Future of OR and Analytics Workshop

Sat 10/15, JW Marriott Indianapolis, White River Ballroom GH

9:00	– 10:30	Session 1: Advances in Robust optimization
10:30	– 11:00	Break
11:00	– 12:30	Session 2: Analytics for a Better World
12:30	– 1:30	Lunch break
1:30	– 3:00	Session 3: Healthcare Analytics
3:00	– 3:30	Break
3:30	– 5:00	Session 4: Machine Learning under a Modern Optimization Lens
5:00	– 5:30	Keynote: Dimitris Bertsimas

Keynote: The Future of OR and Universities

Dimitris Bertsimas

Dimitris Bertsimas is the Boeing Professor of Operations Research and the Associate Dean of Business Analytics at the Sloan School of Management at the Massachusetts Institute of Technology. He is a member of the US National Academy of Engineering, an INFORMS fellow, recipient of the John von Neumann Theory Prize, the Frederick W. Lanchester Prize, the Erlang Prize, finalist of the Franz Edelman Prize four times, and the INFORMS President's Award, among many other research and teaching awards, supervisor of 88 completed and 25 current doctoral theses, co-author of seven graduate level textbooks, 300+ publications, editor of the *INFORMS Journal on Optimization* and co-founder of ten analytics companies and two nonprofit foundations.

Schedule at a glance (morning)

Flash talks are indicated by “[F]” preceding the title.

Session 1: Advances in Robust Optimization (9:00 – 10:30)

9:00	Amedeo Odoni	Advances in Air Transport and Traffic Control
9:15	Vineet Goyal	The Affine Affair
9:30	Chaithanya Bandi	Federated Robust Optimization
9:45	Bradley Sturt	On the Sparsity of Optimal Linear Decision Rules in Robust Optimization
10:00	Bartolomeo Stellato	Mean Robust Optimization
10:15	Aurélie Thiele	Robust Optimization Approaches for Sustainable Energy Infrastructure

Session 2: Analytics for a Better World (11:00 – 12:30)

11:00	Apostolos Fertis	Inspiring the Introduction of Operations Research in the Art of Policy Making
11:15	Phebe Vayanos	Learning Policies for Allocating Scarce Housing Resources to People Experiencing Homelessness
11:30	Dan Iancu	Analytics for Sustainable Sourcing in Commodity Supply Chains
11:45	Andy Sun	Analytics and OR in the Electric Energy Systems Transition
12:00	Margret Bjarnadottir	Closing the Pay Gap with Analytics
12:15	Léonard Boussioux	[F] Multimodality, Models, Algorithms, and Applications to Environmental Sustainability
12:22	Vassilis Digalakis	[F] Decarbonizing OCP

Schedule at a glance (afternoon)

Session 3: Healthcare Analytics (1:30 – 3:00)

1:30	Michael Lingzhi Li	Analytics for High-stakes Applications
1:45	Liangyuan Na	[F] Transforming the Future of Healthcare Operations - Holistic Hospital Optimization at Hartford HealthCare
1:52	Wes Gurnee	[F] Strategic Priorities in Health Security
2:00	Kimberly Villalobos Carballo & Yu Ma	[F] TabText: a Systematic Approach to Aggregate Knowledge Across Tabular Data Structure
2:07	Cynthia Zeng	[F] Holistic Chronic Disease Risk Management using Medical Claims Data
2:15	Periklis Petridis & George Margaritis	Automated Data Extraction for Clinical Databases using Natural Language Processing
2:30	Benjamin Boucher	[F] Optimizing Surgical Blocks: Flexibility, Robustness, and Personalized Care
2:37	Nicholas Johnson	[F] Sparse Plus Low Rank Matrix Decomposition: A Discrete Optimization Approach
2:45	John Silberholz	Combining Pre-approval Clinical Trials And Post-approval Spontaneous Adverse Event Reporting For Improved Safety Signaling

Session 4: Machine Learning under a Modern Optimization Lens

3:30	Nathan Kallus	Learning and Optimization: Separate or Integrated?
3:45	Karthik Natarajan	Optimization with Marginals
4:00	Yannis Paschalidis	Distributionally Robust Multiclass Classification: Robustifying Deep Neural Network Image Classifiers
4:15	Dessislava Pachamanova	Acceptable Set Topic Modeling
4:30	Suleeporn (Yui) Sujichantararat	[F] Simulated Annealing for Classification Trees
4:37	Elijah Pivo	[F] An Optimization Approach to Causal Inference
4:45	Zhen Lin	[F] Towards Improving the Interpretability of Deep Learning
4:52	Angelos Koulouras	[F] Pricing of Resources and Cost Allocation in MIO Problems: a Local Convex Hull Approach

Session 1: Advances in Robust Optimization

Advances in Air Transport and Traffic Control (provisional title)
Amedeo Odoni
Massachusetts Institute of Technology

Abstract TBA

The Affine Affair
Vineet Goyal
Columbia University

I will present some important vignettes from the stream of work analyzing the performance of affine policies for adjustable robust optimization that have fundamentally improved our understanding of why these work well in practice. These have been inspired from Dimitris' past work and collaborations with him.

