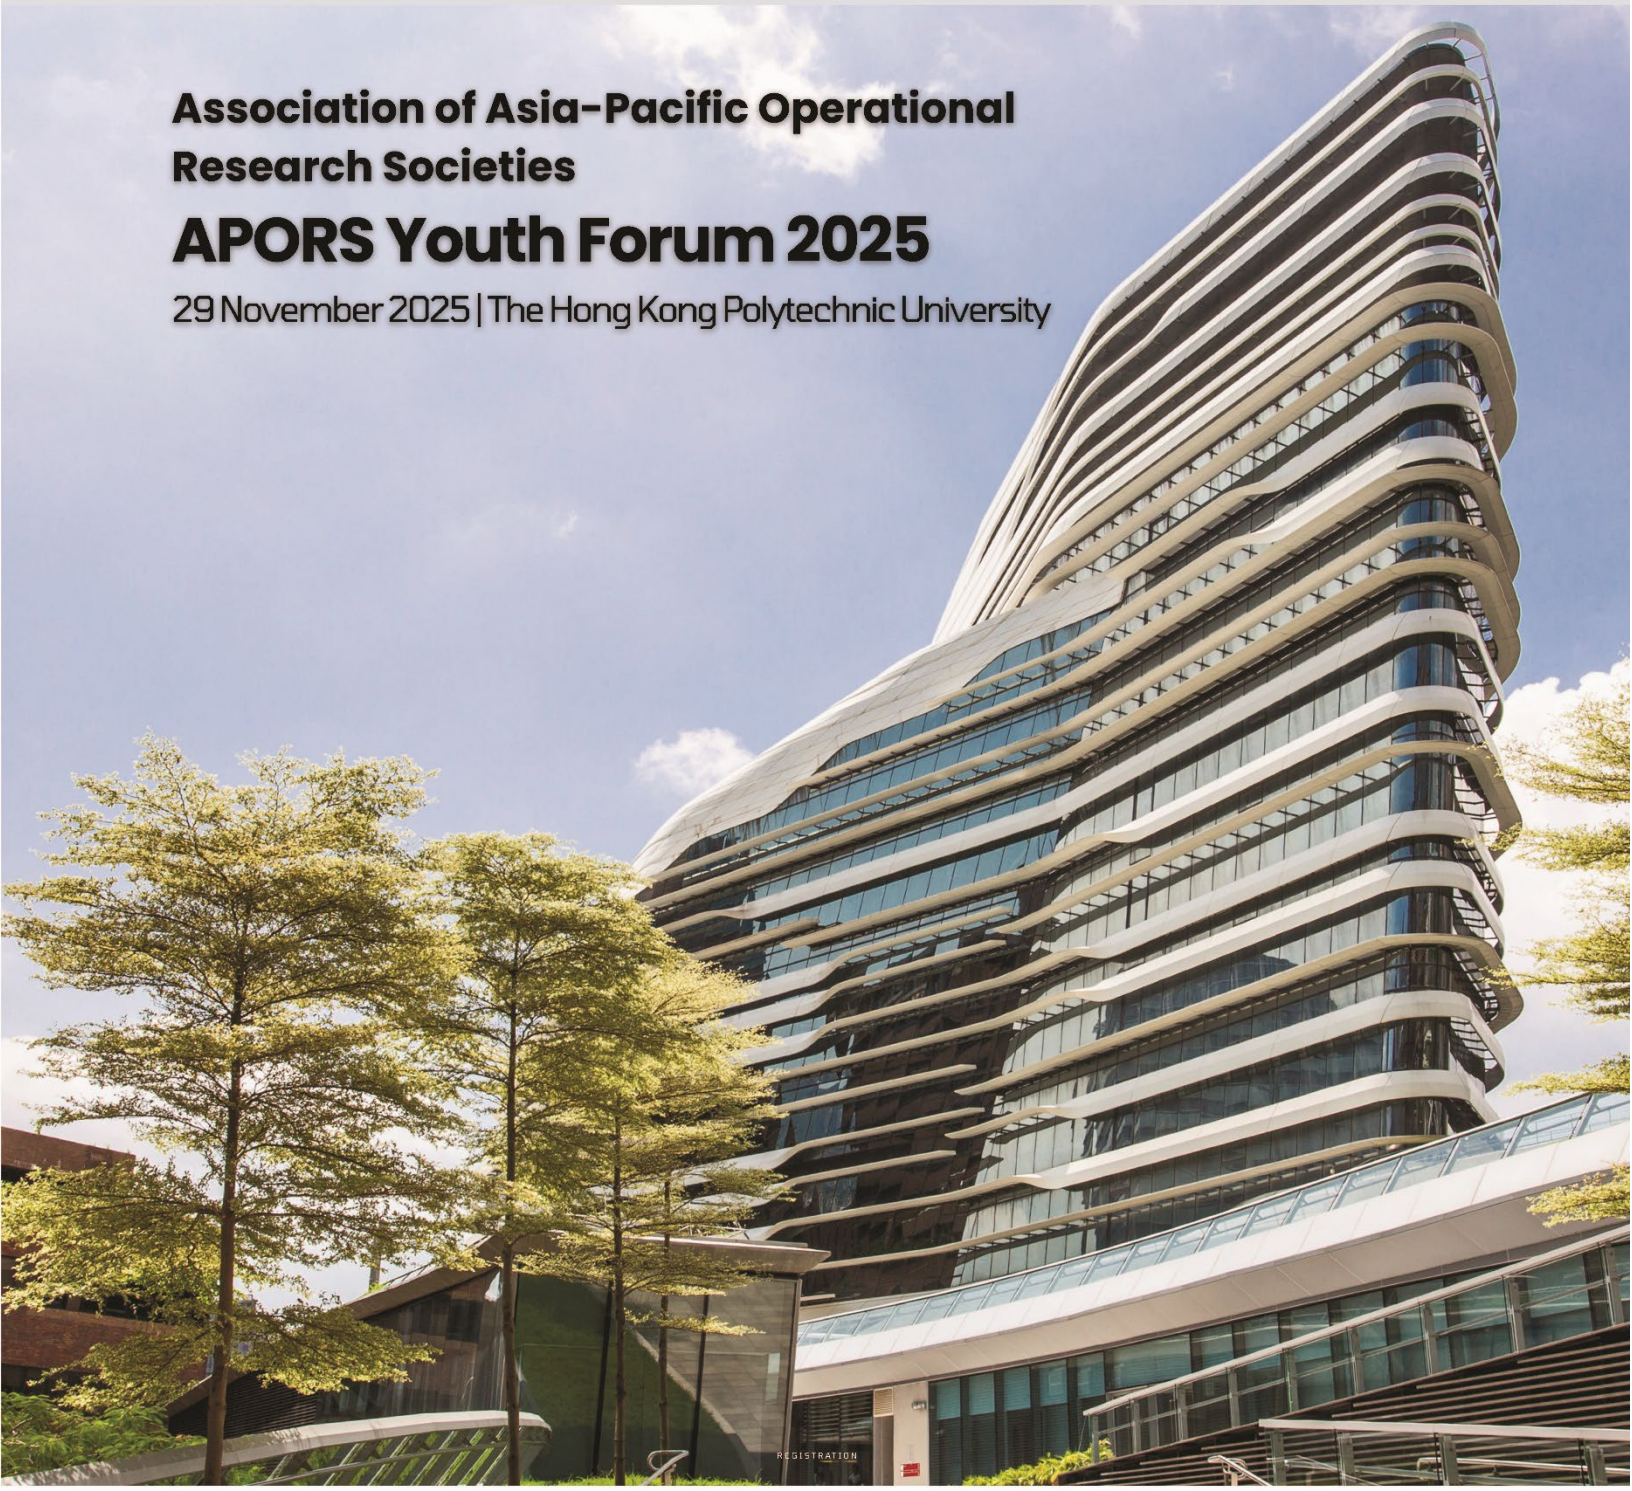

Time: 8:30 to 21:00

Venue: Lecture Theater V322, Jockey Club Innovation Tower (Core V), The Hong Kong Polytechnic University

APORS Youth Forum 2025

Jointly organised by

Association of Asia-Pacific Operational Research Societies (APORS),

Operational Research Society of Hong Kong (ORSHK)

and

Department of Applied Mathematics (AMA)
& Department of Logistics and Maritime Studies (LMS),
The Hong Kong Polytechnic University (PolyU)

Date: 29 November 2025
Time: 08:30 to 21:00
Venue: Lecture Theater V322, Jockey Club Innovation Tower (Core V), The Hong Kong Polytechnic University

Time	Event
08:30 – 09:00	**Registration** Venue: V322
09:00 – 09:15	**Opening Session** Venue: V322 Welcome speeches by *Professor Honglei XU* *President of APORS* *Professor Defeng SUN* *Head of Department, AMA, PolyU*

09:15 – 10:15	**IFORS Tutorial Lecture** Venue: V322 Mathematical Optimization and Operations Research for AI *Professor Yinyu YE* *Professor Emeritus of Stanford and* *Visiting Professors of SJTU, CUHKSZ, HKUST, and SIMIS*
10:15 – 10:30	**Coffee Break**
10:30 – 11:30	**Panel Discussion** Venue: V322 From PhD to Professional: Career Development in Operational Research Panelists: *Professor Yukun CHENG (Jiangnan University)* *Professor Yu-Hong DAI (Chinese Academy of Sciences)* *Professor Janny LEUNG (University of Macau)* *Professor Honglei XU (Curtin University)* *Professor Yinyu YE (Stanford University)* Moderator: *Professor Yancheng YUAN (The Hong Kong Polytechnic University)*
11:30 – 12:00	**APORS Best Paper Award Ceremony** Venue: V322
12:15 – 14:00	**Lunch Break** Venue: Ju Yin House Seafood Restaurant, 4/F, Communal Building, The Hong Kong Polytechnic University
14:00 – 15:45	**Parallel Sessions (I)**
15:45 – 16:00	**Coffee Break**
16:00 – 17:45	**Parallel Sessions (II)**
18:00 – 21:00	**Banquet** Venue: Ju Yin House Seafood Restaurant, 4/F, Communal Building, The Hong Kong Polytechnic University

Contents

Parallel Sessions (I)

Time: 14:00 – 15:45

Optimization Theory and Algorithms I **Venue: V302** **Session Chair: Lexiao Lai**	**Advances in Data Science I** **Venue: V312** **Session Chair: Ruicheng Ao**	**Game Theory I** **Venue: V321** **Session Chair: Junyue Zhang**	**Logistics and Supply Chains I** **Venue: V303** **Session Chair: Abhay Sobhanan**	**Stochastic Models** **Venue: V315** **Session Chair: Wei Liu**	**Strategies and Control Techniques** **Venue: V316** **Session Chair: Jelena Stanković**
A Two-Strategy Acceleration Framework for Phase Retrieval: FAM via Exact Minimization and FRWF via First-Order Approximation Linbin Li and Yong Xia	Graph Topology Adaptive Judgment Against Node Label Noise Mengyao Zhou, Xiao Han, Wei Wei and Guiying Yan	Beyond Regularity: Simple versus Optimal Mechanisms, Revisited Yiding Feng and Yaonan Jin	Dynamic optimization of differentiated pricing for crowdsourced delivery couriers with random utility theory Chunwei Yang and Andy Chow	Exact Analytical Option Pricing Formula Under a Jump-Diffusion Model with Hawkes Intensity Daniel Wei-Chung Miao, Xenos Chang-Shuo Lin and Ruey-An Shiu	Open pit mine planning using optimal control techniques Bowen Zhao, Honglei Xu and Lay Teo Kok
Convex Facial Reduction Algorithm and Strong Extended Duals Ying Lin, Tianxiang Liu and Bruno Lourenço	Optimal pivot path of the simplex method for linear programming based on reinforcement learning Anqi Li, Tiande Guo, Congying Han, Bonan Li and Haoran Li	A Novel Sample-Based Best Response Dynamics for Solving Nash Equilibrium in Continuous Games Zhouming Wu, Xu Xia and Yifen Mu	A Hybrid AI-Driven Forecasting Framework for Hotel Occupancy: Combining Machine Learning, Managerial Insight, and Customer Feedback Tian Tian, Aneng He and Lai Soon Lee	Provable Policy Gradient for Robust Average-Reward MDPs Beyond Rectangularity Qiuhao Wang, Yuqi Zha, Chin Pang Ho and Marek Petrik	Project Monitoring and Control Systems Hall, Zhuoyu Long, Jin Qi and Yuhao Yan

