

I-DISC wishes to thank all the speakers that kindly gave us permission to record and share their presentations with you. These can be viewed by clicking on the presentation title.



Joren Gijsbrechts, Universidade Católica Portuguesa, Portugal

"Can Deep Reinforcement Learning Improve Inventory Management? Performance on Dual Sourcing, Lost Sales and Multi-Echelon Problems"

Abstract: Is Deep Reinforcement Learning (DRL) effective at solving inventory problems? Given that DRL has successfully been applied in computer games and robotics, supply chain researchers and companies are interested in its potential in inventory management. We provide a rigorous performance evaluation of DRL in three classic and intractable inventory problems: lost sales, dual-sourcing and multi-echelon inventory management. We model each inventory problem as a Markov Decision

Process and apply and tune the Asynchronous Advantage Actor Critic (A3C) DRL algorithm for a variety of parameter settings. We demonstrate that the A3C algorithm can match performance of state-of-the-art heuristics and other approximate dynamic programming methods. While the initial tuning was computationally- and time-demanding, only small changes to the tuning parameters were needed for the other studied problems. Our study provides evidence that DRL can effectively solve inventory problems. This is especially promising when problem-dependent heuristics are lacking. Yet generating structural policy insight or designing specialized policies that are (ideally provably) near optimal remains desirable.

BioSketch: Joren Gijsbrechts is an Assistant Professor in Operations Management. He graduated as a Bachelor and Master in Business Engineering at the University of Antwerp in Belgium. Prior to obtaining his doctoral degree from KU Leuven in Belgium, he attained business experience in the Supply Chain and Operations division of Procter and Gamble in Sweden. As a PhD student, he regularly visited Kellogg School of Management and has on-going research projects with renowned institutions. His research centers around data-driven decision making in Operations Management with a strong focus on the recent developments in Machine Learning and Prescriptive Analytics. His models have assisted companies to improve their inventory and transportation management. In addition to research, he is providing guest lectures and company workshops on the recent developments within Analytics.



Swati Gupta, Georgia Tech

"Opportunities for Ethical ML and Supply Chains"

Abstract: In any application area, disparate treatment, disparate impact and resulting unintended outcomes are a reality across the spectrum. These are expected to be further exacerbated by the scale and adoption of cutting-edge technologies (such as optimized ML algorithms, reinforcement learned decisions) and the aftermath of the current pandemic. The interaction of optimization and statistical models with complex socio-economic systems makes development of deployable

machine learning algorithms extremely challenging in practice. In this talk, I will discuss opportunities to develop ethical supply chain systems, to mitigate the harm of socio-economic trends embedded in the data, to ensure that the developed systems are fair, ethical, transparent to the extent possible, accountable under uncertainty in the data, and reduce a deeper propagation of biases in multi-level decisions.

BioSketch: Dr. Swati Gupta is an Assistant Professor and Fouts Family Early Career Professor in the H. Milton Stewart School of Industrial & Systems Engineering at Georgia Institute of Technology. She is the lead of Ethical AI at the NSF AI Institute on AI4OPT (ai4opt.org). She received a Ph.D. in Operations Research from MIT in 2017 and a joint Bachelors and Masters in CS from IIT, Delhi in 2011. Dr. Gupta's research interests are in optimization, machine learning and algorithmic fairness. Her work spans various application domains such as revenue management, energy and quantum computation. She received the JP Morgan Chase Early Career Faculty Award in 2021, Class of 1934 CIOS Honor Roll 2021 and NSF CISE Research Initiation Initiative (CRII) Award in 2019. She was also awarded the prestigious Simons-Berkeley Research Fellowship in 2017-2018, where she was selected as the Microsoft Research Fellow in 2018. Dr. Gupta received the Google Women in Engineering Award in India in 2011. Dr. Gupta's research is partially funded by the NSF and DARPA.



Andrea Lodi, Cornell Tech

"Heuristics for Mixed-Integer Optimization through a Machine Learning Lens"

Abstract: In this talk, we discuss how a careful use of Machine Learning concepts can have an impact in primal heuristics for Mixed-Integer Programming (MIP). More precisely, we consider two applications. First, we design a data-driven scheduler for running both diving and large-neighborhood search heuristics in SCIP, one of the most effective open-source MIP solvers. Second, we incorporate a major learning component into Local Branching, one of the most well-known primal heuristic

paradigms. In both cases, computational results show solid improvements over the state of the art.

