University of Waterloo), Haokun Zhao (University of Waterloo)

We introduce a program for training deep neural networks to compute an approximation of the stationary distribution of the queue length for the GI/G/K queue, which have few analytic solutions. This work is an extension of a recent study. The key contribution of the program is an effective algorithm for generating samples (data) for supervised learning. The algorithm, which is based on the matrix-analytic methods, computes the stationary distributions of the PH/PH/K queues, which are dense in the set of the GI/G/K queues. Specifically, the program brings together the quasi birth-and-death process, the CSFP (count-server-for-phase) method, and the matrix-geometric solution to generate a large number of samples for training and validating neural networks for the GI/G/K queues. The effectiveness of the program is compared to existing asymptotic methods such as simulation, heavy traffic, and heuristic approximations. The work sheds lights on the training of neural networks for complex stochastic systems.

Technical Session 2

May 15 1:00 pm – 2:30 pm

Build Agility and Resilience by Leveraging Machine Learning for Supply Chain Planning

Morris Cohen (University of Pennsylvania), Naren Agrawal (Santa Clara University), Vinayak Deshpande (University of North Carolina)

The global pandemic and other recent geopolitical events have highlighted the importance of supply chain resilience and agility. Unfortunately, many companies have failed to develop strategies that can deliver these capabilities. In particular, the most promising enablers, i.e., artificial intelligence and machine learning, cloud computing, and big data, have fallen short. In this article, we explain reasons why the application of machine learning can fail to deliver success in the context of supply chain planning. We describe Optimal Machine Learning, a new paradigm that we have developed, that leverages the capabilities of advanced technology. We illustrate the power of this new approach for supply chain planning by reporting on two successful implementations in complex global supply chain environments. We then discus a framework to guide efforts to design effective strategies for supply chain resilience and agility. The paper concludes with a review of research applications.

Biases of Humans, of AI, and of Humans with AI

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

Since ChatGPT's release in November 2022, research on decision biases in large language models (LLMs) has rapidly proliferated. LLMs. Although they are mathematical models representing the current frontier of Artificial Intelligence (AI), they are trained on human data. Research has shown that LLMs may mirror human biases, remain unbiased, or even display biases that differ from those of humans. This paper advances the literature by examining whether individuals change their inherent biases when interacting with AI that might be biased similarly or differently from humans. To study this question, we design a novel interactive AI experimental design paradigm and apply it to three representative biases from the literature. We show that AI biases indeed impact human decision-making and discover a new form of nudging via system prompts. Finally, we propose a new take on algorithm aversion by showing that how humans interact with AI matters for improving decision-making.

Cognitive Costs and Ethical Choices: Fairness in Human-Machine Decision Systems

Mohammadreza Shahsahebi (University of Calgary), Osman Alp (University of Calgary), Justin Weinhardt (University of Calgary), Alireza Sabouri (University of Calgary)

Machine Learning (ML) biases pose serious challenges in fields such as healthcare, finance, and human resources. This research studies collaborative decision-making between a rationally inattentive human decision-maker (DM) and a biased ML algorithm. By modeling the DM's limited cognitive capacity and potential Machine biases, we examine how fairness is affected when the DM is aware or unaware of the Machine's bias. Our research reveal explicit thresholds for fairness within belief spaces and highlight systematic advantages or disadvantages conferred on different groups. Moreover, this research shows when the DM remains unaware disadvantaged groups face harsher outcomes, which underscores the need for transparent AI systems that inform DMs about potential algorithmic biases.

Rethinking Algorithmic Fairness for Human-AI Collaboration

Hamsa Bastani (University of Pennsylvania), Osbert Bastani (University of Pennsylvania), Haosen Ge (University of Pennsylvania)

Existing approaches to algorithmic fairness aim to ensure equitable outcomes if human decision-makers comply perfectly with algorithmic decisions. However, perfect compliance is rarely a reality or even desirable in human-AI collaboration. Recent studies show that selective compliance with fair algorithms can amplify discrimination relative to the prior human policy. Ensuring equitable outcomes thus requires fundamentally different design principles that ensure robustness to the decision-maker's (a priori unknown) compliance pattern. We define compliance-robustly fair algorithmic recommendations that guarantee (weakly) improved fairness regardless of compliance behavior. We propose an optimization strategy to identify the best performance-improving compliance-robustly fair policy. However, we show that it may be infeasible to design recommendations that are simultaneously fair in isolation, compliance-robustly fair, and more accurate than the human policy. Thus, improving equity and accuracy in human-AI collaboration may not always require enforcing traditional fairness constraints. We illustrate this on Virginia criminal sentencing data.