On resolution of ℓ_1-norm minimization via a two-metric adaptive projection method Hanju Wu and Yue Xie	Purity Law for Neural Routing Problem Solvers with Enhanced Generalizability Wenzhao Liu, Haoran Li, Congying Han, Zicheng Zhang, Anqi Li and Tiande Guo	Mechanism Design for Exchange Markets Yukun Cheng	Fair Human–Robot Collaborative Logistics in Industry 5.0 Yang Deng and Andy Chow	Multi-Objective Optimization Model for Reserve and Credit Fund Allocation in Banking under Conditional Value-at-Risk Constraints Moch Panji Agung Saputra, Diah Chaerani, Sukono and Mazlynda Md Yusuf	Proficiency Boosts Efficiency: Extended Johnson's Rules for Flow Shop Scheduling with Learning and Forgetting Effects Fenglin Zhang, Yuli Zhang and Zuo-Jun Max Shen
On the diameter of subgradient sequences in o-minimal Structures Lexiao Lai and Mingzhi Song	Optimizing LLM Inference: Fluid-Guided Online Scheduling with Memory Constraints Ruicheng Ao	Arbitrage on Decentralized Exchanges Yutian Zhou, Xuedong He and Chen Yang	Dynamic joint optimization of autonomous vehicle-based mobile warehousing and crowdsourced delivery routing with the use of reinforcement learning Zeyu Zhang, Andy Chow and Chunwei Yang	A Spider-Type Stochastic Subgradient Method For Expectation-Constrained Nonconvex Nonsmooth Optimization Wei Liu and Yangyang Xu	Concentration or diversification? The optimal strategy in R&D project exploration Qian Liu, Tong Lu and Jin Qi
		Last-iterate convergence in zero-sum games: Alternating heterogeneous dynamics to the rescue Xinxiang Guo, Junyue Zhang, Yifen Mu, Xiao Wang and Ioannis Panageas	A Hybrid Genetic Algorithm for the Rural Postman Problem with Multiple Trucks and Drones Abhay Sobhanan, Hadi Charkhgard and Changhyun Kwon		National Innovation System Efficiency Assessed by DEA: Lessons for Policy and Growth Jelena Stanković, Ivana Marjanović and Marina Stanojevic

Parallel Sessions (II)

Time: 16:00 – 17:45

Optimization Theory and Algorithms II **Venue: V302** **Session Chair: Yue Xie**	**Advances in Data Science II** **Venue: V312** **Session Chair: Du Chen**	**Game Theory II** **Venue: V321** **Session Chair: Yiding Feng**	**Logistics and Supply Chains II** **Venue: V303** **Session Chair: Zihao Li**	**Robust Optimization** **Venue: V315** **Session Chair: Jie Wang**	**Transportation** **Venue: V316** **Session Chair: Jinha Hibino**
LLM Serving Optimization with Variable Prefill and Decode Lengths Meixuan Wang, Yinyu Ye and Zijie Zhou	A Logic-Mined, PSO-Tuned Multi-Architecture Ensemble Neural Network Framework for Student Mathematics Learning Potential Prediction in MET-Supported Classrooms Aneng He, Lai Soon Lee and Tian Tian	Screening or Pooling? Dual-channel Competition in Healthcare under Information Asymmetry Fangyuan Cao, Jie Wang, Feng Tian and Yong-Hong Kuo	Reliable Supply Network Design under Correlated Uncertainties Caused by Disruptions Wenjie Li, Miao Song, Jia Shu and Yongzhen Li	Fast Algorithm for Wasserstein Distributionally Robust Markov Decision Processes Ling Dai, Zhuodong Yu and Chin Pang Ho	Adaptive bus bridging strategy under rail transit disruptions via a two-stage reinforcement learning driven approach Junyi Wang, Andy Chow, Zeyu Zhang and Yukun Xu
Adaptively Robust LLM Inference Optimization under Prediction Uncertainty Zixi Chen, Yinyu Ye and Zijie Zhou	Pure Exploration via Frank-Wolfe Self-Play Xinyu Liu, Chao Qin and Wei You	An Evolutionary Game-Theoretic Analysis of Incentive Mechanisms for Cybersecurity in Blockchain-Enabled Smart Home Supply Chains Anu Rao, Amit Singh and Devi Prasad Dash	Data-Driven Robust Inventory Management with Time-Series Demand Zhi Chen, Tong Wang and Fen Xu	Inpatient Spillover: A Distributionally Robust Chance-Constrained Approach Qihao Wu, Yinuo Yu, Abdel Lisser, Ka Chung Abraham Wai and Yong-Hong Kuo	Adversarial multi-agent reinforcement learning for robust train service scheduling under feeder bus service uncertainties Shouyi Wang, Andy Chow and Junyi Wang

Spectral method for general Group Synchronization Ziyue Zhao, Huikang Liu and Man-Chung Yue	Distributionally Robust Federated Learning: An ADMM Algorithm Wen Bai, Yi Wong, Ling Dai, Xiao Qiao and Chin Pang Ho	Robust Multi-Object Auction Design under Knightian Preferences Xiaojin Fu, Anh Quan Nguyen and Jin Qi	Assortment Optimization with Customer-Centric Fairness Yihua He, Qian Liu and Jin Qi	A Tractable Approach for Contextual Distributionally Robust Optimization with Causal Optimal Transport Fenglin Zhang and Jie Wang	A Pricing Model for Enhancing the Utilization of Public Electric Vehicle Facilities: A Case Study in Hong Kong Jiangxue Han and Yong-Hong Kuo
On Tackling High-Dimensional Nonconvex Stochastic Optimization via Stochastic First-Order Methods with Non-smooth Proximal Terms and Variance Reduction Yue Xie	Joint Differential Privacy for Resource Allocation: Protection Capability, Price of Privacy, Algorithm Design, and Policy Implications Du Chen, Geoffrey A. Chua and Victor Richmond R. Jose	Screening with Fairness Consideration: A Unified Framework Zihao Xiang, Man-Chung Yue and Zhi Chen	Optimal Open-Loop Policy in Distributionally Robust Inventory Control with Uncertain Lead Times Xin Chen, Zhuoyu Long and Jingjun Men	Intelligent Dominance Dispatch of Emergency Ambulance Zikun Lin, Daniel Zhuoyu Long and Viet Anh Nguyen	Constructive and local search heuristics for a point-to-point airline network design problem Jinha Hibino, Shinji Imahori and Mihiro Sasaki
		Contracting with a Mechanism Designer Tian Bai, Yiding Feng, Yaohao Liu, Mengfan Ma and Mingyu Xiao	Fully Online Matching with General Stochastic Arrivals and Departures Zihao Li, Hao Wang and Zhenzhen Yan	An Iterative Sampling Approach for Solving Sinkhorn Distributionally Robust Optimization Jindong Jiang and Jie Wang	

IFORS Tutorial Lecture

Venue: V322, Time: 09:15 – 10:15

Mathematical Optimization and Operations Research for AI

Professor Yinyu YE
Professor Emeritus of Stanford and
Visiting Professors of SJTU, CUHKSZ, HKUST, and SIMIS

ABSTRACT

This talk aims to present several mathematical optimization and operations research (OR) problems, models, and algorithms for AI such as the LLM training, tunning and inferencing. In particular, we describe how classic OR models/theories, such as online resource allocation, may be applicable to accelerate and improve the Training/Tunning/Inferencing processes that are popularly used in LLMs. On the other hand, we show breakthroughs in classical Optimization (LP and SDP) Solver developments and their applications aided by AI-related techniques and their implementations on GPU.

BIO

Yinyu Ye, the 2009 John von Neumann Theory Prize recipient and formally the K.T. Li Professor of Stanford University, is now the Visiting Professor of Shanghai Jiao Tong University, Chinese University of Hong Kong at Shenzhen and Hong Kong University of Science and Technology. His current research topics include Continuous and Discrete Optimization, Data Science and Applications, Numerical Algorithm Design and Analyses, Algorithmic Game/Market Equilibrium, Operations Research and Management Science etc.. He was one of the pioneers on Interior-Point Methods, Conic Linear Programming, Distributionally Robust Optimization, Online Linear Programming and Learning, Algorithm Analyses for Reinforcement Learning\Markov Decision Process, nonconvex optimization, and etc. He and his students have received numerous scientific awards, himself including the 2006 INFORMS Farkas Prize (Inaugural Recipient) for fundamental contributions to optimization, the 2009 John von Neumann Theory Prize for fundamental sustained contributions to theory in Operations Research and the Management Sciences, the inaugural 2012 ISMP Tseng Lectureship Prize for outstanding contribution to continuous optimization (every three years), the 2014 Society of Industrial and Applied Mathematics Optimization Prize awarded (every three years), the 2025 Constantin Caratheodory Prize of the Global Optimization Congress, etc.

叶荫宇 (Yinyu Ye), 2009 约翰·冯·诺伊曼理论奖获得者及原斯坦福大学李国鼎讲席教授,现任上海交通大学, 香港中文大学深圳和香港科技大学访问讲席教授。他的主要研究方向为连续和离散优化,数据科学及应用,数字算法设计及分析,算法博弈及市场均衡,运筹及管理科学等; 他和其他科学家开创了内点优化算法,锥规划模型,分布式鲁棒优化,在线线性规划和学习,强化学习和马可夫过程及非凸优化算法分析等。他和他的学生多次获得科学奖项: 包括他自己的 2006 INFORMS Farkas Prize (首位获奖者) ,2009 年约翰·冯·洛伊曼理论奖,国际数学规划 2012 Tseng Lectureship Prize (首位获奖者每三年) ,2014 美国应用数学学会优化奖 (每三年), 2025 Constantin Caratheodory Prize of the 全局优化 Congress, 等。

Session: Optimization Theory and Algorithms I

Venue: V302, Time: 14:00 – 15:45

Session Chair: Lexiao Lai

A Two-Strategy Acceleration Framework for Phase Retrieval: FAM via Exact Minimization and FRWF via First-Order Approximation.

Linbin Li (Beihang University) and Yong Xia (Beihang University).

Abstract. Phase retrieval is a pivotal problem in optical imaging, signal processing, and communications, with Alternating Minimization (AM) and Reshaped Wirtinger Flow (RWF) being two core algorithm families for its solution. However, these two algorithms have limited convergence speed and stability under random initialization, which restricts their practical application. To further improve convergence speed and stability, this work proposes a novel two-strategy acceleration framework for phase retrieval, designed from an optimization perspective—distinct from conventional truncated methods. The framework integrates two improved algorithms: Fast Alternating Minimization (FAM) (optimized from AM) and Fast Reshaped Wirtinger Flow (FRWF) (refined from RWF). Extensive numerical experiments demonstrate that under random initialization, both new algorithms achieve significant improvements in convergence performance and stability. Furthermore, we present preliminary theoretical analysis results regarding convergence. This framework provides flexible solutions for phase retrieval tasks.

Keywords: Phase retrieval, Alternating Minimization, Reshaped Wirtinger Flow

Convex Facial Reduction Algorithm and Strong Extended Duals.

Ying Lin (Department of Data and Systems Engineering, The University of Hong Kong), Tianxiang Liu (Institute of Systems and Information Engineering, Tsukuba University) and Bruno Lourenço (The Institute of Statistical Mathematics).

Abstract. The facial reduction algorithm plays an important role in regularizing any conic convex programming problems and establishing the extended duals that enjoy strong duality without requiring any constraint qualifications. When the cone is nice, the extended dual has a relatively simpler formula. In this work, we generalize both the facial reduction algorithm and the concept of niceness from closed convex cones to convex sets. Our convex facial reduction algorithm can handle a convex problem with a nonlinear objective function and an intersection of two convex sets, and so is able to deal with problems that are inefficiently SDP representable. Based on the convex facial reduction algorithm we establish a class of extended duals that covers the extended duals in conic cases. Similarly, when the two sets are both nice, the formula of the extended dual is relatively simpler.

Keywords: Facial Reduction Algorithm, Strong Duality, Fenchel Dual, Niceness

On resolution of ℓ_1-norm minimization via a two-metric adaptive projection method.

Hanju Wu (University of Hong Kong) and Yue Xie (University of Hong Kong).

Abstract. In this work we propose an efficient two-metric adaptive projection method for solving the ℓ_1-*norm* regularization problem. It is inspired by the two-metric projection method, a simple yet elegant algorithm proposed by Bertsekas to address bound/box-constrained optimization problems. The algorithm's low per-iteration cost and potential for using Hessian information make it a favorable computation method for this problem class and our proposed method inherit these advantages. Previous attempt to extend two-metric projection method for solving ℓ_1 -*norm* minimization problem resorts to the intermediate reformulation as a bound-constrained problem and cause the issue of numerical instability, in sharp contrast to our proposed method. We demonstrate that our method is theoretically sound - it has global convergence. Furthermore, it is an active set method and capable of manifold identification – the potentially low-dimensional structure can be identified in finite number of steps and the method reduces to an unconstrained Newton method in the subspace. It also has superlinear convergence rate under the error bound condition and strict complementarity. Therefore, given sparsity in the solution, the method enjoys superfast convergence in iteration while maintaining scalability, making it desirable for large-scale problems. Numerical experiments are conducted to illustrate the advantages of this algorithm implied by the theory compared to other competitive methods, especially in large-scale scenarios.

Keywords: ℓ_1-*norm* minimization, Two-metric projection method, Manifold identification, Error bound

On the diameter of subgradient sequences in o-minimal structures.

Lexiao Lai (The University of Hong Kong) and Mingzhi Song (The University of Hong Kong).

Abstract. We study subgradient sequences of locally Lipschitz functions definable in a polynomially bounded o-minimal structure. We show that the diameter of sequences is related to the variation in function values, with error terms dominated by a double summation of step sizes. Consequently, we prove that bounded subgradient sequences converge if the step sizes are of order *1/k*. The proof uses Lipschitz *L*-regular stratifications in o-minimal structures to analyze subgradient sequences via their projections onto different strata.

Keywords: Subgradient sequence, o-minimal structures, stratification theory

Session: Advances in Data Science I

Venue: V312, Time: 14:00 – 15:45

Session Chair: Ruicheng Ao

Graph Topology Adaptive Judgment Against Node Label Noise.

Mengyao Zhou (Academy of Mathematics and Systems Science, Chinese Academy of Sciences), Xiao Han (Beihang University), Wei Wei (Beihang University) and Guiying Yan (Academy of Mathematics and Systems Science, Chinese Academy of Sciences).