BioSketch: Andrea Lodi is an Andrew H. and Ann R. Tisch Professor at the Jacobs Technion-Cornell Institute at Cornell Tech and the Technion. He is a member of the Operations Research and Information Engineering field at Cornell University. He received his PhD in System Engineering from the University of Bologna in 2000 and he was a Herman Goldstine Fellow at the IBM Mathematical Sciences Department, NY in 2005–2006. He was a full professor of Operations Research at DEI, the University of Bologna between 2007 and 2015. Since 2015, he has been the Canada Excellence Research Chair in "Data Science for Real-time Decision Making" at Polytechnique Montréal. His main research interests are in Mixed-Integer Linear and Nonlinear Programming and Data Science and his work has received several recognitions including the IBM and Google faculty awards. Andrea is the recipient of the INFORMS Optimization Society 2021 Farkas Prize. He is the author of more than 100 publications in the top journals of the field of Mathematical Optimization and Data Science. He serves as Editor for several prestigious journals in the area. He has been the network coordinator and principal investigator of two

large EU projects/networks, and, since 2006, consultant of the IBM CPLEX research and development team. Andrea Lodi is the co-principal investigator of the project “Data Serving Canadians: Deep Learning and Optimization for the Knowledge Revolution,” recently funded by the Canadian Federal Government under the Apogée Programme and scientific co-director of IVADO, the Montréal Institute for Data Valorization.



Polly Mitchell-Guthrie, Kinaxis

["How to make sure machine learning has an impact on supply chains"](#)

Abstract: The tumult brought on by the pandemic has jumpstarted interest in applying machine learning and automation to address supply chain problems, but adoption rate is still low. Corporate initiatives around AI, machine learning, and data science abound, but measurements of their failure rate range are as high as 87%. The problem space is rich in complexity and worthy of research advancements, so what is required to translate innovation into value on the ground? Solving supply chain problems means thinking end to end across the entire supply network, but

it also means thinking end to end across the analytical life cycle, from business problem framing to modeling to MLOps and maintenance. Anything less will fall short on delivering results. This talk will address practical considerations for incorporating machine learning into supply chains based on experience across many different companies in a variety of industries.

BioSketch: Polly is the VP of Industry Outreach and Thought Leadership at Kinaxis, a supply chain planning and analytics software company. Previously, she was Director of Analytical Consulting Services at the University of North Carolina Health Care System and worked in various roles at SAS, in Advanced Analytics R&D, as Director of the SAS Global Academic Program, and in Alliances. She has an MBA from the Kenan-Flagler Business School of the University of North Carolina at Chapel Hill, where she also received her BA in Political Science as a Morehead Scholar. She has been very active in INFORMS (the leading professional society for operations research and analytics) and co-founded the third chapter of Women in Machine Learning and Data Science (now more than 60 chapters worldwide).



Richard Pibernik, Julius-Maximilians-University, Würzburg, Germany

["From small to large data – how can we leverage synthetic data for ML in Operations & Supply Chain Management"](#)

Abstract: “Data-driven, End-to-End, automated” is the vision underlying many of the new approaches that have recently been developed in the ML-community (e.g. deep reinforcement learning, attention learning). These approaches are typically very “data-hungry”. In many real-life settings in OM and SCM, however, we oftentimes do not have large enough data to exploit these new approaches and to live up to this vision. For inventory and capacity management, for example, we may have a large number of

explanatory variables (“features”), but often only have a very limited number of relevant historical observations, as we take decisions on a daily, weekly, or monthly basis. “Small data” causes a number of theoretical and practical problems that may – at least to some extent – be remedied by new generative ML-techniques. The idea of these techniques is to first implicitly learn the unknown data (demand) generating process, and then to generate large multi-variate artificial data. In this talk we present the first results of our work on synthetic data generation, outline problems and limitations, and discuss the future potential of such approaches.