Poster Session

May 15 2:30 pm – 3:30 pm

OPT2CODE: RAG Framework for Generating Linear Programming Solver Code

Tasnim Ahmed (Queen's University), Salimur Choudhury (Queen's University)

Linear Programming (LP) problems aim to find the optimal solution to an objective under constraints. This study explores the efficiency of Large Language Models (LLMs) in generating solver-specific code from LP problem descriptions. We propose OPT2CODE, a Retrieval-Augmented Generation (RAG) framework that utilizes LLM-as-a-Judge components to transform textual problem descriptions into Gurobi solver code. Empirical studies on our curated dataset show that our framework improves baseline LLM performance by 8.65%. We also provide an analysis of energy consumption for both proposed and existing methods.

The Average Tree of a Random Forest

Rim Hariss (McGill University), Guillaume Pouliot (University of Chicago), Hengguang Mao (University of Chicago)

Random forests are effective in ensemble learning due to their ability to reduce variance while maintaining low bias. However, their lack of interpretability makes them less useful in scenarios requiring transparency. In contrast, decision trees are highly interpretable but suffer from high variance. We introduce the average tree, a meta-learning approach that retains the predictive efficiency of random forests while ensuring transparency.

Previous studies on meta-learner decision trees focused on simplifying ensemble models but neglected stability, making them unreliable under small data perturbations. Our work aims to construct a stable, single-tree representation of a random forest, preserving its predictive power while ensuring consistency. This advancement enhances model transparency, making random forests more interpretable for practical applications.

Job Insecurity in the Age of Algorithms: Employee Reactions to Hiring Decision Aids

Mehnaz Rafi (University of Calgary), Justin Weinhardt (University of Calgary)

This is a 4-experiment project, but we only present one study here due to space limitations. In the presented study, we conducted a randomized experiment to examine employees' reactions to their organization's use of hiring algorithms, outsourcing to hiring experts, and undergraduate student interns. Serial mediation analysis of the data showed that participants (n = 465) experienced higher levels of job insecurity and felt more threatened by the company's use of experts and interns than by algorithms. As a result, they were the least likely to display overconfidence when comparing their performance to an algorithm and more likely to use an algorithm to complete future job tasks than experts or interns.

.

Seeding with Differentially Private Network Information

Amin Rahimian (University of Pittsburgh), Fang-Yu Yu (George Mason University), Yuxin Liu (University of Pittsburgh), Carlos Hurtado (University of Pittsburgh)

In public health interventions such as the distribution of preexposure prophylaxis (PrEP) for HIV prevention, decision makers rely on seeding algorithms to identify key individuals who can amplify the impact of their interventions. In such cases, building a complete sexual activity network is often infeasible due to privacy concerns. Instead, contact tracing can provide influence samples, that is, sequences of sexual contacts without requiring complete network information. This presents two challenges: protecting individual privacy and adapting seeding algorithms to work with incomplete network information. To address these problems, we study privacy guarantees for influence maximization when the social network is unknown and the inputs are samples of prior influence cascades that are collected at random and need privacy protection. Theoretical performance guarantees and simulations on empirically grounded sexual networks reveal the diminishing value of network information with decreasing privacy budget and a graceful decrease in performance with decreasing privacy budget.

ML-based Dynamic Uncertainty Sets to Optimize Hybrid Renewable Energy Systems

Ahmed Saif (Dalhousie University), Ali Keyvandarian (Dalhousie University)

We propose a new class of dynamic uncertainty sets (DUSs) to be used within the robust optimization (RO) framework to capture statistical correlations in the uncertain parameters, thus leading to less-conservative solutions than classical static uncertainty sets. These sets are employed to robustly design hybrid renewable energy systems (HRESs) that combine conventional and renewable energy (RE) generation and storage technologies to serve stand-alone communities' demands. First, DUSs based on time-series models (e.g., AR, MA or ARMA) are designed to independently account for temporal autocorrelations in demand and wind and solar energy supply. Next, the cross-correlation among these parameters is additionally considered by introducing DUSs based on vector autoregressive (VAR) and neural network (NN) models. We show how the adaptive RO models can be tractably reformulated and propose a column-and-constraint-generation algorithm to solve large instances. Numerical experimentation on real-life cases demonstrated the added-value of the DUSs in improving the HRES designs.