Abstract. Graph Neural Networks (GNNs) have exhibited remarkable capabilities in processing graph data. Nevertheless, their performance is highly dependent on the labeled data, making them vulnerable to label noise. Existing methods often improve the robustness of node classification by adding trusted edges to the graph. However, most of them overlook the impact of potential noisy edges on the model's robustness, leading to a notable decline in performance as the average degree of the dataset increases. In this paper, we provide a theoretical explanation for the performance degradation observed in existing methods on datasets with high average degrees. Building on this insight, we propose the Graph Topology Adaptive(GTA) model, which incorporates EdgeBoost Module to add trusted edges based on node latent space similarity and EdgePrune Module to eliminate untrusted edges through optimized screening mechanism. Two modules work collaboratively to adaptively adjust the graph topology and generate final predictions. Both theoretical analysis and extensive experimental results validate the effectiveness of the GTA model.

Keywords: Graph Neural Networks, Label Noise, Robustness, Node Classification

Optimal pivot path of the simplex method for linear programming based on reinforcement learning.

Anqi Li (Nankai University), Tiande Guo (University of Chinese Academy of Sciences), Congying Han (University of Chinese Academy of Sciences), Bonan Li (University of Chinese Academy of Sciences) and Haoran Li (University of Chinese Academy of Sciences).

Abstract. Based on the existing pivot rules, the simplex method for linear programming is not polynomial in the worst case. Therefore, the optimal pivot of the simplex method is crucial. In this paper, we propose the optimal rule to find all the shortest pivot paths of the simplex method for linear programming problems based on Monte Carlo tree search. Specifically, we first propose the SimplexPseudoTree to transfer the simplex method into tree search mode while avoiding repeated basis variables. Secondly, we propose four reinforcement learning models with two actions and two rewards to make the Monte Carlo tree search suitable for the simplex method. Thirdly, we set a new action selection criterion to ameliorate the inaccurate evaluation in the initial exploration. It is proved that when the number of vertices in the feasible region is C_n^m, our method can generate all the shortest pivot paths, which is the polynomial of the number of variables. In addition, we experimentally validate that the proposed schedule can avoid unnecessary search and provide the optimal pivot path. Furthermore, this method can provide the best pivot labels for all kinds of supervised learning methods to solve linear programming problems.

Keywords: simplex method, linear programming, pivot rules, reinforcement learning
Topics: Machine Learning and Data Science

Purity Law for Neural Routing Problem Solvers with Enhanced Generalizability.

Wenzhao Liu (University of Chinese Academy of Sciences),
Haoran Li (University of Chinese Academy of Sciences),
Congying Han (University of Chinese Academy of Sciences),
Zicheng Zhang (University of Chinese Academy of Sciences),
Anqi Li (University of Chinese Academy of Sciences) and
Tiande Guo (University of Chinese Academy of Sciences).

Abstract. Achieving generalization in neural approaches across different scales and distributions remains a significant challenge for routing problems. A key obstacle is that neural networks often fail to learn robust principles for identifying universal patterns and deriving optimal solutions from diverse instances. In this paper, we first uncover Purity Law, a fundamental structural principle for optimal solutions of routing problems, defining that edge prevalence grows exponentially with the sparsity of surrounding vertices. Statistically validated across diverse instances, Purity Law reveals a consistent bias toward local sparsity in global optima. Building on this insight, we propose Purity Policy Optimization (PUPO), a novel training paradigm that explicitly aligns characteristics of neural solutions with Purity Law during the solution construction process to enhance generalization. Extensive experiments demonstrate that PUPO can be seamlessly integrated with popular neural solvers, significantly enhancing their generalization performance without incurring additional computational overhead during inference.

Keywords: Routing Problem, Generalization, Universal Structural Principle, Policy Optimization
Topics: Machine Learning and Data Science, Optimisation

Optimizing LLM Inference: Fluid-Guided Online Scheduling with Memory Constraints.

Ruicheng Ao (MIT).

Abstract. Large Language Models (LLMs) is indispensable in today's applications, but their inference procedure—generating responses by breaking text into smaller pieces and processing them using a memory-heavy element named Key-Value (KV) cache—requires a lot of computational resources, especially when memory is limited. This paper treats LLM inference optimization as a multi-stage online scheduling problem, where prompts arrive sequentially and the incremental expansion of the KV cache during inference renders conventional scheduling algorithms ineffective. To address this challenge, we develop a fluid dynamics approximation to establish a tractable benchmark, providing insights for devising effective scheduling algorithms. Building upon this foundation, we introduce the Waiting for Accumulated Inference Threshold (WAIT) algorithm as a warm-up. This method maintains multiple thresholds to determine the scheduling order of incoming prompts, optimizing resource utilization when output lengths are known at the time of arrival. In practical applications where output lengths are not known at the time of prompt arrival, we extend our method by introducing the Nested WAIT algorithm. This algorithm constructs a hierarchical framework comprising multiple segments, each defined by distinct thresholds, to effectively manage the random prompt arrivals with unknown output lengths. Theoretical analysis shows both algorithms have near-optimal performance compared with the fluid benchmark under heavy traffic limit, balancing throughput, latency, and Time to First Token (TTFT). Numerical experiments conducted with the Llama-7B model on an A100 GPU, utilizing both synthetic and real-world datasets, demonstrate that our approach achieves superior throughput and reduced latency compared to widely adopted baseline methods such as vLLM and Sarathi. This research bridges operations research and machine learning, presenting a theoretically grounded framework for the efficient deployment of large language models under memory constraints.

Keywords: Large Language Model, Key-value cache, Memory Constraint, Online scheduling

Session: Game Theory I

Venue: V321, Time: 14:00 – 15:45

Session Chair: Junyue Zhang

Beyond Regularity: Simple versus Optimal Mechanisms, Revisited.

Yiding Feng (Hong Kong University of Science and Technology) and Yaonan Jin (Huawei TCS Lab).

Abstract. A large proportion of the Bayesian mechanism design literature is restricted to the family of regular distributions [Mye81] or the family of monotone hazard rate (MHR) distributions [BMP63], which has overshadowed this rich and well-developed theory. We (re-)introduce two generalized families: quasi-regular distributions and quasi-MHR distributions. Likewise, the parameterized families of λ-regular (a.k.a. α-strongly regular) distributions [CR14,SS19], which smoothly interpolate regular distributions and MHR distributions, generalize to λ-quasi-regular distributions.

The significance of our new families is manifold. Firstly, their defining conditions are immediate "economic" relaxations of the original defining conditions (e.g., regularity as monotonicity of the virtual value functions), capturing key economic intuitions. Secondly, they satisfy natural mathematical properties (about order statistics) failed for the original families, thus technically more tractable. Thirdly, numerous results (by [BK96, HR09, CD15, DRY15, HR14, AHNPY19, JLTX20, JLQTX19, FLR19, GHZ19, JLX23, LM24] etc) known merely for the original families now can extend to our new families. Many of these extensions incur no quantitative loss, or even improve the state of the art for the original families. Finally, beyond the third point, our new families guide us to entirely new perspectives and thus entirely unknown results. For example, regarding revenue maximization for symmetric versus asymmetric regular buyers, we acquire 0.5- versus 0.1908-approximations for the (less-than-)one-sample prophet inequalities, respectively. To the best of our knowledge, such results are blank in the literature, despite their widely-studied welfare maximization counterparts [CDFS19, RWW20, CCES20, CDFSZ21, CCES24].

Keywords: Mechanism design, Revenue management, Game theory, Distribution robust optimization, Approximation

A Novel Sample-Based Best Response Dynamics for Solving Nash Equilibrium in Continuous Games.

Zhouming Wu (University of Chinese Academy of Sciences),
Xu Xia (State Key Lab of Mathematical Sciences, Academy of Mathematics and Systems Science) and
Yifen Mu (Institute of Systems Science, Academy of Mathematics and Systems Science, Chinese Academy of Sciences).

Abstract. Nash equilibrium (NE) computation is an important problem in systems of distributed artificial intelligence (DAI) with games, i.e, strategic interaction between agents. However, when the action space is continuous and payoff functions are not concave, the pure NE may not exist and there are few algorithms proposed to solve the games. In this paper, we consider NE computation in continuous games with one-dimensional action space for each player. We propose a sample-based best response dynamics (SBRD), in which the best response of each player is approximated via finite sample points. By redefining saddle and sink, as well as investigating the dynamical properties of SBRD, we theoretically prove that SBRD admits an unique solutions from almost all initial states and converges to some critical point. In SBRD, we do not assume any properties of payoff functions or accuracy of the best response. Surprisingly, even when NEs do not exist, SBRD can still find approximate NE. The approach also applies to scenarios with incomplete information, thereby eliminating the limitation of traditional dynamics that require precise payoff knowledge. Besides, this paper bridges DAI and dynamical systems theory by studying the piecewise-continuous dynamical systems induced by SBRD, advancing the application of dynamical system theory for learning dynamics in broader game classes.

Keywords: Learning in Games, Best Response Dynamics, Function Approximation, Convergence, Non-smooth Dynamical Systems

Mechanism Design for Exchange Markets.

Yukun Cheng (Jiangnan University).

Abstract. Exchange markets are a significant type of market economy, in which each agent holds a budget and certain (divisible) resources available for trading. Most research on equilibrium in exchange economies is based on an environment of completely free competition. However, the orderly operation of markets in reality also relies on effective economic regulatory mechanisms. This paper initiates the study of the mechanism design problem in exchange markets, exploring the potential to establish truthful market rules and mechanisms. This task poses a significant challenge as unlike auctioneers in auction design, the mechanism designer in exchange markets lacks centralized authority to fully control the allocation of resources.

In this paper, the mechanism design problem is formalized as a two-stage game. In stage 1, agents submit their private information to the manager, who then formulates market trading rules based on the submitted information. In stage 2, agents are free to engage in transactions within these rules, ultimately reaching an equilibrium. We generalize the concept of liquid welfare from classical budget-feasible auctions and use market liquid welfare as a measure to evaluate the performance of the designed mechanism. Moreover, an extra concept called profitability is introduced to assess whether the market is money-making (profitable) or money-losing (unprofitable). Our goal is to design a truthful mechanism that achieves an (approximate) optimal welfare while minimizing unprofitability as much as possible. Two mechanisms for the problem are proposed. The first one guarantees truthfulness and profitability while approaching an approximation ratio of 1/2 in large markets. The second one is also truthful and achieves 1/2 approximation in general markets but incurs bounded unprofitability. Our aim is for both mechanisms to provide valuable insights into the truthful market design problem.

Keywords: Exchange Market, Budget Constraints, Incentive Compatibility

Arbitrage on Decentralized Exchanges.

Yutian Zhou (The Chinese University of Hong Kong),
Xuedong He (The Chinese University of Hong Kong) and
Chen Yang (The Chinese University of Hong Kong).

Abstract. Decentralized exchanges (DEXs) are alternative venues to centralized exchanges (CEXs) for trading cryptocurrencies and have become increasingly popular. An arbitrage opportunity arises when the exchange rate of two cryptocurrencies in a DEX differs from that in a CEX. Arbitrageurs can then trade on the DEX and CEX to make a profit. Trading on the DEX incurs a gas fee, which determines the priority of the trade being executed. We study a gas-fee competition game between two arbitrageurs who maximize their expected profit from trading. We derive the unique symmetric mixed Nash equilibrium and find that (i) the arbitrageurs may choose not to trade when the arbitrage opportunity and liquidity is small; (ii) the probability of the arbitrageurs choosing a higher gas fee is lower; (iii) the arbitrageurs pay a higher gas fee and trade more when the arbitrage opportunity becomes larger and when liquidity becomes higher; (iv) the arbitrageurs' expected profit could increase with arbitrage opportunity and liquidity. The above findings are consistent with our empirical study.

Keywords: Nash equilibrium, Decentralized exchanges, Arbitrage

Last-iterate convergence in zero-sum games: Alternating heterogeneous dynamics to the rescue.

Xinxiang Guo (School of Mathematical Sciences, University of Chinese Academy of Sciences),
Junyue Zhang (School of Mathematical Sciences, University of Chinese Academy of Sciences),
Yifen Mu (Institute of Systems Science, Academy of Mathematics and Systems Science, Chinese Academy of Sciences),
Xiao Wang (Shanghai University of Finance and Economics) and
Ioannis Panageas (University of California, Irvine).

Abstract. Last-iterate convergence in zero-sum games has received extensive study with many seminal works (Daskalakis et al., 2017, Daskalakis and Panageas, 2018, Mertikopoulos et al., 2018a, Wei et al., 2020, Cai et al., 2022, Gorbunov et al., 2022a,b). Nevertheless, all the aforementioned works focus on simultaneous updates, where both players use the same dynamics. Unfortunately, the picture becomes elusive if the players update in an alternating way, using different dynamics, and little is known about provable convergence guarantees. In this paper, we investigate last-iterate convergence guarantees in repeated two-player zero-sum games under heterogeneous updating rules—that is, when the two players use different dynamics. More specifically, we focus on the setting where the row player employs the mirror descent (MD) algorithm, while the column player knows the strategy of the row player and responds with a best response (BR). We show that when the column player plays BR, the row player achieves best-iterate convergence to their minimax strategy. Second, when the column player adopts smoothed best response (SBR), the stage strategy of the row player exhibits last-iterate convergence. Finally, by using Tikhonov regularization, we prove that the joint strategy profile achieves $O(\frac{1}{\sqrt{T}})$ last-iterate convergence to Nash equilibrium (NE), matching the rate of optimistic and extra gradient from (Cai et al., 2022). We also complement our theoretical results via illustrative simulations. These experiments provide intuitive evidence that heterogeneous learning significantly changes the dynamics in repeated games and enhances the convergence to NE.

