SIGecom Job Market Candidate Profiles 2024

This is the ninth annual collection of profiles of the junior faculty job market candidates of the SIGecom community. The twenty seven candidates for 2024 are listed alphabetically and indexed by research areas that define the interests of the community. The candidates can be contacted individually, or collectively via the moderated mailing list ecom-candidates2024@acm.org.

Shortly before publishing these candidate profiles, we received the shocking news that one of the candidates from this collection, Orestis Papadigenopoulos, unexpectedly passed away. Apart from a great researcher, Orestis was a very kind, pleasant, and fun person and will be sorely missed. His profile remains in this collection in his memory, despite his untimely passing.

–Vasilis Gkatzelis and Jason Hartline

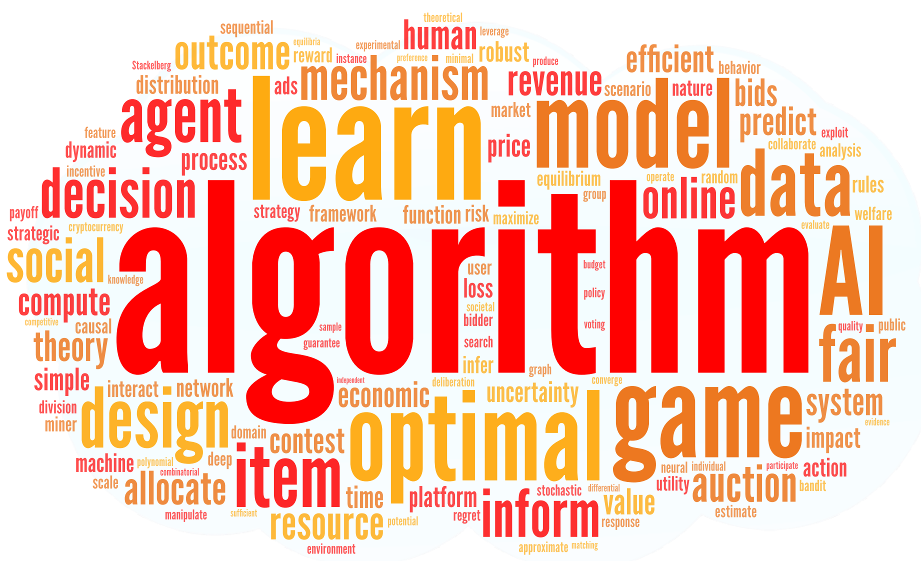

Fig. 1. Generated using the research summaries of the candidates.

Contents

YEGANEH ALIMOHAMMADI (Homepage, CV)

Thesis: Learning and Decision Making using Network Data ('24)

Advisor: Amin Saberi, Stanford University

Brief Biography: Yeganeh is a final-year Ph.D. candidate in Stanford University's Operations Research group. Before joining Stanford, she obtained a B.Sc. in Computer Engineering with a Mathematics minor from Sharif University of Technology. In 2022, she was a research fellow at the Simons Institute and interned at Google Research, hosted by Aranyak Mehta.

Research Summary: Yeganeh's research delves into the analysis of large-scale networks and stochastic systems. She employs tools from applied probability and algorithm design to tackle pressing challenges in business operations.

In today's interconnected world, networks from supply chains to social media play crucial roles. Yeganeh explores how *micro-level* behaviors in these networks cascade into *macro-level impacts*. She has provided insights into how localized data from a handful of individuals can predict epidemics [1,2]; how individual advertising strategies shape system-wide outcomes in online auctions [3]; and how adding a few drivers influences overall efficiency in ridesharing platforms [4].

Another dimension of her research is on creating *sampling algorithms* to streamline the analysis of expansive networks. Yeganeh designs these algorithms to extract a representative subset of nodes and edges that maintain the key topological characteristics of the underlying network and enable fast computation. Beyond using these for epidemic forecasting [1], she has formulated sampling algorithms for enhancing graph neural networks and sampling specific network configurations [5].

Finally, a critical component of her research involves translating theoretical insights into actionable strategies. Using network models, she developed prediction methods that informed the COVID-19 reopening strategies of Los Angeles Unified School District (LAUSD) – the nation's second-largest school district [6].