Federated Robust Optimization
Chaithanya Bandi
National University of Singapore

Federated learning has emerged as a leading machine learning technique that trains an algorithm across multiple decentralized edge devices or servers holding local data samples, without exchanging them. Motivated by a similar setting, we consider the hub-and-spoke model of distributed optimization in which a central authority coordinates the computation of a solution among many agents while limiting communication. In particular, we consider the case where the central authority is interested in min-max optimization. In this talk, we introduce the framework of Federated Robust Optimization (FedRO) and present preliminary results. Our results shows that these methods are provably robust to inexact computation of intermediate local quantities. We apply our approach in the context of a leading fashion e-retailer in South-east Asia and demonstrate the benefits of our methods in practice.

On the Sparsity of Optimal Linear Decision Rules in Robust Optimization
Bradley Sturt
University of Illinois at Chicago

In this talk, we consider the widely-studied class of production-inventory problems from the seminal work of Ben-Tal et al. (2004) on linear decision rules in robust optimization. We prove that there always exists an optimal linear decision rule for this class of problems in which the number of nonzero parameters in the linear decision rule is equal to a small constant times the number of parameters in a static decision rule. This result demonstrates that the celebrated performance of linear decision rules in such robust inventory management problems can be

obtained without sacrificing the interpretability of static decision rules. From a practical standpoint, our result lays a theoretical foundation for the growing stream of literature on harnessing sparsity to develop practicable algorithms for computing optimal linear decision rules in operational planning problems with many time periods.

Mean Robust Optimization

Bartolomeo Stellato
Princeton University

Robust optimization is a tractable and expressive technique for decision-making under uncertainty, but it can lead to overly conservative decisions when pessimistic assumptions are made on the uncertain parameters. Wasserstein distributionally robust optimization can reduce conservatism by being data-driven, but it often leads to large problems with prohibitive solution times. We introduce mean robust optimization, a general framework that combines the best of both worlds by providing a trade-off between computational effort and conservatism. We propose uncertainty sets constructed based on clustered data rather than on observed data points directly, thereby significantly reducing problem size. By varying the number of clusters, our method bridges between robust and Wasserstein distributionally robust optimization. We show finite-sample performance guarantees and explicitly control the potential pessimism introduced by any clustering procedure. We illustrate the benefits of our framework on several numerical examples, obtaining significant computational speedups with little-to-no effect on the solution quality.

Robust Optimization Approaches for Sustainable Energy Infrastructure

Aurélie Thiele
Southern Methodist University

In this talk we present robust optimization approaches for sustainable energy infrastructure. We develop robust portfolio models of clean and renewable energy production over time to meet local and state governments' short-term and long-term clean energy goals, incorporating concepts from project finance to finance the building of critical new infrastructure in the presence of demand uncertainty and price uncertainty. We model and mitigate a wide range of risks both pre-completion and post-completion to create a comprehensive strategy to achieve renewable energy goals.

Session 2: Analytics for a Better World

Inspiring the Introduction of Operations Research in the Art of Policy Making

Apostolos Fertis
Government of the Hellenic Republic

Evidence-informed policy making has come recently into focus, given the severe and diverse challenges that humanity faces, ranging from climate change, energy efficiency and natural disasters to pandemics and population aging. The design and implementation of informed policies based on observed facts and scientific reasoning is attracting the attention of multiple policy making and scientific communities, that can see clearly the benefit of improving the mechanisms of policies to avoid taking ad hoc decisions that only temporarily solve problems. During our recent experience in the General Secretariat of Coordination within the Government of the Hellenic Republic, we implemented a new set of methodologies for policy planning, monitoring and evaluation, as well as promoted the use of data in decision making in multiple cases, as for example the fight against Covid-19. This experience has convinced us that Operations Research and Quantitative Analytics can offer much in this direction, especially because they encompass a structured approach for tackling problems and discovering solutions that has been proved fruitful in a wide spectrum of areas. Concepts like quantitative modelling of complex systems, optimal design and decision making, data uncertainty, robustness and fairness are at the core of thinking for the Operations Research community. Their incorporation into policy making can be insightful for creating new innovative methodologies to design efficient policies.