Keywords: Last-iterate convergence, Heterogeneous dynamics, Alternating updates, Zero-sum games

Session: Logistics and Supply Chains I

Venue: V303, Time: 14:00 – 15:45

Session Chair: Abhay Sobhanan

Dynamic optimization of differentiated pricing for crowdsourced delivery couriers with random utility theory.

Chunwei Yang (City University of Hong Kong) and Andy Chow (City University of Hong Kong).

Abstract. Crowdsourced delivery offers urban logistics the potential of reducing operating costs through using external resources, while enabling flexible workforce management and the identification of profitable demand. It has become a popular approach for delivering services with the increasingly available digital platforms and mobile apps. Nevertheless, crowdsourced delivery also induces uncertainties in the availability and willingness of couriers to take tasks. To address these challenges, this paper proposes a dynamic optimization framework for crowdsourced delivery platforms that jointly optimizes courier–order assignment and delivery fee under stochastic customer demand and uncertain courier behavior. The platform's decision-making process is formulated as a Markov Decision Process (MDP), where courier acceptance is modeled using a utility-based random choice model. The proposed framework integrates a reinforcement learning framework through two components: a planning model that solves a value-aware combinatorial optimization problem at each decision epoch by jointly selecting the optimal courier–order pairings and personalized delivery fees to maximize expected platform gains; and a learning model that estimates the long-term value of system states using a neural network trained via off-policy temporal difference (TD) learning. This hybrid structure enables forward-looking decision-making that balances immediate profit with future system performance.

We evaluate the proposed framework in a simulated crowdsourced delivery environment based on a 3 km × 3 km urban network in Shenzhen, China. Computational experiments show that the proposed method effectively adapts to spatial and temporal demand fluctuations, improving both platform profit and order fulfillment quality. Methodologically, this study contributes to intelligent logistics management by combining reinforcement learning with exact combinatorial optimization, thus avoiding the exploration explosion typically encountered in end-to-end RL over large combinatorial decision spaces. From a practical perspective, the results underscore the value of jointly optimizing matching and pricing decisions, and highlight how personalized incentives can serve as a powerful lever to boost courier participation, leading to improved platform efficiency and profit.

Keywords: Crowdsourced delivery, Task assignment, Dynamic pricing, Markov decision process, Reinforcement learning

A Hybrid AI-Driven Forecasting Framework for Hotel Occupancy: Combining Machine Learning, Managerial Insight, and Customer Feedback.

Tian Tian (Universiti Putra Malaysia), Aneng He (Universiti Putra Malaysia) and Lai Soon Lee (Universiti Putra Malaysia).

Abstract. In the digitally driven hospitality landscape, online travel agencies (OTAs) have transformed hotel marketing and distribution, yet their quantifiable impact on hotel performance remains ambiguous. This study investigates the relative influence of pricing strategies and online customer reviews on hotel occupancy by developing a novel artificial intelligence (AI) framework that prioritizes interpretability and decision support. Utilizing advanced text mining and sentiment analysis on OTAs review data, the model integrates active learning to optimize training efficiency and particle swarm optimization to enhance predictive accuracy. Importantly, SHapley Additive exPlanations (SHAP) and feature selection techniques are employed to ensure model transparency, enabling hotel managers to understand and trust the model's predictions. The resulting framework synthesizes customer sentiment, pricing history, and market dynamics into a multidimensional occupancy forecasting tool. Empirical evaluation using real hotel data demonstrates the model's high accuracy and practical utility in revenue management and strategic planning. By bridging AI innovation with interpretability, this research provides a rigorous, data-driven approach to demand forecasting and contributes to the broader operations research discourse in hospitality management.

Keywords: machine learning, revenue management, entropy weight, active learning, sentiment analysis

Fair Human–Robot Collaborative Logistics in Industry 5.0.

Yang Deng (Tsinghua University) and Andy Chow (City University of Hong Kong).

Abstract. In the era of Industry 5.0, integrating human-centric values with advanced automation is pivotal, yet most existing research still isolates machine decision-making from human work practices, which leaves a gap between theoretical experiments and practical deployment. We address a realistic distribution planning problem that requires collaboration between human distributors and autonomous trucks within a manufacturing setting, tightly coupled with upstream production decisions. We formulate the system as a partially observable Markov decision process (POMDP) and propose a multi-agent reinforcement learning (MARL) framework that unifies production control with distributor planning. At the upper level, two agents jointly determine production quantities for pre-ordered and just-in-time products while selecting distribution channels under volatile market signals and uncertain retail capabilities. At the lower level, a decentralized execution policy coordinates human and robotic distributors to fulfill dynamic customer demand, accounting for travel times, safety constraints, and vehicle capacities. A communication-enabled centralized training method with differentiable inter-agent messaging (recurrent encoders for belief state updates) aligns local actions with global objectives and ensures that automation augments rather than undermines the human workforce, which is the core of Industry 5.0. Methodologically, we (i) introduce a hierarchical structure with POMDP that captures partial observability in retail service levels, distributor availability, and market shocks; (ii) design a communication-based MARL algorithm with a shared critic and message gating to fuse production, channel selection, and routing decisions; (iii) embed human-centric objectives—fair workload and stability of assignments—via soft constraints and shaped rewards (e.g., exponential penalties for inventory overflow). Computational experiments on a steel-manufacturing case study in mainland China show that the proposed framework outperforms other benchmarks and conventional RL baselines without communication. We observe higher profit, lower stockouts, lower average tardiness, and markedly reduced assignment churn, while maintaining equitable workload across human distributors. The framework is model-free and deployable with standard telemetry, offering a practical way for Industry 5.0 applications for production–distribution co-optimization.

Keywords: Partially observable Markov decision process, Multi-agent reinforcement learning, Inventory control, Industry 5.0

Dynamic joint optimization of autonomous vehicle-based mobile warehousing and crowdsourced delivery routing with the use of reinforcement learning.

Zeyu Zhang (CityU), Andy Chow (City University of Hong Kong) and Chunwei Yang (City University of Hong Kong).

Abstract. This study proposes an innovative urban logistics framework that integrates autonomous vehicles (AVs) as mobile replenishment warehouses with crowdsourced delivery riders. A dual-layer stochastic optimization problem is formulated to address the coordinated operation of mobile inventory management through AVs and last-mile distribution by crowdsourced riders under dynamic urban conditions. The problem is structured as a Markov Decision Process (MDP) incorporating a hierarchical decision-making structure to simultaneously optimize AV routing for inventory replenishment and rider dispatching for delivery operations. To handle the stochasticity from real-time order arrivals and spatiotemporal delivery constraints, we develop a model-based stochastic optimization framework integrating Deep Q-Network (DQN) reinforcement learning. Numerical experiments demonstrate that the hybrid approach achieves higher system efficiency compared to the conventional method.

Keywords: Markov decision process, Reinforcement learning, Urban logistics, Vehicle routing problem, Mobile warehousing

A Hybrid Genetic Algorithm for the Rural Postman Problem with Multiple Trucks and Drones.

Abhay Sobhanan (Indian Institute of Management Bangalore),
Hadi Charkhgard (University of South Florida) and Changhyun Kwon (KAIST).

Abstract. Arc Routing Problems are fundamental to logistics and field operations, with applications spanning infrastructure inspection, waste collection, and urban patrolling. A prominent variant, the Rural Postman Problem (RPP), focuses on servicing a designated subset of network edges. Although recent advances in collaborative truck–drone logistics have shown promising results in node-routing problems, their extension to arc-routing contexts remains largely unexplored despite the clear practical relevance. This paper addresses this gap by introducing the Rural Postman Problem with multiple Trucks and Drones (RPP-mTD), a novel and challenging generalization aimed at minimizing the overall operational makespan. The RPP-mTD poses significant combinatorial complexity due to the simultaneous coordination of multiple truck–drone fleets, the allowance for drones to service multiple arcs per trip, and the need to manage flexible rendezvous strategies between vehicles. To tackle this NP-hard problem, we develop a Hybrid Genetic Algorithm (HGA) explicitly designed for multi-fleet coordination. The HGA incorporates a two-part chromosome structure that decouples arc sequencing from vehicle assignment, a segment-preserving crossover operator that protects the integrity of truck routes, and a suite of local search heuristics to enhance intensification within the solution space. Computational experiments demonstrate both the competitiveness and scalability of the approach. On benchmark single-truck, single-drone instances, the HGA achieves an average optimality gap of only 2.25% while running substantially faster than the state-of-the-art heuristic. On large-scale instances involving up to 500 nodes and 100 required arcs, the method continues to deliver robust performance. From a managerial perspective, the results reveal that increasing the number of trucks yields significant reductions in makespan, whereas adding more drones provides diminishing improvements; extending drone endurance proves highly advantageous but computationally demanding; and relaxing rendezvous constraints offers consistent yet gradually saturating benefits. The proposed HGA establishes itself as an effective and scalable decision-support tool for mixed-fleet logistics in arc-routing applications, bridging a key gap between theoretical problem models and practical operational needs.

Keywords: Arc routing problem, Hybrid genetic algorithm, Mixed-fleet logistics, Truck-drone routing

Session: Stochastic Models

Venue: V315, Time: 14:00 – 15:45

Session Chair: Wei Liu

Exact Analytical Option Pricing Formula Under a Jump-Diffusion Model with Hawkes Intensity.

Daniel Wei-Chung Miao (National Taiwan University of Science and Technology), Xenos Chang-Shuo Lin (National Taiwan University of Science and Technology) and Ruey-An Shiu (National Taiwan University).

Abstract. In this paper, we extend the jump-diffusion model by allowing the jump intensity to follow a Hawkes process. The self-exciting property of the Hawkes process captures the serial correlation of successive jump events, reflecting the empirical observation that jumps tend to occur in clusters during periods of market instability. While this specification is more realistic, analyzing the distribution of jump counts presents technical challenges, which in turn complicates the derivation of option prices under the proposed framework. We address this issue by deriving exact analytical expressions for the full probability distribution of the number of jumps over a given time horizon. These results yield a recursive procedure for efficiently computing the jump-count distribution, which can then be applied directly to the pricing of European options. As the results are exact and fully analytical, numerical implementation is both accurate and computationally efficient. Through numerical experiments, we demonstrate that the recursive method produces results consistent with Monte Carlo simulations. Furthermore, using our pricing formula, we investigate how Hawkes process parameters affect European option prices and implied volatilities.

Keywords: Jump-diffusion model, Hawkes process, Self-exciting, European option, Implied volatility

Provable Policy Gradient for Robust Average-Reward MDPs Beyond Rectangularity.

Qiuhao Wang (Research Institute for Digital Economy and Interdisciplinary Sciences, Southwestern University of Finance and Economics),
Yuqi Zha (Department of Data Science, City University of Hong Kong),
Chin Pang Ho (Department of Data Science, City University of Hong Kong) and
Marek Petrik (Department of Computer Science, University of New Hampshire).

Abstract. Robust Markov Decision Processes (MDPs) offer a promising framework for computing reliable policies under model uncertainty. While policy gradient methods have gained increasing popularity in robust discounted MDPs, their application to the average-reward criterion remains largely unexplored. This paper proposes a Robust Projected Policy Gradient (RP2G), the first generic policy gradient method for robust average-reward MDPs (RAMDPs) that is applicable beyond the typical rectangularity assumption on transition ambiguity. In contrast to existing robust policy gradient algorithms, RP2G incorporates an adaptive decreasing tolerance mechanism for efficient policy updates at each iteration. We also present a comprehensive convergence analysis of RP2G for solving ergodic tabular RAMDPs. Furthermore, we establish the first study of the inner worst-case transition evaluation problem in RAMDPs, proposing two gradient-based algorithms tailored for rectangular and general ambiguity sets, each with provable convergence guarantees. Numerical experiments confirm the global convergence of our new algorithm and demonstrate its superior performance.

Keywords: Markov Decision Processes, Average-Reward Markov Decision Processes, Policy Gradient, Robust Optimization

Multi-Objective Optimization Model for Reserve and Credit Fund Allocation in Banking under Conditional Value-at-Risk Constraints.

Moch Panji Agung Saputra (Doctoral Student of Mathematics Study Program, Faculty of Mathematics and Natural Science, Universitas Padjadjaran),
Diah Chaerani (Department of Mathematics, Faculty of Mathematics and Natural Sciences, Universitas Padjadjaran),
Sukono (Department of Mathematics, Faculty of Mathematics and Natural Sciences, Universitas Padjadjaran)
and Mazlynda Md Yusuf (Faculty of Science and Technology, Universiti Sains Islam Malaysia (USIM)).

Abstract. In the realm of financial management, optimizing the allocation of funds in banking companies is vital to their operational efficiency. Banks manage their funds by allocating them into reserve and credit funds as their main financial activities. Optimizing these allocations ensures that all assets are effectively utilized. This study develops a multi-objective optimization model based on linear programming with a lexicographic approach for the allocation of reserve and credit funds in banking. The model integrates two sequential objectives: first, minimizing reserve funds to avoid excessive non-productive assets, and second, maximizing returns from credit funds to enhance profitability. The lexicographic order is determined by regulatory requirements from the financial services authority, which mandate that banks must meet adequate reserve fund provisions before being allowed to operate. Therefore, reserve fund minimization is addressed as the primary objective, followed by credit fund allocation optimization to achieve maximum returns. To address the risk dimension of banking operations, Conditional Value-at-Risk (CVaR) is applied as a key constraint, ensuring that potential losses from financial and credit risks remain within acceptable limits. The resulting optimization model produces efficient allocation strategies for reserve and credit funds, providing banks with a practical tool to achieve regulatory compliance while improving financial resilience and profitability.