Evaluating Drone-Delivery Efficiency in Different Urban Settings Using Graph Neural Networks

Bahar Viniche (York University), Opher Baron (University of Toronto), Mehdi Nourinejad (York University)

Last-mile delivery is a critical yet costly phase of logistics, requiring innovative solutions to enhance efficiency in urban and suburban areas. Drone delivery optimizes last-mile logistics by reducing delivery times and overcoming challenges such as parking limitations in dense urban regions and ensuring timely deliveries to remote locations. Evaluating its effectiveness across diverse urban settings is essential. To achieve this, we use Graph Neural Networks (GNNs), leveraging city infrastructure’s networked nature to assess drone delivery efficiency. GNNs apply deep learning to graph-structured data, modeling urban structures and routing metrics for accurate evaluation. Effectiveness is measured through travel time reductions and environmental impacts, such as fuel consumption. We developed the Drone Sidekick Tool, an interactive platform that collects and analyzes data on travel distance, time, and environmental factors. Combined with our algorithmic approach, this tool offers insights into drone delivery’s potential, aiding the development of more efficient and adaptable logistics solutions.

Lightning Session 2

May 15 4:00 pm – 5:00 pm

Inference-Aware Policy Optimization

Osbert Bastani (University of Pennsylvania), Hamsa Bastani (University of Pennsylvania), Bryce McLaughlin (University of Pennsylvania)

There has been a surge of interest in deploying automated decision making policies to target interventions. However, these policies often fail to improve performance according to rigorous evaluation methodologies due to high variance in the quality of decisions across individuals. The key issue is that these policies are typically trained to improve expected performance, ignoring variance in outcomes. To bridge this gap, we propose a novel policy learning algorithm that optimizes for statistical significance. We show how to use machine learning to approximate this objective; while the resulting objective is highly non-convex, we show how it can be rewritten as a family of convex problems that can be solved in an efficient manner. We provide experiments validating our methodology. Our work represents a key step towards policy learning algorithms designed to support downstream evaluation procedures.

Data-Driven Contextual Stochastic Optimization for Adaptive Capacity Planning

Borzou Rostami (University of Alberta), Saleh Farham (University of Alberta), Fahimeh Rahimi (University of Alberta), S. Davod Hosseini (Saint Mary's University)

This study develops a data-driven Contextual Stochastic Optimization (CSO) framework for adaptive capacity planning under demand uncertainty. Traditional stochastic optimization models treat estimation and decision-making as separate processes, relying on unconditional probability distributions, which often overlook external factors influencing demand variability. In contrast, CSO integrates contextual information (such as temporal, spatial, and environmental factors) directly into the optimization process to enhance decision-making accuracy. We focus on capacity planning problems commonly encountered in transportation networks, supply chains, and resource allocation systems, where capacity decisions must adapt to fluctuating demand. Formally, we propose a two-stage CSO model, where first-stage strategic decisions and second-stage capacity decisions are optimized based on contextual covariates correlated with uncertain demand. Using historical data, we develop several approximation methods leveraging machine learning techniques. Furthermore, we establish theoretical performance guarantees and conduct a comprehensive computational study to evaluate the effectiveness and practical relevance of the proposed approaches

Generative AI for Operations Management: Applications in Simulation, Qualitative Analysis, and Causal Graphs

Mohammad Jalali (Harvard University)

Generative AI, particularly Large Language Models (LLMs), is transforming Operations Management (OM) by enhancing efficiency in simulation modeling, qualitative analysis, and causal inference. Despite their potential, challenges remain in ensuring rigor, interpretability, and reliability. This talk explores three AI-driven applications conducted in our lab over the last year: (1) using LLM to streamline problem articulation and documentation in simulation modeling of COVID-19's economic impact, (2) applying AI in qualitative research on obesity prevention, where it improves efficiency but struggles with contextual understanding, and (3) developing novel causal graph comparison methods using AI-generated datasets to address inconsistencies in automated causal modeling. Across these applications, we emphasize AI's role as an augmentative tool rather than a replacement for human judgment, providing insights into best practices for integrating AI into OM research while maintaining analytical depth and decision-making accuracy.