BioSketch: Richard Pibernik is a Full Professor of Logistics & Quantitative Methods at the University of Würzburg in Germany. In Würzburg he heads a research group dedicated to data-driven Operations and Supply Chain Management. Richard is also an Adjunct Professor at the Zaragoza Logistics Center and a Visiting Professor at the Malaysia Institute for Supply Chain Innovation – both institutions are part of MIT’s Global Scale Network. Richard has published his research in numerous renowned international journals and has been responsible, as a principal investigator, for many projects that were funded by industry and public funding agencies. His team is currently working with numerous companies on projects focused on data-driven Supply Chain Management, Supply Chain analytics, and exploitation of big data in Supply Chain Management.



Cynthia Rudin, Duke University

"Almost Matching Exactly for Observational Causal Inference"

Abstract: I will present an approach that aims to match a current situation with almost identical situations from the past, in order to use these past situations to predict the future. This approach has proven invaluable in the study of complex systems where causal effects can easily be confused with correlations. The matching framework I will present, called Almost Matching Exactly, is useful for causal inference in the potential outcomes setting. This framework has several important elements: (1) Its algorithms create matched groups that are interpretable. The goal is to match

treatment and control units as closely as possible, or "almost exactly." (2) Its algorithms create accurate estimates of individual treatment effects. This is because we use machine learning on a separate training set to learn which features are important for matching. Variables that are important are “stretched” so that the matched groups agree closely on these variables. (3) Our methods are fast and scalable. In summary, these methods rival black box machine learning methods in their treatment effect estimation accuracy but have the benefit of being interpretable and easier to troubleshoot. Our lab website is here: <https://almost-matching-exactly.github.io>

BioSketch: Cynthia Rudin is a professor of computer science, electrical and computer engineering, statistical science, and biostatistics & bioinformatics at Duke University. She directs the Interpretable Machine Learning Lab, whose goal is to design predictive models with reasoning processes that are understandable to humans. Her lab applies machine learning in many areas, such as healthcare, criminal justice, and energy reliability. She holds an undergraduate degree from the University at Buffalo, and a PhD from Princeton University. She is the recipient of the 2021 Squirrel AI Award

for Artificial Intelligence for the Benefit of Humanity from the Association for the Advancement of Artificial Intelligence (the "Nobel Prize of AI"). She is a fellow of the American Statistical Association and a fellow of the Institute of Mathematical Statistics. Her work has been featured in many news outlets including the NY Times, Washington Post, Wall Street Journal, and Boston Globe.



Zuo-Jun Max Shen, UC-Berkeley / University of Hong Kong

"A Practical End-to-End Inventory Management Model with Deep Learning"

Abstract: We investigate a data-driven multi-period inventory replenishment problem with uncertain demand and vendor lead time (VLT), with accessibility to a large quantity of historical data. Different from the traditional two-step predict-then-optimize (PTO) solution framework, we propose a one-step end-to-end (E2E) framework that uses deep-learning models to output the suggested replenishment amount directly from input features without any intermediate step. The E2E model

is trained to capture the behavior of the optimal dynamic programming solution under historical observations, without any prior assumptions on the distributions of the demand and the VLT. By conducting a series of thorough numerical experiments using real data from one of the leading e-commerce companies, we demonstrate the advantages of the proposed E2E model over conventional PTO frameworks. We also conduct a field experiment with JD.com and the results show that our new algorithm reduces holding cost, stockout cost, total inventory cost and turnover rate substantially compared to JD's current practice. For the supply-chain management industry, our E2E model shortens the decision process and provides an automatic inventory management solution with the possibility to generalize and scale. The concept of E2E, which uses the input information directly for the ultimate goal, can also be useful in practice for other supply-chain management circumstances.

BioSketch: Zuo-Jun Max Shen is the Vice-President and Pro-Vice-Chancellor (Research) and the Chair Professor in Logistics and Supply Chain Management at the University of Hong Kong. He is on leave from the University of California, Berkeley, where he is a Chancellor's Professor in the Department of Industrial Engineering and Operations Research and the Department of Civil and Environmental Engineering. He received his Ph.D. from the Department of Industrial Engineering and Management Sciences at Northwestern University. He has been active in the following research areas: integrated supply chain design and management, operations management, data driven optimization algorithms and applications, energy systems, and transportation system planning and optimization. Max has extensive research collaborations with government agencies as well as private companies (both US and international). He is currently the Chief Supply Chain Scientist for JD.com. Max is serving as the president-elect for the Production and Operations Management Society, a Department Editor for Production and Operations Management, and Associate Editors for leading journals such as Operations Research, Management Science, Manufacturing & Service Operations Management, Decision Sciences, Naval Research Logistics, and IISE Transactions. Max received the CAREER award from National Science

Foundation, the Franz Edelman Laureate Award from INFORMS, won several best paper awards, and was elected Fellow of INFORMS in 2018.