Keywords: Multi-objective optimization, Lexicographic, Conditional Value-at-Risk, Banking allocation, Reserve funds, Credit funds

A Spider-Type Stochastic Subgradient Method for Expectation-Constrained Nonconvex Nonsmooth Optimization.

Wei Liu (Hong Kong PolyU) and Yangyang Xu (RPI).

Abstract. Many real-world problems, such as those with fairness constraints, involve complex expectation constraints and large datasets, necessitating the design of efficient stochastic methods to solve them. Most existing research focuses on cases, which have no constraint, or have constraints that admit an easy projection, or only have deterministic constraints. In this paper, we consider weakly convex stochastic optimization problems with expectation constraints, for which we build an exact penalty model. We first show the relationship between the penalty model and the original problem. Then on solving the penalty problem, we present a SPIDER-type stochastic subgradient method, which utilizes the subgradients of both the objective and constraint functions, as well as the constraint function value at each iteration. Under the Slater-type constraint qualification (CQ), we establish an iteration complexity result of $O(\epsilon^{-4})$ to reach a near-ϵ stationary point of the penalized problem in expectation, matching the lower bound for such tasks. Building on the exact penalization, an (ϵ, ϵ)-KKT point of the original problem is obtained. For a few scenarios, our complexity of either the objective sample subgradient or the constraint sample function values can be lower than the state-of-the-art results in [10, 27, 46] by a factor of ϵ^{-2}. Moreover, on solving two fairness-constrained problems and a multi-class Neyman-Pearson classification problem, our method is significantly (up to 466 times in terms of data pass) faster than the state-of-the-art algorithms.

Keywords: stochastic, subgradient, expectation constraints, weakly convex, fairness-constrained

Session: Strategies and Control Techniques

Venue: V316, Time: 14:00 – 15:45

Session Chair: Jelena Stanković

Open pit mine planning using optimal control techniques.

Bowen Zhao (Curtin University), Honglei Xu (Curtin University) and Lay Teo Kok (Sunway University).

Abstract. Traditional mathematical models for open pit mine planning with the purpose of maximizing the system's economic benefits are constructed as integer programming or mixed integer programming problem based on the economic block model. These systems are large and complex, making these approaches difficult to balance both complexity and dynamics. Optimal control technology studies the optimality of dynamic systems from an optimization perspective.
In this presentation, we reformulate the open pit mine planning problem within the framework of optimal control technology and propose new solutions to the ultimate pit limit problem and the strategic planning problem of open pit mines. Furthermore, we provide preliminary consideration and discussion of scenarios involving complex geological conditions.

Keywords: Open pit mine planning, Optimal control, Nonlinear programming, Gradient-type method

Project Monitoring and Control Systems.

Nicholas Hall (Fisher College of Business, The Ohio State University),
Zhuoyu Long (Department of Systems Engineering and Engineering Management, Chinese University of Hong Kong),
Jin Qi (Department of Industrial Engineering and Decision Analytics, Hong Kong University of Science and Technology) and
Yuhao Yan (Department of Industrial Engineering and Decision Analytics, Hong Kong University of Science and Technology).

Abstract. The global use of professional project management to deliver innovative products and services has not eliminated widely observed problems in delivering projects within schedule and within budget. Hence, accurate monitoring of project progress and cost is essential, both for communication with stakeholders and for project control through expediting decisions. Two systems for the monitoring and control of projects, Earned Value Management and Earned Schedule, are available. We model these systems and develop algorithms to minimize the expected total cost of a project under task time and cost uncertainty. This enables us to make structural comparisons between the optimal expediting decisions and expected costs of the two systems. This in turn reveals that the time profile of project workload is an important factor in the relative costs of the two systems. An extensive computational study validates these theoretical results, and shows that the expected loss from choosing the wrong monitoring and control system is typically significant. Our work provides theoretical support for the empirical performance of the two systems. We thus provide project managers with a theoretically justified, simple, and effective way to choose the best monitoring and control system at the start of a project, based on project characteristics that are generally known.

Keywords: project management, monitoring and control systems, expected makespan and total cost, expediting, earned value system, earned schedule

Proficiency Boosts Efficiency: Extended Johnson's Rules for Flow Shop Scheduling with Learning and Forgetting Effects.

Fenglin Zhang (The Chinese University of Hong Kong (Shenzhen) / Beijing Institute of Technology), Yuli Zhang (Beijing Institute of Technology) and Zuo-Jun Max Shen (The University of Hong Kong).

Abstract. Motivated by operations at a large-scale personal computer manufacturing factory assembly line, we investigate a flow shop scheduling problem with category-dependent learning and forgetting effects (CLFEs), where workers' proficiency affects processing efficiency. First, we propose a general CLFEs model to characterize the monotonicity between proficiency and efficiency, including existing specific learning curve formulations as special cases. Second, for the two-machine scheduling problem, we propose a Batch-wise Johnson's (BJ) rule to construct an optimal batch-wise schedule (i.e., a schedule where jobs of the same category are processed consecutively), which extends the well-known Johnson's rule. Next, we establish optimal scheduling rules for three specific two-machine scheduling scenarios: for problems with a strong forgetting effect, we prove the global optimality of the BJ rule; for problems without a forgetting effect, we propose an optimal Divide-and-Conquer Johnson's rule; for general two-machine scheduling problems, we provide sufficient conditions for the optimality of the BJ rule. Finally, for general multi-machine scheduling problems, we develop a Multi-machine Batch-wise Johnson's (MBJ) algorithm based on the principles of the BJ rule. Numerical studies show that the proposed BJ rule for the two-machine scheduling problem is nearly optimal with an average optimality gap of 0.1%, and the MBJ algorithm consistently finds optimal batch-wise schedules for large-scale multi-machine instances. Furthermore, our computational results reveal that the average optimal makespan gap between problems with and without CLFEs is approximately 10%, highlighting the practical importance of considering CLFEs in production planning.

Keywords: Permutation flow shop scheduling, Learning and forgetting effects, Extended Johnson's rules, Optimality analysis, Approximation algorithms

Concentration or diversification? The optimal strategy in R&D project exploration.

Qian Liu (HKUST), Tong Lu (HKUST) and Jin Qi (HKUST).

Abstract. We consider an R&D project that can be completed through multiple directions but no one knows which directions are feasible. In this paper, we build a model to characterize the organization's optimal exploration process under two policies: concentration and diversification. Success in any direction leads to project success and the organization gains rewards. Exploration costs are incurred during exploration, and the organization can choose when to stop. Our goal is to analyze how the organization should choose between the two policies to achieve a higher expected payoff. We find that, in a homogeneous setting, diversification is more efficient, but concentration is better when the project is easier to succeed. We also find that in a heterogeneous setting, diversification performs worse as the differences between directions increase. We also extend our model to cases with correlation among different directions and under budget constraints.

Keywords: resource allocation, project exploration, dynamic programming, Bayesian updating

National Innovation System Efficiency Assessed by DEA: Lessons for Policy and Growth.

Jelena Stanković (University of Niš, Faculty of Economics),
Ivana Marjanović (Faculty of Economics, University of Niš) and
Marina Stanojevic (University of Nis, Faculty of Economics).

Abstract. Innovation has become a key driver of long-term economic growth, making the efficiency of national innovation systems an essential subject of analysis. This paper applies Data Envelopment Analysis (DEA) as a multi-criteria method to evaluate the efficiency of knowledge creation, commercialization, and the overall innovation process in both developed and developing economies. Focusing on the most recent available year, the study highlights significant differences in efficiency levels across countries. Developed economies demonstrate a higher ability to transform research and development inputs into valuable innovation outputs, while many developing economies show persistent structural weaknesses. Nonetheless, several countries stand out as exceptions, achieving innovation efficiency comparable to advanced economies. The results confirm DEA's suitability for benchmarking national innovation systems and emphasize its potential to guide policymakers in designing strategies that improve innovation capacity, strengthen resource allocation, and stimulate sustainable economic growth.

Keywords: national innovation system, Data Envelopment Analysis, innovation efficiency, economic growth

Session: Optimization Theory and Algorithms II

Venue: V302, Time: 16:00 – 17:45

Session Chair: Yue Xie

LLM Serving Optimization with Variable Prefill and Decode Lengths.

Meixuan Wang (Tsinghua University), Yinyu Ye (HKUST, IEDA) and Zijie Zhou (HKUST, IEDA).

Abstract. We study the problem of serving LLM (Large Language Model) requests where each request has heterogeneous prefill and decode lengths. In LLM serving, the prefill length corresponds to the input prompt length, which determines the initial memory usage in the KV cache. The decode length refers to the number of output tokens generated sequentially, with each additional token increasing the KV cache memory usage by one unit. Given a set of n requests, our goal is to schedule and process them to minimize the total completion time. We show that this problem is NP-hard due to the interplay of batching, placement constraints, precedence relationships, and linearly increasing memory usage. We then analyze commonly used scheduling strategies in practice, such as First-Come-First-Serve (FCFS) and Shortest-First (SF), and prove that their competitive ratios scale up sublinearly with the memory limit—a significant drawback in real-world settings where memory demand is large. To address this, we propose a novel algorithm based on a new selection metric that efficiently forms batches over time. We prove that this algorithm achieves a constant competitive ratio. Finally, we develop and evaluate a few algorithm variants inspired by this approach, including dynamic programming variants, local search methods, and an LP-based scheduler, demonstrating through comprehensive simulations that they outperform standard baselines while maintaining computational efficiency.

Keywords: LLM Inference, Online Scheduling, Approximation Algorithms

Adaptively Robust LLM Inference Optimization under Prediction Uncertainty.

Zixi Chen (School of Mathematical Sciences, Peking University),
Yinyu Ye (Department of Management Science and Engineering, Stanford University) and
Zijie Zhou (Department of Industrial Engineering and Decision Analytics, HKUST).

Abstract. We study the problem of optimizing Large Language Model (LLM) inference scheduling to minimize total latency. LLM inference is an online and multi-task service process and also heavily energy consuming by which a pre-trained LLM processes input requests and generates output tokens sequentially. Therefore, it is vital to improve its scheduling efficiency and reduce the power consumption while a great amount of prompt requests are arriving. A key challenge in LLM inference scheduling is that while the prompt length is known upon arrival, the output length, which critically impacts memory usage and processing time, is unknown. To address this uncertainty, we propose algorithms that leverage machine learning to predict output lengths, assuming the prediction provides an interval classification (min-max range) for each request.

We first design a conservative algorithm, $\mathcal{A}_{max}$, which schedules requests based on the upper bound of predicted output lengths to prevent memory overflow. However, this approach is overly conservative: as prediction accuracy decreases, performance degrades significantly due to potential overestimation. To overcome this limitation, we propose $\mathcal{A}_{min}$, an adaptive algorithm that initially treats the predicted lower bound as the output length and dynamically refines this estimate during inferencing. We prove that $\mathcal{A}_{min}$ achieves a log-scale competitive ratio. Through numerical simulations, we demonstrate that $\mathcal{A}_{min}$often performs nearly as well as the hindsight scheduler, highlighting both its efficiency and robustness in practical scenarios. Moreover, $\mathcal{A}_{min}$ relies solely on the lower bound of the prediction interval—an advantageous design choice since upper bounds on output length are typically more challenging to predict accurately.

Keywords: Robust and Online Scheduling, LLM Inference, Operations Management in AI with Prediction

Spectral method for general Group Synchronization.

Ziyue Zhao (The University of Hong Kong), Huikang Liu (Shanghai Jiao Tong University)
and Man-Chung Yue (The University of Hong Kong).

Abstract. Synchronization problems over group $\tilde{\mathcal{G}}$aims to recover group elements $G_i^* \in \tilde{\mathcal{G}}$ for $1 \le i \le n$ from some pairwise noise measurements $G_i^* \circ \tilde{\mathcal{G}}_j^{*-1}$. The problem has various applications in signal processing, machine learning, and robotics. The common approach via its least square estimator results in a nonconvex and nonlinear constrained optimization problem. In this paper, we consider the group synchronization problem over the general subgroup in the Euclidean group. First, we propose a novel spectral method for synchronization over the general subgroup, which encompasses any subgroup of the orthogonal group and the special Euclidean group. Second, we provide performance guarantees for our general group synchronization problem. We conduct numerical experiments on synthetic data and the multiple-point set registration problem, demonstrating the superiority of our spectral method over state-of-the-art spectral-type methods.

Keywords: Group synchronization, Spectral method, Special Euclidean group, Estimation Error bound

On Tackling High-Dimensional Nonconvex Stochastic Optimization via Stochastic First-Order Methods with Non-smooth Proximal Terms and Variance Reduction.

Yue Xie (University of Hong Kong).