Representative Papers:

[1] Epidemic Forecasting on Networks: Bridging Local Samples with Global Outcomes (submitted to Operations Research)
with C. Borgs, R. van der Hofstad, and A. Saberi

[2] Locality of Random Digraphs on Expanders (The Annals of Probability, 2023)
with C. Borgs and A. Saberi

[3] Incentive Compatibility in the Auto-Bidding World (EC, 2023)
with A. Mehta and A. Perlroth

[4] The Value of Excess Supply in Spatial Matching Markets (EC, 2022)
with M. Akbarpour, S. Li, and A. Saberi

[5] Sequential Importance Sampling for Estimating Expectations over the Space of Perfect Matchings (The Annals of Applied Probability, 2023)
with P. Diaconis, M. Roghani, and A. Saberi

[6] Relative-Risk and the Assessment of School Safety in the COVID-19 Pandemic (Health Management, Policy, and Innovation, 2021)
with K. Shiragur, R. Johari, K. Schulman, and K. Staudenmayer

HEDYEH BEYHAGHI (Homepage, CV, Google Scholar)

Thesis: Approximately-Optimal Mechanisms in Auction Design, Search Theory, and Matching Markets ('19)

Advisor: Éva Tardos

Brief Biography: Hedyeh Beyhaghi is a postdoctoral research associate in the School of Computer Science at Carnegie Mellon University, hosted by Nina Balcan. Hedyeh received her PhD in Computer Science from Cornell University, advised by Éva Tardos. During her PhD studies, she was a long-term visitor at the Simons Institute for the Theory of Computing for the Economics and Computation program, an intern at Google, and an Ivy-Plus Exchange Scholar at Princeton University hosted by Matt Weinberg. Before joining CMU, Hedyeh was a postdoctoral research fellow at Toyota Technological Institute at Chicago (TTIC) and Northwestern University, hosted by Avrim Blum, Jason Hartline, and Samir Khuller.

Research Summary: My research primarily revolves around decision-making in the context of strategic behavior and uncertainty. This involves creating solutions when the input is derived from strategic agents, devising rules to counteract self-interested agents, and developing algorithms with limited knowledge about future events or based on limited data. Such problems are prevalent in numerous real-world situations, ranging from learning classification rules and designing auctions to online stochastic optimization.

Decision-making under strategic behavior and uncertainty poses several challenges. The solutions must ensure the agent's actions align with the overall good outcome, even when they are driven by their own self-interest. Agents may manipulate their characteristics to achieve a better outcome, and there may not be access to any unmanipulated data. The agents' behavior can vary in response to decision-making parameters. Furthermore, while the decisions made are often irreversible in a sequential scenario, the solutions need to be robust and resilient, able to perform well in unforeseen future scenarios.

To address these complex issues, I utilize a combination of techniques from mechanism design, online optimization, and learning theory. By leveraging these approaches, I design efficient algorithms that perform competently in uncertain and strategic environments. My contributions include: (1) Providing novel characterizations and solutions to online optimization problems under uncertainty that relax assumptions for more realistic, more general scenarios. (2) Successful adaptation of classic learning algorithms for strategic scenarios, including the first high-performance learning algorithms in several settings. (3) Designing algorithms with enhanced societal aspects, e.g., improving social development and fairness, in addition to traditional optimization criteria.

Representative Papers:

[1] Pandora's Problem with Nonobligatory Inspection: Optimal Structure and a PTAS (STOC 2023) with L. Cai

[2] The Strategic Perceptron (EC 2021) with S. Ahmadi, A. Blum, and K. Naggita

[3] Improved Revenue Bounds for Posted-Price and Second-Price Mechanisms (Operations Research 2021) with N. Golrezaei, R. Paes Leme, M. Pal, and B. Sivan

JOHANNES BRÜSTLE (CV, Google Scholar)

Thesis: The Competition Complexity of Online Mechanisms ('24)

Advisors: Paul Dütting (Google Research Zurich), László Végh (London School of Economics and Political Science)

Brief Biography: I am currently in my fourth year of pursuing a Ph.D. at the London School of Economics and Political Science, under the guidance of Professors Paul Dütting and László Végh. My research is supported by an LSE PhD Studentship.