Racing with Generative AI: Impact of Conversational Agents on Mental Health Counselors

Seokjun Youn (University of Arizona), Chenxi Guo (Beijing Institute of Technology), Xunyu Chen (Virginia Commonwealth University), Karen Xie (University of Connecticut), Wei Chen (University of Connecticut)

Using Generative Artificial Intelligence (GenAI) for mental health is a major healthcare advancement, yet its disruption to incumbent counselors is understudied. We compare counselors exposed to a GenAI chatbot with an unaffected group and find that counselor appointments rose 24.7%, driven by a competition effect. Counselors amplified engagement—29.5% more replies, 14.9% more endorsements, 9.04% deeper interactions—especially when GenAI improved, indicating a "catfish effect" where the competition spurs counselors to improve. LLM-driven content analysis reveals that, post-GenAI launch, counselors increased informational support and reduced emotional support, contrary to the intuitive assumption that human counselors may leverage their emotional strengths to differentiate from AI. Counselors also mirror GenAI's support style, offering more targeted assistance when the agent does the same. Our findings highlight how counselors adapt to GenAI's rise, offering strategic insights for digital mental healthcare, technology management, and the evolving future of work in mental health services.

Participants

Name	Affiliation	E-mail
Naren Agrawal	Santa Clara University	nagrawal@scu.edu
Tasnim Ahmed	Queen's University	tasnim.ahmed@queensu.ca
Osman Alp	University of Calgary	osman.alp@haskayne.ucalgary.ca
Hamsa Bastani	University of Pennsylvania	hamsab@wharton.upenn.edu
Osbert Bastani	University of Pennsylvania	obastani@seas.upenn.edu
Mehmet Begen	Western University	mbegen@ivey.uwo.ca
Marco Bijvank	University of Calgary	marco.bijvank@haskayne.ucalgary.ca
Shi Chen	University of Washington	shichen02@gmail.com
Morris Cohen	University of Pennsylvania	cohen@wharton.upenn.edu
Giovani da Silveira	University of Calgary	giovani.dasilveira@haskayne.ucalgary.ca
Tinglong Dai	Johns Hopkins University	dai@jhu.edu
Serasu Duran	University of Calgary	serasu.duran@haskayne.ucalgary.ca
Pnina Feldman	University of Virginia	feldmanp@darden.virginia.edu
Mike Galbreth	University of Tennessee	galbreth@utk.edu
Berk Gorgulu	McMaster University	gorgulub@mcmaster.ca
Rim Hariss	McGill University	rim.hariss@mcgill.ca
Qi-Ming He	University of Waterloo	q7he@uwaterloo.ca
Yichun Hu	Cornell University	yh767@cornell.edu
Mohammad Jalali	Harvard University	msjalali@mgh.harvard.edu
Tracy Jenkin	Queen's University	tracy.jenkin@queensu.ca
Donald Lee	Emory University	donald.lee@emory.edu
Sun Ju Lee	Georgia Institute of Technology	julee@gatech.edu
Michael Lingzhi Li	Harvard University	mili@hbs.edu
Yang Li	Western University	yali@ivey.ca
Alys Liang	McGill University	alys.liang@mcgill.ca
Richard Liu	University of Calgary	runqian.liu@ucalgary.ca
Narges Mohammadi	Imperial College London	n.mohammadi19@imperial.ac.uk
Ken Moon	University of Pennsylvania	kenmoon@upenn.edu
Hossein Piri	University of Calgary	hossein.piri@ucalgary.ca
Mehnaz Rafi	University of Calgary	mehnaz.rafi@ucalgary.ca
Amin Rahimian	University of Pittsburgh	rahimian@pitt.edu
Borzou Rostami	University of Alberta	borzou@ualberta.ca
Alireza Sabouri	University of Calgary	alireza.sabouri@haskayne.ucalgary.ca
Ahmed Saif	Dalhousie University	ahmed.saif@dal.ca
Saied Samiedaluie	University of Alberta	samiedal@ualberta.ca
Amirhossein Shabani	University of Calgary	amirhossein.shabani@ucalgary.ca

Name	Affiliation	E-mail
Mohammadreza Shahsahebi	University of Calgary	seyedmohammadreza.sh@ucalgary.ca
Mahdi Shakeri	University of Calgary	mahdi.shakeri@ucalgary.ca
Pengyi Shi	Purdue University	shi178@purdue.edu
Soroush Shojaeimehr	University of Calgary	soroush.shojaeimehr@ucalgary.ca
Sina Taheria	University of Calgary	sina.taheria@ucalgary.ca
Bahar Viniche	York University	bahardv@yorku.ca
Zhenggao Wu	University of Waterloo	z365wu@uwaterloo.ca
Seokjun Youn	University of Arizona	syoun@arizona.edu
Christos Zacharias	University of Miami	czacharias@bus.miami.edu

Organizing Committee

Dr. Marco Bijvank

Dr. Osman Alp

Dr. Hossein Piri