David B. Shmoys, Cornell University, NY

["Algorithmic Tools for US Congressional Districting: Fairness via Analytics"](#)

Abstract: The American winner-take-all congressional district system empowers politicians to engineer electoral outcomes by manipulating district boundaries. To date, computational solutions mostly focus on drawing unbiased maps by ignoring political and demographic input, and instead simply optimize for compactness and other related metrics. However, we maintain that this is a flawed approach because compactness and fairness are orthogonal qualities; to achieve a

meaningful notion of fairness, one needs to model political and demographic considerations, using historical data. We will discuss two papers that explore and develop this perspective. In the first (joint with Wes Gurnee), we present a scalable approach to explicitly optimize for arbitrary piecewise-linear definitions of fairness; this employs a stochastic hierarchical decomposition approach to produce an exponential number of distinct district plans that can be optimized via a standard set partitioning integer programming formulation. This enables the largest-ever ensemble study of congressional districts, providing insights into the range of possible expected outcomes and the implications of this range on potential definitions of fairness. In the second paper (joint with Nikhil Garg, Wes Gurnee, and David Rothschild), we study the design of multi-member districts (MMDs) in which each district elects multiple representatives, potentially through a non-winner-takes-all voting rule (as currently proposed in H.R. 4000). We carry out large-scale analyses for the U.S. House of Representatives under MMDs with different social choice functions, under algorithmically generated maps optimized for either partisan benefit or proportionality. We find that with three-member districts using Single Transferable Vote, fairness-minded independent commissions can achieve proportional outcomes in every state (up to rounding), and this would significantly curtail the power of advantage-seeking partisans to gerrymander. We believe that this work opens up a rich research agenda at the intersection of social choice and computational redistricting.

BioSketch: David Shmoys is the Laibe/Acheson Professor and Director of the Center for Data Science for Enterprise & Society at Cornell University. He obtained his PhD in Computer Science from the University of California at Berkeley in 1984, and held postdoctoral positions at MSRI in Berkeley and Harvard University, and a faculty position at MIT before joining the faculty at Cornell University. He was Chair of the Cornell Provost's "Radical Collaborations" Task Force on Data Science and was co-Chair of the Academic Planning Committee for Cornell Tech. His research has focused on the design and analysis of efficient algorithms for discrete optimization problems, with applications including scheduling, inventory theory, computational biology, computational sustainability, and most recently, data-driven decision-making in the sharing economy. His work has highlighted the central role that linear programming plays in the design of approximation algorithms for NP-hard problems. His book (co-authored with David

Williamson), The Design of Approximation Algorithms, was awarded the 2013 INFORMS Lanchester Prize and his work on bike-sharing (joint with Daniel Freund, Shane Henderson, and Eoin O'Mahony) was awarded the 2018 INFORMS Wagner Prize. He is a Fellow of the ACM, INFORMS, and SIAM, and was an NSF Presidential Young Investigator.



Larry Snyder, Lehigh University "[Chips and Beer: Using Machine Learning to Optimize Inventory](#)"

Abstract: I will discuss two applications of machine learning in inventory optimization. The first uses deep neural networks (DNN) to optimize the order quantity in a newsvendor-type problem in which the demand distribution depends on exogenous features. Our approach uses DNN to optimize the order quantity directly from data, as opposed to the classical “forecast-then-optimize” approach. (I’ll illustrate this idea using an example based on potato chips.) In the second application, we develop a deep reinforcement learning (DRL) agent to play the beer game, a popular classroom activity that demonstrates certain aspects of inventory management. Our DRL agent learns near-optimal performance when its computerized “teammates” follow a base-stock policy. More interestingly, it outperforms the best-known policy when its teammates emulate (irrational) human players, suggesting that we might be able to learn from how the DRL agent plays the game. We demonstrate using a computerized beer game that we developed in collaboration with Opex Analytics (now Coupa).