Research Summary: I am broadly interested in algorithmic game theory, mechanism design and the asymptotics of online decision making. During my PhD, I have been investigating the prophet inequality setting with additional bidders for the online algorithm. Furthermore, I have been lucky to be able to work on different subjects within algorithmic game theory, such as robust mechanisms and their connection to learning as well as fair division of indivisible goods.

In [2], we study the problem of comparing the performance of the optimal posted price mechanism for single item online arrival to that of the prophet. We explore a new direction compared to the well known results in this area by giving the online algorithm an additional number of i.i.d bidders. How many such additional resources are necessary for the online algorithm to catch up to the performance of the prophet? Surprisingly, we find that we can analyze the corresponding optimization problem exactly. In a working paper [1] we find that by a very different approach, we recover the same asymptotic behavior even for the independent non-i.i.d setting.

I am also interested in learning within algorithmic game theory. More precisely, consider the classical problem of helping an auctioneer obtain large expected revenue through an auction in which the bidders have incentive to participate and to report their values truthfully for any bundles of items offered. In [3], we show that robust mechanisms allow us to push the boundary of learning (approximately) optimal multi-item auctions to the important setting of item dependence.

Another area I have worked on is fair division of indivisible goods. In [4], we want to get an upper bound on the total amount of money that has to be given to agents on top of their allocations in order to eliminate envy. We achieve the tight upper bound on total subsidy which is $(n-1)$ dollars, where the most valued item is normalized to 1 dollar.

Representative Papers:

[1] The Competition Complexity of Prophet Inequalities (working paper)
with J. Correa, P. Dütting, T. Ezra, M. Feldman and V. Verdugo

[2] The Competition Complexity of Dynamic Pricing (EC 2022, MOR 2023)
with J. Correa, P. Dütting, and V. Verdugo

[3] Multi-Item Mechanisms without Item-Independence: Learnability via Robustness (EC 2020) with Y. Cai, and C. Daskalakis

[4] One Dollar Each Eliminates Envy (EC 2020)
with J. Dippel, V.V. Narayan, M. Suzuki, and A. Vetta

LINDA CAI (Homepage, CV, Google Scholar)

Thesis: Strategic Decision Making in the Wild: Information Acquisition, Resource Augmentation and Practical Irrationality ('24)

Advisor: Matt Weinberg, Princeton University

Brief Biography: I am currently a fourth (and final) year PhD student in Computer Science at Princeton University advised by Prof. Matt Weinberg. My research interest includes a broad set of problems that relate to online optimization, incentive, learning, and their intersection. I am fortunate to be the recipient of a number of scholarships and awards, including School of Engineering and Applied Science Award for Excellence, Siebel Scholar, Chainlink Fellowship and Francis Robbins Upton Fellowship.

Research Summary: Recently, my work focuses on understanding the impact of relaxing or changing key assumptions in economic models. For instance, I have studied the effect of resource augmentation, existence of non-rational (possibly learning) agents, and the role of inspection in optimal search setting.

In [1], we study the classical setting of optimal search amidst search friction, but with the tweak that pre-selection information acquisition isn't mandatory. Our results reveal a structured optimal policy that requires minimal order adaptation, and where computational hardness arises primarily from determining the initial inspection order. We also resolve a longstanding open problem regarding the approximability of optimal utility in this setting by providing a polynomial approximation scheme (PTAS).

In the repeated mechanism design setting, my interest lies in scenarios involving not fully rational participants, such as those using learning algorithms from past interactions, and their susceptibility or resistance to exploitation by mechanism designers. For instance, in [2], we show that auctioneers can extract the maximum possible revenue (close to the optimal social welfare) against multiple unsophisticated no-regret learners, but face complexities and reduced revenue extraction when learners cap their bids to item values.

I have also investigated the impact of resource augmentation in the traditional one-shot auction design setting. While a complex auction may extract more revenue from a fixed bidder pool, its complexity might dissuade participation. The auctioneer hence need to consider the trade-off between simplicity and optimality. Our work in [3] shows that any constant revenue competitive simple auction, when scaled with a constant multiplicative increase in bidders, surpasses the original optimal revenue, except when another simple auction (second-price auction) already achieves near-optimal revenue.