Learning Policies for Allocating Scarce Housing Resources to People Experiencing Homelessness

Phebe Vayanos
University of Southern California

We study the problem of allocating scarce housing resources of different types to individuals experiencing homelessness based on their observed covariates. We leverage administrative data collected in deployment to design an online policy that maximizes mean outcomes while satisfying budget and fairness requirements. We propose a policy in which an individual receives the resource maximizing the difference between their mean treatment outcomes and the resource bid price, or roughly the opportunity cost of using a resource. Our approach has nice asymptotic guarantees and is easily interpretable. We show results on real data from the Homeless Management Information System.

Analytics for Sustainable Sourcing in Commodity Supply Chains

Dan Iancu
Stanford University

We discuss recent work dealing with problems emerging in global commodity supply chains (including wool, palm oil, cocoa, and beef) where data analytics and a structured approach to decision-making have shown great promise.

Analytics and OR in the Electric Energy Systems Transition

Andy Sun
Massachusetts Institute of Technology

Electric energy systems around the world are experiencing rapid changes -- from fossil fuel centered systems with largely passive load to clean-power centered systems with active demand and increasing electrification of transportation and industrial processes. In this talk, we will give some snapshots of challenges in this transition, where analytics and OR could offer its help.

Closing the Pay Gap with Analytics

Margret Bjarnadottir
University of Maryland

Pay equity and demographic pay gaps are a topic of increasing discussion in the boardroom, in the media and amongst policymakers. Despite this increased interest, there has been a lack of methodology to ensure pay equity at scale. Using analytics rooted in optimization and statistics, we derive an approach to close demographic pay in the fairest way possible. We show that by intelligently increasing the wages of selected workers, we can not only close pay gaps but correct underlying biases. We also discuss some of the unintended consequences of applying optimization to close the gap without considerations of the pay-structure, fairness, or operational constraints. The approach has been implemented around the world impacting hundreds of thousands of people.

[Flash talk] Multimodality, Models, Algorithms, and Applications to Environmental Sustainability

Léonard Boussioux
Massachusetts Institute of Technology

This talk illustrates how combining optimization, machine learning, and deep learning can help solve sustainability challenges. We investigate how to transform high-dimensional, messy, and unstructured data into impactful predictive and prescriptive analytics. We propose a novel data-driven framework - Gather, Extract, Predict - that generalizes state-of-the-art machine learning and deep learning tools into efficient and practical feature extractors. We develop data pre-processing, mining, and fusion pipelines to process large-scale multimodal data sources such as text, images, time series, and tables. We present exciting opportunities to apply the framework to ecosystem modeling and understand biodiversity's resilience to climate change.

[Flash talk] Decarbonizing OCP

Vassilis Digalakis
Massachusetts Institute of Technology

We present our collaboration with OCP, one of the world's largest producers of phosphate and phosphate-based products, in support of a green initiative designed to significantly reduce the company's greenhouse gas emissions. We study the problem of decarbonizing OCP's electricity supply by installing a mixture of solar panels and batteries to minimize their time-discounted investment cost plus the cost of satisfying their remaining demand via the national grid. We adopt a data-driven robust optimization approach which averages across uncertainty sets constructed from historical solar capacity factors and prevents overfitting through distributional robustness. Our model scales to solve problems with tens of millions of decision variables in minutes or hours without resorting to decomposition schemes. The proposed methodology reduces by over 60% the emissions which arise from OCP's energy needs while lowering their time-adjusted operational costs. Further, we show that the proposed investment and operational strategies are robust to perturbations in the data, guaranteeing that the decarbonization process will be profitable even if OCP's estimates of future fertilizer demand or future electricity prices are inaccurate.

Session 3: Healthcare Analytics

Analytics for High-stakes Applications

Michael Lingzhi Li
Harvard University

We present a brief overview of the unique challenges in utilizing analytics and data science for applications where the cost of failure is large, including healthcare, public policy, finance, and more. In particular, the challenges include interpretability, stability, and causality. We discuss a few research directions that are aimed to solve these challenges.