BioSketch: Larry Snyder is a Professor of Industrial and Systems Engineering and Director of the Institute for Data, Intelligent Systems, and Computation (I-DISC) at Lehigh University in Bethlehem, PA. He received his Ph.D. in Industrial Engineering and Management Sciences from Northwestern University. Dr. Snyder’s research interests include modeling and solving problems in supply chain management and energy systems, particularly when the problem exhibits significant amounts of uncertainty. He is co-author of the textbook Fundamentals of Supply Chain Theory, published in 2011 by Wiley, which won the IIE/Joint Publishers Book-of-the-Year Award in 2012; a second edition was published in 2019. He has served on the editorial boards of Transportation Science, IIE Transactions, OMEGA, and the Wiley Series on Operations Research and Management Science. He previously served as a Senior Research Fellow—Optimization for Opex Analytics. For more information, visit coral.ise.lehigh.edu/larry.



Jiankun Sun, Imperial College London, UK

"Predicting Human Discretion to Adjust Algorithmic Prescription: A Large-Scale Field Experiment in Warehouse Operations"

Abstract: Conventional optimization algorithms that prescribe order packing instructions (which items to pack in which sequence in which box) focus on box volume utilization yet tend to overlook human behavioral deviations. We observe that packing workers at the warehouses of Alibaba Group deviate from algorithmic prescriptions for 5.8% of packages, and these deviations increase packing time and reduce operational efficiency. We posit two mechanisms and demonstrate that they result in two types of deviations: (1) information deviations stem from workers having more information and in turn better solutions than the algorithm; and (2) complexity deviations result from workers' aversion, inability or discretion to precisely implement algorithmic prescriptions. We propose a new "human-centric bin packing algorithm" that anticipates and incorporates human deviations to reduce deviations and improve performance. It predicts when workers are more likely to switch to larger boxes using machine learning techniques and then pro-actively adjusts the algorithmic prescriptions of those "targeted packages." We conducted a large-scale randomized field experiment with the Alibaba Group. Orders were randomly assigned to either the new algorithm (treatment group) or Alibaba's original algorithm (control group). Our field experiment results show that our new algorithm lowers the rate of switching to larger boxes from 29.5% to 23.8% for targeted packages and reduces the average packing time of targeted packages by 4.5%. This idea of incorporating human deviations to improve optimization algorithms could also be generalized to other processes in logistics and operations.

BioSketch: Jiankun Sun is an Assistant Professor of Operations Management at Imperial College Business School, Imperial College London. Her research interest is in digital platform operations, especially how digitalization and artificial intelligence reshapes operations in an organization and impacts consumer behavior. She applies both data analytics and theoretical modeling techniques to study practical problems in digital platform operations. Jiankun Sun obtained her Ph.D. in Operations Management from Kellogg School of Management, Northwestern University, and her B.E. in Industrial Engineering from Tsinghua University.



Martin Takac, Lehigh University / Mohamed bin Zayed University of Artificial Intelligence, UAE

"Data-driven stochastic Vehicle Routing Problem and Job Shop using Reinforcement Learning"

Abstract: Reinforcement Learning (RL) is a subfield of machine learning that focuses on sequential decision making. The RL agent can be trained to maximize reward/minimize cost in stochastic environments. For example, in the stochastic Vehicle Routing Problem, the demand for items and travel times could be stochastic, not necessary IID. In this talk, we focus on an entirely data-driven approach to train RL. We will assume that a few

external features will drive the randomness (e.g., demand for ice cream is influenced by outside temperature, precipitation, and many other factors). We conclude the talk with preliminary experiments that demonstrate the superior performance of the RL approach when compared to many state-of-the-art heuristics.

BioSketch: Martin Takac is an Associate Professor and deputy department chair of Machine Learning department at Mohamed bin Zayed University of Artificial Intelligence (MBZUAI), UAE. Before joining MBZUAI, he was an Associate Professor in the Department of Industrial and Systems Engineering at Lehigh University, where he has been employed since 2014. He received his B.S. (2008) and M.S. (2010) degrees in Mathematics from Comenius University, Slovakia, and Ph.D. (2014) degree in Mathematics from The University of Edinburgh, United Kingdom. He received several awards during this period, including the Best Ph.D. Dissertation Award by the OR Society (2014), Leslie Fox Prize (2nd Prize; 2013) by the Institute for Mathematics and its Applications, and INFORMS Computing Society Best Student Paper Award (runner up; 2012). His current research interests include the design and analysis of algorithms for machine learning, applications of ML, optimization, HPC. Martin received funding from various U.S. National Science Foundation programs, including through a TRIPODS Institute grant awarded to him and his collaborators at Lehigh, Northwestern, and Boston University. He currently serves as an Associate Editor for Mathematical Programming Computation, Journal of Optimization Theory and Applications, and Optimization Methods and Software and is an area chair at machine learning conferences like ICML, NeurIPS, ICLR, and AISTATS.