Representative Papers:

[1] Pandora's Problem with Nonobligatory Inspection: Optimal Structure and a PTAS. (STOC 2023) with Hedyeh Beyhaghi

[2] Selling to Multiple No-Regret Buyers (WINE 2023)
with S. Matthew Weinberg, Evan Wildenhain and Shirley Zhang

[3] 99% Revenue with Constant Enhanced Competition. (EC 2021)
with Raghuvansh R. Saxena

YURONG CHEN (Homepage, CV, Google Scholar)

Thesis: Games of Private Information over Learning Agents ('24)

Advisor: Xiaotie Deng, Peking University

Brief Biography: Yurong is a final-year Ph.D. candidate at Peking University, advised by Xiaotie Deng. During grad school, she visited the University of Hong Kong in Spring 2023, hosted by Zhiyi Huang. Before that, she completed her B.S. degree in mathematics in Hua Luogeng Honors Class at Beihang University, where she achieved top rank in class. Yurong is a recipient of the Best Student Paper Award at WINE 2022.

Research Summary: Yurong mainly focused on understanding, learning, and computing how agents make use of their information advantage during interaction with learning agents, from a game theorist's point of view. For example, machine learning enables the extraction of real-world game parameters from data, e.g., strategy sets and payoff functions. This makes game theory possible to be applied to real-world scenarios. However, agents with information advantage can provide fake data, manipulating the learning results and game outcomes to benefit themselves. Players are already gaming with their private information during the learning process before the learned game starts. It even remains a question whether machine learning can produce credible results under such manipulation. Therefore, studying how agents utilize their information advantage is crucial for both learning and game theory to be more robustly applied to real-world situations.

Private Information Misreporting in Stackelberg equilibria [1,3]. [1] studies how the follower optimally misreports his payoff function in extensive-form games. In four settings considered, they fully characterize all the game outcomes inducible through the follower's misreporting payoffs, and provide polynomial-time algorithms to find the optimal way of manipulation. [3] studies the query version of this problem in bimatrix games: the follower does not know the leader's payoff function, but has to learn to misreport optimally. They show that a polynomial number of queries and operations is sufficient for optimally misreporting.

Coordinated Auto-Bidding in Online Ad Auctions [2]. In online ad auctions, advertisers are delegating bidding tasks to bidding agencies. While most studies on online bidding focus on bidding for one bidder in one campaign, the bidding agency actually has information on all her clients and can coordinate bids to benefit everyone. [2] proposes coordinated online bidding strategies in repeated second-price auctions with budgets that theoretically and experimentally guarantee everyone a better utility than the best they get under independent bidding.

Representative Papers:

[1] Optimal Private Payoff Manipulation against Commitment in Extensive-form Games (WINE 2022, Best Student Paper)
with X. Deng, and Y. Li

[2] Coordinated Dynamic Bidding in Repeated Second-Price Auctions with Budgets (ICML 2023)
with Q. Wang, Z. Duan, H. Sun, Z. Chen, X. Yan and X. Deng

[3] Learning to Manipulate a Commitment Optimizer (Working Paper)
with X. Deng, J. Gan, and Y. Li

MICHAEL CURRY (Homepage, CV, Google Scholar)

Thesis: Learning and Robustness With Applications to Mechanism Design ('22)

Advisor: John Dickerson & Tom Goldstein, University of Maryland

Brief Biography: Michael Curry is a postdoc splitting time between Sven Seuken's group at the University of Zurich and David Parkes' group as Harvard. His research is at the intersection of machine learning and mechanism design, with a particular focus on using machine learning techniques to search through complex spaces of mechanisms and strategies.

Research Summary:

It's long been observed that mechanism design given samples from the valuation distribution is in essence a learning problem. And mechanisms are just functions, so why not represent them using the computational tools and rich function approximators that modern deep learning provides? This is the pitch of differentiable economics, a recent thread of work introduced in a number of papers, notably "Optimal Auctions Through Deep Learning" by Dütting et al., which trains mechanisms represented directly as deep neural networks.