[Flash talk] Transforming the Future of Healthcare Operations - Holistic Hospital Optimization at Hartford HealthCare

Liangyuan Na
Massachusetts Institute of Technology

Healthcare systems strive to improve patient outcomes with financial and operational burdens as well as the stress of meeting uncertain patient demand. Opportunities have arisen to transform the future of hospital operations with analytics and operations research. We present Holistic Hospital Optimization (H2O) to support strategic, tactical, and operational decision making for key stakeholders in hospital systems. We develop a system of practical predictive and optimization models that unify staff scheduling, patient flows, and bed assignments. Besides creating models, our goal is also to implement the innovations in the real world. We work closely with Hartford HealthCare, the largest hospital network in Connecticut. We deploy an end-to-end software in production, currently benefiting hundreds of users and supporting daily operational decisions at all 7 hospitals at Hartford HealthCare. The implementation projects annual revenue uplift of millions of dollars, and demonstrates significant potential to scaling to other hospital systems in the world.

[Flash talk] Strategic Priorities in Health Security

Wes Gurnee
Massachusetts Institute of Technology

The COVID-19 pandemic has illustrated the consequences of failing to contain a highly-contagious infectious disease in the modern era. Despite these costs, comparatively little has been done to address the underlying strategic gaps in pandemic prevention and preparedness. This deficiency is made even more alarming when considering the rapid advances in biotechnology enabling the creation of engineered pathogens of existential severity. While addressing these risks requires a radically multidisciplinary approach, OR and analytics should play a central role in pandemic prevention and preparedness. Example applications include designing an early-warning observation network, better disease models and forecasts, facility location and inventory management of the national strategic stockpile, vaccine {manufacturing, allocation, and distribution}, and hardening other critical supply chains for

testing, PPE, etc. Such problems require data-driven, robust, adaptive, and prescriptive solutions to shift the strategic balance towards the defensive.

[Flash talk] TabText: a Systematic Approach to Aggregate Knowledge Across Tabular Data Structure

Kimberly Villalobos Carballo
Massachusetts Institute of Technology

Despite recent advancements in the analysis of images, videos and natural language, tabular data still remains the most prominent format in today's data intensive environments. Yet, the lack of a systematic approach to combine, process and analyze tabular information effectively poses a major challenge in the machine learning community. In this work, we develop TabText, a unified framework to represent and standardize tabular data as language. This methodology leverages the flexibility of language to represent data without imposing any structural constraints, and thus provides a new perspective to encode tabular data from different table structures and time periods. Across two different healthcare datasets, we show that features extracted via TabText outperform those extracted with traditional processing methods by 1-2%. We also analyze the sensitivity of our framework against different choices for sentence representations of missing values, meta information, and language descriptiveness. The generalizability and simplicity of our framework could potentially offer a promising gateway for the integration of tabular data in various application settings such as healthcare.

[Flash talk] Holistic Chronic Disease Risk Management using Medical Claims Data

Cynthia Zeng
Massachusetts Institute of Technology

According to the CDC, 90% of the US annual health care expenditure is spent on chronic and mental health conditions. Therefore, identification of individuals at risk is an important challenge to allow early interventions to reduce costs. In this work, we developed a machine learning approach using medical claims data to develop personal risk scores for major chronic diseases. Using data of more than 60 thousand unique individuals' in the US, we conducted experiments to systematically predict the onset risk for 50 major diseases for the next 1-year, 3-year, 5-year time horizon. Some most predictable diseases include Type II Diabetes, achieving 0.76-0.82 ROC-AUC score, and Hypertension, achieving 0.75-0.79 for different time horizons. Our work demonstrates the potential of leveraging medical claims data to manage chronic diseases in public health.

Automated Data Extraction for Clinical Databases using Natural Language Processing

Periklis Petridis and George Margaritis
Massachusetts Institute of Technology

Cardiothoracic surgery departments and clinics in the US rely on central agencies, such as the Society of Thoracic Surgeons (STS), to evaluate their operational performance compared to their peer institutions. Specifically, 97% of U.S. adult cardiac surgery programs collaborate and transfer their patient data to the STS National Database for quality improvement and risk assessment. However, as patient records primarily consist of unstructured text reports, such data transfers involve a large operational overhead, requiring manual data extraction by teams of experienced data managers. Utilizing recent breakthroughs in Natural Language Processing, we propose an end-to-end machine learning pipeline that automatically extracts all patient data (either structured or unstructured) from multiple sources, over multiple patient visits, and for multiple target outcomes. Our pipeline can be extended to new sources or outcomes, as hospitals may use different data conventions from each other. Preliminary results on Massachusetts General Hospital data show promise and our methodology achieves up to 98% AUC in common diagnoses and up to 85-95% in more challenging ones. We believe using an automated pipeline for data extraction has four benefits: i) reduce operational overhead and costs for institutions when transferring their data, ii) increase data consistency and quality, while reducing variation from human errors, iii) allow institutions to transfer even richer data that would otherwise require additional human effort, and iv) create a unified framework for end-to-end medical predictions and diagnoses from any type and format of patient medical reports.