Abstract. When the nonconvex problem is complicated by stochasticity, the sample complexity of stochastic first-order methods may depend linearly on the problem dimension, which is undesirable for large-scale problems. To alleviate this linear dependence, we adopt non-Euclidean settings and propose two choices of non-smooth proximal terms when taking the stochastic gradient steps. This approach leads to stronger convergence metric, incremental computational overhead, and potentially dimension-insensitive sample complexity. We also consider further acceleration through variance reduction which achieves near optimal sample complexity and, to our best knowledge, is the first such result in the ℓ_1/ℓ_∞ setting. Since the use of non-smooth proximal terms is unconventional, the convergence analysis deviates much from algorithms in Euclidean settings or employing Bregman divergence, providing tools for analyzing other non-Euclidean choices of distance functions. Efficient resolution of the subproblems in various scenarios is also discussed and simulated. We illustrate the dimension-insensitive property of the proposed methods via preliminary numerical experiments.

Keywords: Nonconvex Large-scale Optimization, Stochastic First-order Methods, Non-Euclidean distances, Sample Complexity, Variance Reduction

Session: Advances in Data Science II

Venue: V312, Time: 16:00 – 17:45

Session Chair: Du Chen

A Logic-Mined, PSO-Tuned Multi-Architecture Ensemble Neural Network Framework for Student Mathematics Learning Potential Prediction in MET-Supported Classrooms.

Aneng He (Universiti Putra Malaysia), Lai Soon Lee (Universiti Putra Malaysia) and Tian Tian (Universiti Putra Malaysia).

Abstract. To address the gap in predictive learning analytics regarding the estimation of individual students' potential in specific learning environments, this study proposes a technically rigorous framework for predicting mathematical learning potential across schools, grades, and classes under diverse Mathematics Educational Technology (MET) settings. A dataset of 3,108 observations from 444 students in three schools was collected at multiple time points. Based on Bloom's Taxonomy, students' test scores, classroom performance, and homework were categorized into academic performance attributes (APA) and non-academic performance attributes (NAPA). The learning environment (LE) was defined by the Feldman–Silverman model describing mathematics classroom types in conjunction with recorded MET usage. From these dimensions, an academic attribute matrix and a teacher expectation matrix were constructed and integrated into a student attribute matrix incorporating expert judgment. This matrix follows a normal distribution, enabling cumulative recording of student state attributes over time. Guided by educational experts and logic mining, three domain-specific logical rules were embedded into an Ensemble Neural Network (ENN) comprising ANN, BPNN, CNN, DNN, DQN, and GAN models, with hyperparameters optimized via Particle Swarm Optimization (PSO). The ENN partitions the data based on the LE, grouping similar or identical categories into sub-datasets, within which training and test sets are determined using a time-window strategy. For each sub-dataset, the best-performing model is selected from the sub-training set and evaluated on its corresponding sub-test set. Model performance was assessed using Accuracy, F1-score, AUC, and Brier Score. Results indicate that the ENN with embedded logic rules achieved statistically significant improvements over predictions by three senior teachers in MET contexts, outperforming them across all evaluation metrics. These findings demonstrate that integrating multi-architecture ENNs, PSO-based hyperparameter tuning, and expert-informed logic mining yields a robust, interpretable, and highly accurate approach to predicting individual learning potential in MET-supported educational environments.

Keywords: Predictive Learning Analytics, Learning Attribute Modeling, Expert-guided Rules, Logic Mining, Ensemble Neural Network

Pure Exploration via Frank-Wolfe Self-Play.

Xinyu Liu (The Hong Kong University of Science and Technology), Chao Qin (Stanford University) and Wei You (The Hong Kong University of Science and Technology).

Abstract. We propose a unified approach to pure exploration in structured stochastic multi-armed bandits. Structure introduces sharp pathologies that complicate algorithm design and analysis: our linear-bandit case study exhibits non-unique optima, optimal designs with zero mass on the best arm, bilinear corners, and boundary nonsmoothness. We address these challenges through a game-theoretic reformulation. For a broad class of tasks, asymptotic analyses often reduce to a maximin optimization that admits a two-player zero-sum game interpretation between an experimenter and a skeptic: an experimenter allocates measurements to rule out alternatives while a skeptic proposes them. We propose to reformulate the game by allowing the skeptic adopt mixed strategy, leading to a concave--convex saddle-point problem. This viewpoint yields Frank–Wolfe Self-Play (FWSP): a projection-free, regularization-free, tuning-free method whose one-hot updates on both sides match the bandit sampling paradigm.

We prove convergence of the game value under FWSP. Our analysis proceeds via a continuous-time limit: a differential inclusion with a Lyapunov function that decays exponentially, implying a vanishing duality gap and convergence to the saddle value. We then embed the discrete updates into a perturbed flow and show that the discrete game value converges as well. Building on FWSP, we also propose a learning algorithm based on posterior-sampling. Numerical experiments demonstrates value convergence with a vanishing duality gap.

Keywords: pure exploration, zero-sum games, mixed-strategy equilibrium, Frank-Wolfe, differential inclusion

Distributionally Robust Federated Learning: An ADMM Algorithm.

Wen Bai (The Chinese University of Hong Kong), Yi Wong (City University of Hong Kong), Ling Dai (City University of Hong Kong), Xiao Qiao (City University of Hong Kong) and Chin Pang Ho (City University of Hong Kong).

Abstract. Federated learning (FL) aims to train machine learning (ML) models collaboratively using decentralized data, bypassing the need for centralized data aggregation. Standard FL models often assume that all data come from the same unknown distribution. However, in practical situations, decentralized data frequently exhibit heterogeneity. We propose a novel FL model, Distributionally Robust Federated Learning (DRFL), that applies distributionally robust optimization to overcome the challenges posed by data heterogeneity and distributional ambiguity. We derive a tractable reformulation for DRFL and develop a novel solution method based on the alternating direction method of multipliers (ADMM) algorithm to solve this problem. Our experimental results demonstrate that DRFL outperforms standard FL models under data heterogeneity and ambiguity.

Keywords: Federated Learning, Distributionally Robust Optimization, First Order Method

Joint Differential Privacy for Resource Allocation: Protection Capability, Price of Privacy, Algorithm Design, and Policy Implications.

Du Chen (Nanyang Technological University Singapore),
Geoffrey A. Chua (Nanyang Technological University Singapore) and
Victor Richmond R. Jose (Georgetown University).

Abstract. Resource allocation is a fundamental problem in management science. Recent growth in privacy concerns has made this problem more challenging because private data might be inadvertently revealed through the optimal allocation decisions. While differential privacy (DP) limits data leakage in resource allocation, it overprotects individual data necessary for targeted allocation, thus compromising decision utility substantially. To address this, we adopt joint differential privacy (JDP), another data privacy standard, and provide a new comprehensive understanding of this concept. We formally show that JDP provides privacy protection as strong as DP if recipients do not disclose their allocations. Moreover, JDP achieves this with substantially lower utility loss, with per-recipient loss vanishing as the system scales. However, the direct application of existing JDP algorithms may not reap the full benefits of lower utility loss and could empirically underperform compared to DP. To address this, we exploit the primal-dual relationship and develop the Noisy Dual Mirror Descent (NDMD) algorithm, a minimax near-optimal approach for a large class of convex resource allocation problems. NDMD offers a practical solution for privacy-preserving allocation problems, including the long-debated public policy problem of allocating Title I US education grants with privacy. Our simulations using real-world data show that shifting from DP to JDP reduces nationwide misallocation by as much as 52% and significantly mitigates the disproportionate impact of privacy protection on historically disadvantaged minority groups. Our results highlight the importance of operational knowledge such as disclosure discretion, downstream operational tasks, and algorithm design in balancing privacy with utility

Keywords: joint differential privacy, offline resource allocation, privacy-utility trade-off, public policy, Title I grant allocation

Session: Game Theory II

Venue: V321, Time: 16:00 – 17:45

Session Chair: Yiding Feng

Screening or Pooling?
Dual-channel Competition in Healthcare under Information Asymmetry.

Fangyuan Cao (The University of Hong Kong),
Jie Wang (Southeast University School of Economics and Management),
Feng Tian (The University of Hong Kong) and Yong-Hong Kuo (The University of Hong Kong).

Abstract. Healthcare platforms, such as HealthTap and Dingxiang Doctor, have gained popularity in recent years due to their convenience and flexibility, competing with in-person hospitals within the market for medical consultations. Although platforms can verify the qualifications of physicians, physicians' actual performance remains unknown to the public. In this research, we consider that, to strive for excellence, the online healthcare provider offers physicians two types of contracts for the service provided: a screening contract and a pooling contract. We develop a sequential game theory model to (i) analyze the optimal contract type to physicians on an online revenue-focused healthcare provider where physician efficiency is private information and (ii) study its effects on healthcare providers and patients in dual-channel competition. We find that, interestingly, the pooling contract may be more beneficial for an online healthcare provider when the failure cost is low, the offline hassle cost is sufficiently high, or it is more likely for a patient to encounter a low-type physician. Our contributions lie in the modeling framework and insights derived for the situation faced by online healthcare providers in managing information asymmetry about physicians, while also considering competition from traditional offline healthcare providers.

Keywords: telemedicine, healthcare competition, information asymmetry, contract design, healthcare operations management

An Evolutionary Game-Theoretic Analysis of Incentive Mechanisms for Cybersecurity in Blockchain-Enabled Smart Home Supply Chains.

Anu Rao (School of Management and Entrepreneurship (SME), IIT Jodhpur),
Amit Singh (Amrut Mody School of Management, Ahmedabad University) and
Devi Prasad Dash (School of Management and Entrepreneurship (SME), IIT Jodhpur).

Abstract. In the digital era, the smart home market faces significant cybersecurity inefficiencies, resulting in increased risks of data breaches, financial losses, and negative impacts on consumer trust, which in turn affect its smart home supply chain stakeholders. This research proposes that these challenges can be addressed through strategic collaboration among stakeholders, who can share information and resources to create a more secure ecosystem. We develop a cooperative framework where a platform operator (P), service providers (SPs), and end-users (U) collaborate through a blockchain-enabled platform to achieve a greater network-level security resilience and deliver trustworthy information flow. This framework is based on an evolutionary game model where the platform offers incentives to motivate SPs and users to participate in information sharing. Each boundedly rational player independently decides actions to maximize their utility under information asymmetry. This blockchain-enabled cybersecurity framework formalizes the interplay between the platform's incentive strategy, the SP's information-sharing strategy, and the user's trust strategy. Further, we employ a replicator dynamics model to simulate strategic evolution and derive the evolutionary stable strategies (ESS) to obtain the optimal operational decisions. Next, a stability analysis is conducted to identify the conditions required for a beneficial and secure equilibrium.

Further, we compare the stability conditions of our collaborative framework with a non-collaborative baseline where no incentives are offered. The analysis reveals that under a baseline model, the players bear the full cost of their security measures and risk exposure, which is a significant barrier for achieving systemic security in real-life settings. In contrast, without compromising each player's core information, our framework offers a more analytically tractable structure for decision makers, leading to a higher network-level security resilience for the entire smart home supply chain.

Keywords: Smart Home Supply Chain, Evolutionary Game Theory, Information Sharing, CyberRisk Management

Robust Multi-Object Auction Design under Knightian Preferences.

Xiaojin Fu (Noah's Ark Lab, Huawei, Hong Kong SAR, China),
Anh Quan Nguyen (Hong Kong University of Science and Technology) and
Jin Qi (Hong Kong University of Science and Technology).

Abstract. Auction theory relies on the critical assumption that the distribution of bidder valuations is common knowledge among agents, an assumption that is often considered unrealistic. There has been a surge of interest in robust mechanism design, which seeks mechanisms that perform well without precise distributional knowledge. Although there have been many successes in recent years, much of the literature still focuses on simple settings, especially in the one-object, one-bidder scenario, which limits its potential applications.

In this work, we extend a classic two-object, two-bidder model by introducing Knightian preferences. We assume bidders' beliefs are not a single prior but are restricted to an ambiguity set of probability distributions, and that incentive compatibility and individual rationality constraints must hold for all possible distributions within this set. This robust formulation, which safeguards against both distributional uncertainty and bidders' attitudes towards it, is shown to be equivalent to a tractable linear program. We further analyze the case of an uncertain seller who also exhibits Knightian preferences, applying the concept of finite adaptability from robust optimization—a novel approach in this context that yields simple, less conservative mechanisms.

Our analysis reveals several key findings. First, standard optimal bundling auctions remain optimal under some conditions even under significant uncertainty, which is a big contrast to many results in the literature that suggest selling items separately is optimal. Second, the choice between independent (separate) sales and bundling depends on the degree of ambiguity and the correlation of valuations. Third, when the seller is permitted to withhold objects, we find that the propensity to do so increases as the ambiguity level rises, especially when the low valuation is small relative to the potential gains from trade. Our findings provide a more nuanced understanding of multi-object sales and highlight the value of integrating analytical tools from robust optimization into mechanism design.

Keywords: robust mechanism design, robust optimization, finite adaptability, Knightian preferences

Screening with Fairness Consideration: A Unified Framework.

Zihao Xiang (The University of Hong Kong), Man-Chung Yue (The University of Hong Kong) and Zhi Chen (The Chinese University of Hong Kong).

Abstract. The screening problem aims to design an expected profit-maximizing mechanism for selling a product to a single buyer assuming the knowledge of her valuation distribution. In practice, we often want to impose a fairness constraint so that the mechanism is not overly favourable to any sub-population with a certain advantage. Due to the infinite-dimensional nature of the optimization problem, integrating fairness consideration in the screening problem while maintaining theoretical and computational tractability is highly non-trivial. In this work, we develop a unified framework for modelling, analyzing, and solving fair screening problems. A central theoretical result of our framework is the characterization of the extreme points of the fair screening problem, which allows us to reduce the infinite-dimensional to a finite-dimensional optimization problem. Furthermore, an efficient customized algorithm is devised to solve the resulting optimization problem. Several applications will be presented to illustrate the strength and versatility of our framework.