Barrett Thomas, University of Iowa

"[Same-Day Delivery with Fair Customer Service](#)"

Abstract: The demand for same-day delivery (SDD) has increased rapidly in the last few years and has particularly boomed during the COVID-19 pandemic. The fast growth is not without its challenge. In 2016, due to low concentrations of memberships and far distance from the depot, certain minority neighborhoods were excluded from receiving Amazon's SDD service, raising concerns about fairness. In this paper, we study the problem of offering fair SDD-service to customers. The service area is partitioned into different regions. Over the course of a day, customers request for SDD service, and the timing of requests and delivery locations are not known in advance. The dispatcher dynamically assigns vehicles to make deliveries to accepted customers before their delivery deadline. In addition to overall service rate (utility), we maximize the minimal regional service rate across all regions (fairness). We model the problem as a multi-objective Markov decision process and develop a deep Q-learning solution approach. We introduce a novel transformation of learning from rates to actual services, which creates a stable and efficient learning process. Computational results demonstrate the effectiveness of our approach in alleviating unfairness both spatially and temporally in different customer geographies. We also show this effectiveness is valid with different depot locations, providing businesses with opportunity to achieve better fairness from any location. Further, we consider

the impact of ignoring fairness in service, and results show that our policies eventually outperform the utility-driven baseline when customers have a high expectation on service level.

BioSketch: Barry Thomas is the Senior Associate Dean of the Tippie College of Business at the University of Iowa and the Gary C. Fethke Research Professor of Business Analytics. He previously served as the Departmental Executive Officer (Department Head) in the Department of Business Analytics at the University of Iowa. Barry received his PhD and MS in Industrial and Operations Engineering from the University of Michigan and holds BA degrees in both Mathematics and Economics from Grinnell College. He has over 50 peer-reviewed research publications, many focused on the application of machine learning to last-mile logistics problems. His research has been sponsored by the National Science Foundation and private industry. Barry is currently co-Area Editor for the Routing and Logistics Area of the journal Transportation Science. He has previously served as President of the INFORMS Transportation and Logistics Society, an association of over 1000 transportation science researchers, and as a Vice President for INFORMS, the largest association of analytics and operations research professionals in the world. After 15 years, Barry recently retired from the Grinnell College's Board of Trustees. As a Trustee, he had served as Chair of the Audit & Assessment and Governance Committees and had been Vice Chair of the Board.



Michael Watson, Coupa Software

"AI/ML Applications for Industry: The Proven and Future Direction"

Abstract: This talk will cover how companies using AI/ML to improve their supply chains. We'll highlight this with plenty of industry examples. We'll also cover a new trend we are seeing—using community data. In the consumer world, community data shows up in apps like Waze—everyone shares their data and gets better driving directions. This trend is moving to the business community, unlocking new supply chain value in risk management and understanding costs. To implement AI solutions, it is also important to define what we mean with the term AI and how companies

are building AI/ML teams.

BioSketch: Mike is VP of AI for Coupa Software. Prior to Coupa, he was the co-founder and CEO of Opex Analytics, an AI firm that grew to 140 people from 2013 to 2019 when it was acquired by LLamasoft. At LLamasoft he was the General Manager of the Opex Analytics Business Unit. LLamasoft was acquired by Coupa in late 2020. He is also an adjunct professor at Northwestern University. At Northwestern, he teaches a class on Operations Management (in the Masters in Engineering program) and one on Optimization (in the Masters in Analytics program). He is the lead author of the books "Managerial Analytics" and "Supply Chain Network Design." He holds a PhD from Northwestern University in Industrial Engineering and Management Sciences. (<https://www.linkedin.com/in/michael-watson-07600a1>)