[Flash talk] Optimizing Surgical Blocks: Flexibility, Robustness, and Personalized Care

Benjamin Boucher
Massachusetts Institute of Technology

Hospital beds are a crucial resource for the healthcare system - both in the costs they entail and in the key role they play in a patient's recovery. Flattening the census of recovering patients in hospital beds effectively reduces the number of beds used, and would increase the hospital's capacity for treating patients without expensive investments. We propose a mixed integer program to schedule and staff Operating Blocks while flattening the census of recovering patients post-surgery. We adopt a novel, robust approach to the problem to tackle the uncertainty of the times of rests of patients. Our approach respects the personal patient-provider relation without changing the usual work hours of a surgeon, though both can be individually relaxed to enhance results.

[Flash talk] Sparse Plus Low Rank Matrix Decomposition: A Discrete Optimization Approach

Nicholas Johnson
Massachusetts Institute of Technology

We study the Sparse Plus Low Rank decomposition problem (SLR), which is the problem of decomposing a corrupted data matrix into a sparse matrix containing the perturbations plus a low rank matrix. SLR is a fundamental problem in Machine Learning arising in many applications such as data compression, latent semantic indexing and collaborative filtering. We

introduce a novel formulation for SLR that directly models the underlying discreteness of the problem. For this formulation, we develop an alternating minimization heuristic to compute high quality solutions and a novel semidefinite relaxation that provides meaningful bounds. We further develop a custom branch and bound routine that solves small instances of SLR to certifiable near-optimality. Our heuristic can scale to n=10000 in hours, our relaxation can scale to n=200 in hours, and our branch and bound algorithm can scale to n=25 in minutes. Our numerical results demonstrate that our approach outperforms existing state-of-the-art approaches.

Combining Pre-approval Clinical Trials And Post-approval Spontaneous Adverse Event Reporting For Improved Safety Signaling

John Silberholz
University of Michigan

A classical question in pharmacovigilance is how to combine pre-approval RCTs and post-approval surveillance data to increase the power for side effect detection. A key step is to learn the degree to which the observational data is biased before one can combine it with unbiased clinical trial data. In this work, we propose a model that uses information about common toxicities to help de-bias the observational data on rare toxicities through the correlation of bias among different toxicities (e.g., correlation due to co-reported drugs, indications, and patient health). Using Bayesian statistics, we analyze the benefit of "cross"-debiasing and identify the situation where such benefit is largest. Numerical experiments using real data from the FDA Adverse Event Reporting System (FAERS) suggest significant values of using cross-debiasing to improve drug safety signaling.

Session 4: Machine Learning under a Modern Optimization Lens

Learning and Optimization: Separate or Integrate?

Nathan Kallus
Cornell University

Predictive side information (aka "context") makes optimization under uncertainty less uncertain, but leveraging it requires we learn a potentially complex predictive relationship. We might use off-the-shelf ML methods for this, but their training process ignores the downstream optimization task where the model would be plugged in. Instead, we can train the model end-to-end to optimize the downstream costs of the decisions it would induce. In this talk I will tackle the question, which is better? That is, should we separate or integrate the learning and optimization tasks? I show that in linear problems, where we only care to learn the mean-prediction function, the naive separated approach has regret that surprisingly vanishes orders faster than end-to-end methods -- a consequence of real problems not having arbitrarily bad near-dual-degeneracy. However, for general (nonlinear) contextual optimization, which involves the hard-to-learn conditional-probability prediction function, it may be better to integrate the tasks and directly learn a decision policy. I show how to do this tractably near-optimally for tree-based policies and build scalable decision forests.

Optimization with Marginals

Karthik Natarajan
Singapore University of Technology and Design

I will discuss two applications of optimization problems where uncertainty plays an important role. In these applications, by focusing on just the marginal distributions and accounting for extremal dependence among the uncertainties, we obtain useful predictive and prescriptive performance. The two applications are in discrete choice models and influence maximization.