Keywords: Mechanism design, Fairness, Screening, Infinite-dimensional optimization

Contracting with a Mechanism Designer.

Tian Bai (The University of Hong Kong),
Yiding Feng (Hong Kong University of Science and Technology),
Yaohao Liu (University of Electronic Science and Technology of China),
Mengfan Ma (Central China Normal University) and
Mingyu Xiao (University of Electronic Science and Technology of China).

Abstract. This paper explores the economic interactions within modern crowdsourcing markets. In these markets, employers issue requests for tasks, platforms facilitate the recruitment of crowd workers, and workers complete tasks for monetary rewards. Recognizing that these roles serve distinct functions within the ecosystem, we introduce a three-party model that distinguishes among the principal (the requester), the intermediary (the platform), and the pool of agents (the workers). The principal, unable to directly engage with agents, relies on the intermediary to recruit and incentivize them. This interaction unfolds in two stages: first, the principal designs a profit-sharing contract with the intermediary; second, the intermediary implements a mechanism to select an agent to complete the delegated task.

We analyze the proposed model as an extensive-form Stackelberg game. Our contributions are fourfold: (1) We fully characterize the subgame perfect equilibrium and show that the intermediary's optimal mechanism coincides with the Bayesian revenue-optimal mechanism parameterized by the principal's contract. This characterization further reduces the principal's contract design problem to a novel auction-theoretic formulation we term "virtual value pricing", and reveals that "linear contracts are optimal" even when the task have multiple outcomes and agents' cost distributions are asymmetric. (2) To quantify the principal's utility loss from delegation and information asymmetry, we introduce the "price of double marginalization" (PoDM) and the classical "price of anarchy" (PoA), and derive tight or nearly tight bounds on both ratios under regular and monotone hazard rate (MHR) distributions. (3) We further examine these two ratios in a natural setting where the intermediary is restricted to anonymous pricing mechanisms, and show that similar qualitative insights continue to hold. (4) Finally, we extend our results on both PoDM and PoA to a robust framework that accommodates scenarios in which the principal lacks precise information about the market size.

Keywords: contract design, Bayesian auction design, price of anarchy, price of double marginalization, crowdsourcing

Session: Logistics and Supply Chains II

Venue: V303, Time: 16:00 – 17:45

Session Chair: Zihao Li

Reliable Supply Network Design under Correlated Uncertainties Caused by Disruptions.

Wenjie Li (The Hong Kong Polytechnic University),
Miao Song (The Hong Kong Polytechnic University),
Jia Shu (University of Electronic Science and Technology of China) and
Yongzhen Li (Southeast University).

Abstract. Supply chain disruptions bring significant uncertainties into network design. These uncertainties are usually correlated, particularly when disruptions are caused by natural disasters or systemic threats. This paper studies the problem of supply network design under disruptions that can cause various uncertainties and affect different aspects of the network. The realization of these uncertainties depends on the disruption scenarios, which are characterized by an unknown joint distribution. As the true distribution is unknown and difficult to predict accurately in practice, we assume it lies within a moment-based ambiguity set. To address this challenge, we formulate a two-stage distributionally robust model that simultaneously minimizes fixed location cost, inventory pre-positioning cost, allocation cost, and the expected transportation cost under the worst-case distribution.

Keywords: Supply network design, Correlated uncertain disruptions, Distributionally robust optimization, Cutting plane algorithm

Data-Driven Robust Inventory Management with Time-Series Demand.

Zhi Chen (The Chinese University of Hong Kong),
Tong Wang (The City University of Hong Kong) and
Fen Xu (The Chinese University of Hong Kong).

Abstract. We study the multi-period stochastic inventory management problem with time-series demand in a data-driven setting. When historical data is limited, the estimate-then-optimize method often suffers from overfitting and poor out-of-sample performance. To address this, we propose a data-driven robust optimization approach that constructs a Wasserstein ambiguity set capturing demand correlation and uncertainty across the entire planning horizon. We identify that this approach enables a recursive solution via robust dynamic programming, and we show that the resulting robust value functions are piecewise-linear and jointly convex, thus a state-dependent base-stock policy is robustly optimal. We also characterize the worst-case distribution as having a threshold form, which reduces the computational complexity of identifying these distributions from exponential to linear in sample size. For the positively correlated first-order autoregressive model, we prove submodularity of the value function, revealing strategic complementarity between inventory levels and realized demand, and establishing monotonicity of both the worst-case distribution's threshold and the base-stock level. Statistically, we derive finite-sample performance guarantees for the data-driven robust policy relative to the full-information optimal policy, extending existing results by explicitly accounting for demand correlation and distributional uncertainty. Numerical experiments demonstrate the superior out-of-sample performance of our data-driven robust policy, particularly with limited data, and underscore the importance of modeling general time-series demand.

Keywords: inventory management, time series, distributionally robust optimization, Wasserstein distance

Optimal Open-Loop Policy in Distributionally Robust Inventory Control with Uncertain Lead Times.

Xin Chen (Georgia Institute of Technology), Zhuoyu Long (Chinese University of Hong Kong) and
Jingjun Men (Chinese University of Hong Kong).

Abstract. In this paper, we consider a multi-periods inventory control problem with uncertain lead times, which present significant challenges in both practice and academic. We consider the distributionally robust optimization (DRO) setting where only the marginal distribution is known and we aim at minimizing the worst-case expected cost. We focus on open-loop policies, which have attracted substantial attention recently due to their asymptotic optimality and appealing numerical performance. By showing the supermodularity property of the cost function, we develop an efficient algorithm that identifies the worst-case distribution and reformulate the DRO inventory problem into a Mixed Integer Linear Programming (MILP) problem. Further, we identify necessary and sufficient conditions under which there are no order-crossovers. When such conditions are satisfied, our problem can be further simplified into a shortest path problem on a directed acyclic network, and thus can be solved within polynomial time. Numerical examples are provided to demonstrate the practical effectiveness of our approach.

Keywords: inventory, uncertain lead time, distributionally robust optimization, open-loop policy

Assortment Optimization with Customer-Centric Fairness.

Yihua He (HKUST), Qian Liu (HKUST) and Jin Qi (HKUST).

Abstract. The rapid growth of e-commerce platforms has enabled businesses to implement customized assortment strategies, tailoring product offerings to different customer groups. While these strategies can significantly enhance profitability, they often raise concerns about fairness, particularly when certain customer groups receive fewer options. Such practices can provoke customer backlash and damage a company's reputation. Motivated by these issues, we propose a fairness-aware framework for assortment optimization that balances revenue maximization with equitable treatment of customer groups. By introducing cardinality-based fairness constraints, our framework ensures that sellers provide a similar number of products to different groups, subject to varying levels of fairness tolerance. Our analysis demonstrates that such fairness constraints compel sellers to offer more products to disadvantaged groups, as well as enhancing their customer surplus, even though the seller does not explicitly target this outcome. However, despite this benefit, the sellers have no incentive to exceed the minimum fairness requirements. Building on the structural properties, we develop efficient polynomial algorithms to solve the problem exactly when customers are categorized into two categories and for the multi-type scenarios with semi-independent tolerance levels. We also propose an approximate algorithm for the most general multi-type scenarios with a guaranteed performance bound. Our numerical results further show that fairness constraints have a limited impact on seller revenue, making them less likely to resist such regulations. However, policymakers must carefully weigh the trade-offs as fairness constraints can reduce overall customer surplus in some cases.

Keywords: fairness, assortment, discrete choice model, revenue management

Fully Online Matching with General Stochastic Arrivals and Departures.

Zihao Li (National University of Singapore),
Hao Wang (University of Science and Technology of China) and
Zhenzhen Yan (Nanyang Technological University).

Abstract. We study a fully online matching problem with general stochastic arrivals and departures. In this model, each online arrival follows a known identical and independent distribution over a fixed set of agent types. Its sojourn time is unknown in advance and follows type-specific distributions with known expectations. The goal is to maximize the weighted reward from successful matches. To solve this problem, we propose a linear program (LP)-based algorithm whose competitive ratio is lower bounded by 0.192 under mild conditions. To demonstrate the challenges of the problem, we further establish several hardness results. In particular, we show that no online algorithm can achieve a competitive ratio better than $\frac{1}{2}$ in this model, and if using our LP as a benchmark for competitive ratio analysis, no algorithm can achieve a better ratio than $\frac{1}{3}$.

When no assumptions are made regarding the sojourn time distributions, we demonstrate that it is impossible to achieve a positive competitive ratio for the general case using our LP as a benchmark for competitive ratio analysis.

We further extend our model to accommodate general sojourn times under Poisson arrivals and demonstrate a better competitive ratio compared to state-of-the-art results derived under Poisson arrivals and departures, a special case of our general settings.

Finally, we demonstrate the effectiveness and efficiency of our algorithm numerically.

Keywords: Fully Online Matching, Randomized Algorithm, Competitive Ratio

Session: Robust Optimization

Venue: V315, Time: 16:00 – 17:45

Session Chair: Jie Wang

Fast Algorithm for Wasserstein Distributionally Robust Markov Decision Processes.

Ling Dai (City University of Hong Kong), Zhuodong Yu (City University of Hong Kong) and Chin Pang Ho (City University of Hong Kong).

Abstract. Markov Decision Processes (MDPs) are widely used for dynamic decision-making under uncertainty, but they suffer from unknown transition probabilities in practice. Distributionally robust MDPs have been proposed to address this issue, but solving them remains computationally challenging, especially with Wasserstein ambiguity sets. Existing methods for solving such problems are either limited to specific cases or do not efficiently recover the optimal policy. We propose a new algorithmic framework that supports general Wasserstein ambiguity sets of any order. Our method combines an approximate trisection technique with a bisection scheme to solve the Bellman updates efficiently. To recover the optimal policy, we develop a new algorithm based on a sub-differential characterization of the optimal policy. The resulting algorithms achieve quasi-linear time complexity in the problem size for both computing the Bellman update and recovering the optimal policy from the optimal value. Consequently, the overall complexity of computing the optimal value function and the optimal policy matches the order of that for solving a nominal MDP. This suggests that our methods offer a high level of scalability while accounting for uncertainty in the transition probabilities. Numerical experiments demonstrate that our approach significantly improves computational efficiency while maintaining high-quality solutions.

Keywords: Markov Decision Processes, Distributionally Robust Optimization, Wasserstein Ambiguity Set, Computational Complexity

Inpatient Spillover: A Distributionally Robust Chance-Constrained Approach.

Qihao Wu (The University of Hong Kong), Yinuo Yu (Georgia Institute of Technology), Abdel Lisser (CentraleSupélec), Ka Chung Abraham Wai (The University of Hong Kong) and Yong-Hong Kuo (The University of Hong Kong).

Abstract. Effective and efficient inpatient management is crucial for providing quality healthcare services. One of the main challenges is the imbalance between patient demands and inpatient resources across various units. When the intended inpatient ward is fully occupied, deciding whether to hold the admission in queue for an available bed or to redirect the patient to another ward (a less preferred option) presents a dilemma -- this depends on the costs of waiting and redirecting patients, and the potential future consequences. Based on a Markov decision process framework, our research proposes a distributionally robust model with chance constraints and optimizes hospital overflow and transfer decisions. We then design state aggregation strategies to address the curse of dimensionality and measure the costs in an uncertainty set with a known mean. The original model can be reformulated as a second-order cone program that can be solved efficiently. The structural properties further facilitate computational efficiency. Our optimal policies are calibrated with real-world historical data and pooling policies. Our experiments suggest that our proposed approach reduces accumulated waiting time compared status quo. Moreover, off-service placements and total costs are lower than those under a typical pooling policy. Our experimental results further indicate that the proposed solution is robust and practically tractable.

Keywords: Inpatient Flow Management, Distributionally Robust Optimization, Chance-Constrained Markov Decision Processes

A Tractable Approach for Contextual Distributionally Robust Optimization with Causal Optimal Transport.

Fenglin Zhang (The Chinese University of Hong Kong (Shenzhen)) and Jie Wang (The Chinese University of Hong Kong (Shenzhen)).

Abstract. This paper studies a contextual distributionally robust optimization (DRO) framework that leverages covariate information to enhance decision-making. We first introduce the causal Sinkhorn discrepancy (CSD), an entropy-regularized causal Wasserstein distance that facilitates continuous transport plans while preserving the conditional dependencies between random variables. Then, we formulate a DRO model with the CSD-based ambiguity set, termed causal-SDRO. We derive its tractable dual reformulation, revealing its equivalence to a stochastic multi-level compositional optimization program and characterizing its worst-case distribution as a continuous mixture of Gibbs distributions. To solve this problem for parametric decision rules, we analyze the sample complexity of sample average approximation method of the dual problem, and develop a stochastic compositional gradient algorithm, which converges to an ε-stationary point at the same rate as the standard stochastic gradient descent method. Numerical experiments on synthetic and real-world datasets validate the efficacy of the proposed approach.

Keywords: Contextual distributionally robust optimization, Causal optimal transport, Sinkhorn discrepancy, Stochastic compositional optimization

Intelligent Dominance Dispatch of Emergency Ambulance.

Zikun Lin (The Chinese University of Hong Kong),
Daniel Zhuoyu Long (The Chinese University of Hong Kong) and
Viet Anh Nguyen (The Chinese University of Hong Kong).