One thread of my research on differentiable economics has involved overcoming a significant limitation of these approaches – they are not perfectly strategyproof. In our paper "Certifying Strategyproof Auction Networks", we adapt neural auction architectures to make them certifiable – we can *exactly* compute the degree to which strategyproofness is really violated. In our IJCAI-2023 paper "Differentiable Economics for Randomized Affine Maximizer Auctions", we present a new and high-performing architecture that is also perfectly strategyproof for any setting of the parameters. (This paper has itself been further improved upon by a separate group from ours in "A Scalable Neural Network for DSIC Affine Maximizer Auction Design" by Duan et al., appearing this year at NeurIPS.)

Other work with collaborators has included devising a practical method for identifying only the most important transplants for pre-screening in paired kidney exchange. Additionally, while interning at Salesforce Research, I worked with the AI Economist team on a project to use GPU-accelerated simulations to model a simple economy involving hundreds of learning agents.

In ongoing work at Sven Seuken's group, we aim to find high-performing *dynamic* mechanisms in very general MDP settings; in a separate project, we also use RL to find good bidding strategies in iterative combinatorial auctions. And in ongoing work with David Parkes and Zhou Fan, we study a particular case of multi-good automated market making under adverse selection. We have shown that in such settings, the problem of finding a profit-maximizing automated market maker is dual to an optimal transport problem, and that differentiable economics can be used to search for optimal mechanisms.

Representative Papers:

[1] Differentiable Economics for Randomized Affine Maximizer Auctions (IJCAI '23) with M. Curry, J. Dickerson, T. Sandholm

[2] Certifying Strategyproof Auction Networks (NeurIPS '20) with M. Curry, P.Y. Chiang, T. Goldstein, and J. Dickerson

SULAGNA DASGUPTA (Homepage, CV)

Thesis: Screening Knowledge ('24)

Advisor: Ben Brooks, University of Chicago

Brief Biography: Sulagna Dasgupta is a final year PhD candidate at the University of Chicago. She is a microeconomic theorist with interests in mechanism design and information economics. She has three broad, interconnected research agendas – (A) screening knowledge, (B) the interaction of hard and soft information in strategic communication and information design, and (C) matching theory.

Research Summary: Tests of knowledge are ubiquitous – job interviews, exams, standardized tests etc. The goal of such tests is to discern some underlying quality of the test-taker, which enables them to do well on the test. My job market paper asks: How to optimally design such tests to maximize this learning? To answer it, I set up a model of screening knowledge. In order to focus on the forces of incentive new to this problem, I use its most elemental version, where the subject matter is simply a binary fact. I show that even though the test-designer can set up any complicated grading scheme by rewarding test-takers according to the strength of their conviction in the correct answer, optimally, the tests take the simple True/False or True/False/"I don't know" forms.

The optimal tests feature two notable features. First, they may reward the admission of ignorance, by way of the "I don't know" option. This result can be interpreted as providing a basis for the "guessing penalty" used in many multiple-choice tests such as the SAT (till 2016). Secondly, they may pass incorrect answers or fail correct answers. This "unfairness" of the optimal tests is the most salient consequence of agency issues particular to the knowledge screening problem. I show that only "obvious" answers can be failed in spite of being correct and only "counterintuitive" answers can be passed in spite of being incorrect. An obvious (respectively, counterintuitive) answer is the a priori likely (respectively, unlikely) answer, when the prior is sufficiently extreme. This reflects the common feature of real world evaluation schemes which sometimes attach greater penalty to getting "obvious" questions wrong than to getting "trick" questions wrong.

In two follow up papers (works in progress) I explore a case where (1) the test-taker also needs to "show his work" and (2) another, where the subject matter is more complex – namely, a finitely-valued state. I model the first as a mechanism design with evidence problem. Using a symmetric setting, I show that the optimal test passes (respectively, fails) the test-taker regardless of his answer, if the amount of evidence provided ("steps shown") is sufficiently high (respectively, sufficiently low), and passes if and only if his answer is correct, when the amount of evidence provided is intermediate. In the second, again using a symmetric setting, I show that when there are potentially many questions and/or the answer possibilities of each question are many, the optimal test allows the test-taker to pick a fixed number of "correct" answers, and passes him if the actual correct answer is one of them.