Distributionally Robust Multiclass Classification: Robustifying Deep Neural Network Image Classifiers

Yannis Paschalidis
Boston University

I will present a Distributionally Robust Optimization (DRO) formulation for Multiclass Logistic Regression (MLR). The DRO framework uses a probabilistic ambiguity set defined as a ball of distributions that are close to the empirical distribution of the training set in the sense of the Wasserstein metric. We relax the DRO formulation into a regularized learning problem whose regularizer is a norm of the coefficient matrix. We establish out-of-sample performance guarantees for the solutions to our model, offering insights on the role of the regularizer in controlling the prediction error. We apply the proposed method in rendering deep Vision

Transformer (ViT) and Convolutional Neural Network (CNN)-based image classifiers robust to both random and even adversarial attacks. (joint work with Ruidi Chen and Boran Hao)

Acceptable Set Topic Modeling

Dessislava Pachamanova
Babson College

Topic modeling is a significant branch of natural language processing and machine learning focused on inferring the generative process of text. Traditionally, algorithms for estimating topic models have relied on Bayesian inference and Gibbs sampling. We propose a novel “acceptable set” framework for formulating topic modeling problems inspired by ideas from discrete component analysis and data-driven robust optimization. Our approach not only simplifies the design and inference of topic models, but also allows for extensions and generalizations that are challenging to integrate into traditional approaches, and naturally controls perplexity. Through extensive computational experiments, we show that our methods have improved solution quality compared to baseline methods and reconstruct more reliably the underlying generative models. Our framework overcomes known vulnerabilities of traditional topic modeling algorithms: our methods are effective in low-data settings, register good out-of-sample performance, and perform well for a variety of initial assumptions on input parameter values.

[Flash talk] Simulated Annealing for Classification Trees

Suleeporn (Yui) Sujichantararat
Massachusetts Institute of Technology

Optimal Classification Tree (OCT) was developed to improve on CART (Classification And Regression Trees) optimality by utilizing multiple random warm starts and local search to iteratively improve each warm start until a locally optimal classification tree is found. However, local search does not guarantee global optimality if the number of random warm starts does not exceed the number of local minima. Hence, we propose a framework utilizing simulated annealing to further improve OCT by probabilistically allowing a tree to move to a tree with a worse objective value as some worse transformations could lead to a better final model. We report computational results that show that in 61 real-world datasets from UCI machine learning repository. Simulated annealing improves the average out-of-sample Gini impurity score on the 61 datasets by 1.30% comparing to OCT and 6.25% comparing to CART.

[Flash talk] An Optimization Approach to Causal Inference

Elijah Pivo
Massachusetts Institute of Technology

In causal inference we typically quantify the average causal effect of applying a treatment to some unit. For example, we may wish to estimate the average causal effect of smoking on life expectancy or of a training program on wage. Current practice is to use flexible machine learning methods like neural networks or random forests to model the relationships between

treatment and outcome under the doubly robust framework. In this work, we consider the case where the outcome is a decision and propose using optimization problems to model the relationship between treatment and outcome. We compare the performance of machine learning methods and optimization models for quantifying average causal effects in several examples.

[Flash talk] Towards Improving the Interpretability of Deep Learning

Zhen Lin
Massachusetts Institute of Technology

We develop low depth (depth 4 to 5) classification trees with hyperplanes to closely approximate neural networks, and achieve accuracies in the mid 90's. In this way, we contribute in increasing the interpretability of neural networks. We develop an optimization-based adaptive learning method to generate additional data for training the trees, which further improves the performance of the trees. We also show that this optimization method performs better than its randomization-based variant. In addition to closely approximating the target neural networks, we also demonstrate that the trees trained as described achieve high accuracy on the original real world classification datasets used to train the target neural networks. We report comprehensive computational results on 59 real world classification datasets, with neural networks of different sizes.

[Flash talk] Pricing of Resources and Cost Allocation in MIO Problems: a Local Convex Hull Approach

Angelos Koulouras
Massachusetts Institute of Technology

Cost allocation in mixed-integer optimization problems is a popular problem in OR literature with important financial, social and political consequences. In this work, we study the problem of pricing and cost allocation in mixed-integer optimization problems with demand constraints (e.g. facility location problems) under a duality lens. A key idea of the paper is that a variety of cutting plane approaches, such as mixed-integer Gomory cuts, provide an exact representation of the “local convex hull”, that is, a polyhedral representation at optimality. So, we utilize insights from LO duality theory to provide both a theory of pricing of resources as well as cost allocation among customers or users. Finally, we test our method on formulations inspired by the unit commitment setting and provide insights based on the results.