Abstract. Timely ambulance response is paramount in life-threatening emergencies such as out-of-hospital cardiac arrest (OHCA). An effective dispatch policy requires balancing two conflicting criteria: (i) minimizing response time to improve survival rates, and (ii) ensuring judicious use of the limited ambulance fleet. We propose a novel bilevel representation learning approach and a dominance analysis-based dispatch framework to effectively model uncertain traffic conditions. Besides dispatching a single ambulance along the nominally fastest path identified at the time of the call, our approach also assesses whether this choice remains sufficiently advantageous under travel time uncertainty. If an alternative path could potentially offer an earlier arrival under certain scenarios, we dispatch a second ambulance from the same depot or an alternative one, following a distinct path, to mitigate the risk. We validate our approach using real-world records of approximately 1,800 OHCA cases from Northwest Hong Kong Island. The advantage of the dominance analysis-based dispatch system lies in its use of alternative dispatching paths, which are selected according to urban structures and traffic conditions. Our work presents a new framework for dynamic resource allocation under uncertainty, providing practical insights for optimizing emergency services.

Keywords: ambulance dispatch, robust optimization, bilevel programming

An Iterative Sampling Approach for Solving Sinkhorn Distributionally Robust Optimization.

Jindong Jiang (CUHK-SZ) and Jie Wang (CUHK-SZ).

Abstract. Distributionally robust optimization (DRO) has emerged as a powerful paradigm for reliable decision-making under uncertainty. This paper focuses on DRO with ambiguity sets defined via the Sinkhorn discrepancy—an entropy-regularized Wasserstein distance—referred to as Sinkhorn DRO. Existing work primarily addresses Sinkhorn DRO from a dual perspective, leveraging its formulation as a conditional stochastic optimization problem, for which many stochastic gradient methods are applicable. However, the theoretical analyses of such methods often rely on the boundedness of the loss function. In contrast, we study Sinkhorn DRO from the primal perspective. Our approach proceeds iteratively: in each step, we first sample from the worst-case distribution, then construct a gradient estimator with both low bias and low variance to update the decision variable. This strategy has two key advantages: (i) it removes the boundedness requirement on the loss function; and (ii) it naturally generates samples from the evolving worst-case distribution, providing additional interpretation on how the decision sequence shapes the worst-case scenario. We establish theoretical guarantees for the proposed method and demonstrate its effectiveness in numerical study.

Keywords: Distributionally robust optimization, optimization, machine learning

Session: Transportation

Venue: V316, Time: 16:00 – 17:45

Session Chair: Jinha Hibino

Adaptive bus bridging strategy under rail transit disruptions via a two-stage reinforcement learning driven approach.

Junyi Wang (City University of Hong Kong), Andy Chow (City University of Hong Kong), Zeyu Zhang (City University of Hong Kong) and Yukun Xu (City University of Hong Kong).

Abstract. Disrupted urban rail transit services often rely on bus bridging operations to provide temporary connecting services. However, deriving adaptive bus bridging plans subject to prevailing operational conditions and constraints in a computationally effective way remains a technical challenge. This paper presents a two-stage adaptive bus bridging design and operation framework with the use of mathematical optimization and reinforcement learning approaches. Upon a disruption occurring, a set of candidate bridging routes is first derived via column generation with consideration of the topology of the disrupted rail transit network and the spatial distribution of demand for the bridging services from the stranded passengers. The bridging service schedules, routes, and use of vehicles are then determined with respect to the prevailing demand and system status through a Markov Decision Process (MDP). The bus bridging plans aim to minimize the passengers' travel times and waiting times, subject to the availability of vehicle resources. For computational effectiveness, the MDP is to be solved by a reinforcement learning (RL) driven approach. Numerical experiments are conducted using real-world data collected from the Hong Kong Light Rail Transit (LRT) network. The proposed method delivers significant advantages in terms of both optimization performance and computational time when compared with other approaches and control settings. The present work contributes to robust transit operations in real time.

Keywords: Rail transit disruptions, adaptive bus bridging, column generation, Markov decision process, reinforcement learning

Adversarial multi-agent reinforcement learning for robust train service scheduling under feeder bus service uncertainties.

Shouyi Wang (City University of Hong Kong), Andy Chow (City University of Hong Kong) and Junyi Wang (City University of Hong Kong).

Abstract. This paper presents a robust and adaptive control system for network-wide train service scheduling under feeder bus service uncertainties through an adversarial multi-agent deep reinforcement learning approach. To enhance system robustness against the worst-case feeder bus service disturbance, we formulate a Markov game with one attacker agent and multiple defender agents. Within this attacker-defender framework, the attacker agent increases the number of feeder bus services on the designated bus route in response to the increased passenger demand caused by various exogenous factors such as weather conditions, holidays, special events, or other demand-generating scenarios. In contrast, the defender agents manage the dispatching headways, dispatching routes, and assigned fleet sizes for train services departing from specific depots to minimize associated passenger waiting times and operating costs under adversarial disturbances. To address the inherent computational challenge of this robust optimization problem, we implement a multi-agent policy optimization solution framework wherein each agent employs artificial neural networks (ANNs) to approximate a decentralized policy function and a centralized value function approximation (VFA). Furthermore, we develop a multi-agent proximal policy optimization (MAPPO) algorithm to train the policies and VFAs of all defender agents and the attacker agent. The proposed framework is validated using real-world scenarios and operational data from the Hong Kong Light Rail Transit network. Computational results demonstrate the robustness and adaptability of the developed framework under various degrees of feeder bus service disturbances. This study contributes to robust and real-time rail train service scheduling in network-wide rail transit systems by integrating advanced control and optimization technologies.

Keywords: Adaptive train scheduling, Feeder bus uncertainties, Markov decision process, Multi-modal transit systems, Multi-agent deep reinforcement learning

A Pricing Model for Enhancing the Utilization of Public Electric Vehicle Facilities: A Case Study in Hong Kong.

Jiangxue Han (The University of Hong Kong) and Yong-Hong Kuo (The University of Hong Kong).

Abstract. With the wide acceptance of electric vehicles (EVs) and the significant increase in EVs, greater attention has been paid to developing public EV charging facilities. However, issues such as constrained charging spaces, EV overcharging, and inconvenient charging locations might lead to difficulties in achieving a high utilization of public EV charging facilities. Existing research primarily focuses on enhancing the spatial planning of EV charging stations. In this research, we provide a pricing perspective as a more flexible and adaptive complement to optimize the utilization of the established EV charging infrastructure.

We focus on the case of Hong Kong in this work. While the Government has adopted an (almost) uniform pricing policy, we aim to develop a more effective pricing scheme for the public EV charging services in the city. Based on the primary factors that may influence drivers' choice in EV charging facilities, we build an optimization model to explore the mutual influences between these factors, with the objectives to improve the utilization of public EV chargers and to satisfy the EV charging demands. Numerical experiments are conducted based on real data to examine the performance of the proposed pricing mechanism.

Keywords: Electric vehicle, Charging facility, Pricing model

Constructive and local search heuristics for a point-to-point airline network design problem.

Jinha Hibino (Chuo University), Shinji Imahori (Chuo University) and Mihiro Sasaki (Nanzan University).

Abstract. Airline network design models are fundamentally classified into two types: a hub-and-spoke model and a point-to-point model. The latter model offers high flexibility in network design; however, research on its mathematical models is scarce. In 2017, Sasaki and Furuta presented the point-to-point airline design problem (PPANP) and solved the problem using a naive constructive heuristic.

The PPANP seeks an optimal network with a given number of edges to maximize an airline's total captured revenue. This problem assumes that the locations of airports and the revenue between each pair of airports are given. In this talk, we present: (1) a constructive heuristic to find an initial high-quality solution, and (2) local search heuristics to improve solutions further.

We have implemented effective heuristics that run in a reasonable time for large-scale instances with up to 100 airports. The presented constructive heuristic is based on the greedy heuristic presented by Sasaki and Furuta. It develops a network by iteratively adding one edge at a time. This method always selects an additional edge that provides the largest increase in revenue. The computational performance is enhanced by focusing only on OD pairs with potential revenue increase. As a result, we can efficiently identify the best edge among the candidates to add. Furthermore, we employ the same technique to accelerate the edge addition process within our local search heuristics.

To improve the initial solution generated by the constructive heuristic, we propose three local search heuristics: BestR&A, BestR→BestA, and BestA→BestR. These heuristics seek a better network. First, BestR&A identifies a pair of edges to swap that achieves the largest net revenue increase. While BestR&A exhaustively searches all possible pairs, BestR→BestA and BestA→BestR explore a smaller neighborhood. Both heuristics find a pair of edges: an edge to remove, and an edge to add. The difference between them is their search order: BestR→BestA first finds the edge to remove and then the edge to add. BestA→BestR finds the two edges in the reverse order of BestR→BestA.

We evaluate the performance of our heuristics from the perspectives of computation time and solution quality. Our experiments demonstrate that the constructive heuristic rapidly generates high-quality initial solutions, which are subsequently refined by the local search heuristics. The local search heuristics exhibit a clear trade-off. Although BestR&A is the most time-consuming per iteration, it finds better solutions more frequently. In contrast, BestR→BestA and BestA→BestR have faster iterations but a lower frequency of improvements.

Keywords: Airline network design problem, Point-to-point design model, Local search heuristic, Efficient implementation

Notes to Speakers in Parallel Sessions

- Each presentation is allocated 15 minutes. Additional 4 minutes are designated for Questions and Answers (Q & A). Q & A timing will be handled flexibly by the session chair, depending on available time.
- A computer will be available in each session room.
- Speakers may use their own laptops or the computer provided at the venue for their presentation.
- Speakers should arrive at the designated presentation venue 10 minutes before their session for setup and testing.
- If using the computer at the venue, the presentation file must be provided to the helper before the session begins.
- Presentation files should be saved on a USB device.
- The session chair will remind presenters when 5 minutes and 1 minute remain for their presentation.

Notes to Session Chairs in Parallel Sessions

- Session chairs are expected to arrive at the presentation venue at least 10 minutes prior to the session start time.
- Please liaise with the student helper at the session before the session begins and confirm the presence of presenters and verify their names.
- Session chairs may start the session on time with a brief self-introduction and session title.
- Each presentation is allocated 15 minutes for the speaker's presentation, with an additional 4 minutes allocated for Questions and Answers (Q & A). Q & A can be handled flexibly by the session chair, depending on the number of speakers and available time.
- The session chair should closely monitor each speaker's time. Please provide time signals (e.g., 5 minutes, 1 minute, and Time's up) to remind speakers of the time, and ask speakers to stop when their allocated time is up.
- If a presentation exceeds the allocated time, the Q & A session time will be reduced accordingly.
- The session chair will make announcements to the audience as needed.

About APORS

The Association of Asia Pacific Operational Research Societies, abbreviated as APORS, is one of the regional groupings under the International Federation of Operational Research Societies (IFORS). Its objectives are to develop operational research (OR) as a unified science and to advance it in the Asian Pacific nations through sponsoring international meetings, exchange information on operational research, encourage AP nations to establish their national societies, and promoting OR education and applications. By 2011, APORS has 12 member societies and becomes actively involved in all IFORS activities.

APORS, like other regional groupings within IFORS, does not support a secretariat for daily operation. Instead, its president and his/her national society provide secretariat function in communicating with other member societies. So far, APORS does not have any funding source. Therefore, the council meeting, generally during APORS and IFORS conferences, plays critical part in discussing some important matters and making decisions on election, triennial conference host, and collaborative activities, etc.

More information about APORS is available at http://apors.org/

About ORSHK

The Operational Research Society of Hong Kong (ORSHK) has been established since 1980. The objectives of the Society are to promote the dissemination of knowledge and information relating to Operational Research and Management Science by means of meetings, publications, awards and related activities. We also aim to promote mutual interaction and cooperation among researchers, practitioners, and professional organisations, related to Operational Research and Management Science discipline, in Hong Kong.

Since 1983, ORSHK has been affiliated to the International Federation of Operational Research Societies (IFORS) which is currently composed of 54 national societies.

More information about ORSHK is available at https://orshongkong.wixsite.com/home

Organising Committee

- Prof. Miao SONG (OC co-chair), The Hong Kong Polytechnic University
- Dr. Yancheng YUAN (OC co-chair), The Hong Kong Polytechnic University
- Dr. Andy H.F. CHOW (PC co-chair), City University of Hong Kong
- Dr. Yong-Hong KUO (PC co-chair), The University of Hong Kong
- Prof. Yukun CHENG, Jiangnan University
- Dr. Clint C.P. HO, City University of Hong Kong
- Dr. Sin C. HO, The Chinese University of Hong Kong
- Prof. Janny M.Y. LEUNG, The University of Macau
- Dr. Lishuai LI, City University of Hong Kong
- Dr. Carrie Ka Yuk LIN, City University of Hong Kong
- Dr. Wei LIU, The Hong Kong Polytechnic University
- Dr. Daniel Z. LONG, The Chinese University of Hong Kong
- Dr. Jin QI, The Hong Kong University of Science and Technology

Venue

Lecture Theater V322, Jockey Club Innovation Tower (Core V), The Hong Kong Polytechnic University

Getting to PolyU Campus

- **MTR.** Get off at Hung Hom station on the East Rail Line or Tuen Ma Line, take the footbridge at Exit A1 to reach PolyU main campus.
- **Bus.** For public buses, get off at "Cross Harbour Tunnel Toll Plaza" stop (entrance/exit of the Tunnel on the Kowloon side).
- **Taxi.** Drop-off at the entrance on Cheong Wan Road.