Representative Papers:

[1] Optimal Test Design for Knowledge-based Screening (EC '23)
[2] Communication via Hard and Soft Information (EC '23)
[3] Information Design in One-sided Matching Markets (R&R, JME)

TOM DEMEULEMEESTER (Homepage, CV, Google Scholar)

Thesis: Fairness Through Randomization: an Operations Research Perspective ('24)

Advisor: Roel Leus (KU Leuven), Dries Goossens (Ghent University)

Brief Biography: I am a final year Ph.D. student at KU Leuven. While my educational background lies in operations research, I enjoy interdisciplinary research on the intersections with theoretical economics and theoretical computer science.

Research Summary: The main focus of my research has been to study settings where fairness can only be obtained by introducing randomization. More specifically, I have studied how to obtain probabilistic solution concepts with compelling fairness properties, and how these probabilistic outcomes can be written as a lottery over deterministic outcomes that all satisfy a range of properties that are deemed desirable for the problem at hand ("best of both worlds" approach).

A first problem that I have been particularly interested in is the one-sided matching problem (school choice, house allocation...). In [1], for example, we propose an algorithmic framework to decompose a given probabilistic assignment over deterministic matchings in such a way that the risk of ending up with an "undesirable" outcome is minimized, for various interpretations of undesirability. Alternatively, in [2], we introduce a new probabilistic mechanism for the classical assignment problem which is concerned with the egalitarianism of the final assignment. The resulting *Rawlsian assignment* lexicographically maximizes the well-being of the worst-off agents. We study its axiomatic properties, and validate its performance on data from housing cooperatives in Uruguay.

Second, we study how to fairly select one of the (possibly many) optimal solutions of an integer programming formulation in [3]. While solvers traditionally return one of the optimal solutions deterministically, we argue why a randomized approach is desirable for high-impact problems, such as kidney exchange, matching variants, or sortition. We propose solution methods to implement, for example, the Nash maximum welfare, the leximin or the Random Serial Dictatorship rules without requiring a full enumeration of the optimal solutions.

In general, I aim to obtain a holistic understanding of the problems that I study by collaborating with researchers from various backgrounds. This has resulted, for example, in collaborations on coalition formation [4] and committee voting [5].

Representative Papers:

[1] A Pessimist's Approach to One-Sided Matching (European Journal of Operational Research) with D. Goossens, B. Hermans, and R. Leus

[2] Rawlsian Assignments (under review) with J. S. Pereyra

[3] Fair Integer Programming Under Dichotomous Preferences (major revision, European Journal of Operational Research) with D. Goossens, B. Hermans, and R. Leus

[4] Relaxed Core Stability in Hedonic Games with Size-Dependent Utilities (MFCS 2023) with J. Peters

[5] Strategy-Proofness and Proportionality in Party-Approval Multi-Winner Elections (AAAI 2023) with T. Delemazure, M. Eberl, J. Israel, and P. Lederer

ZHUN DENG (Homepage, CV)

Thesis: Robustness, Generalization and Fairness in Learning: Analysis and Design ('22)

Advisor: Cynthia Dwork, Harvard University

Brief Biography: Zhun Deng is a postdoctoral researcher at Columbia University, and also part of Simons Collaboration on the Theory of Algorithmic Fairness. Previously,he completed his Ph.D. in the Theory of Computation group at Harvard University, advised by Cynthia Dwork. He is also fortunate to work with David Parkes, Weijie Su, and James Zou on various projects. His papers have won multiple honors such as Spotlight and Oral Presentation at flagship machine learning conferences, including ICML, NeurIPS, ICLR, and AISTATS.

Research Summary: Modern digital systems powered by artificial intelligence (AI) are facing pivotal challenges standing in the way of reliable deployment: (1) Effectively *quantifying the quality of predictions* and *enabling the control of catastrophic outcomes*, as AI is deployed in high-impact, risk-sensitive domains like medicine and autonomous driving. (2) Reducing the potential for *negative social impact* of AI systems, as they inadvertently amplify existing social biases or blindly push users towards unhealthy consumption on social platforms.

My research aims to develop the next generation of principled methods that guide the design and deployment of AI systems in a responsible and societally beneficial way. In pursuit of this goal, my research agenda centers around introducing *rigorous statistical tools* in *uncertainty estimation* and *stochastic decision-making* to ensure reliable *real-world practice* of AI systems regarding prediction, decision-making, and societal value alignment. I am particularly keen on applying these methods to *interdisciplinary domains* such as healthcare, digital sociology, and large language models.

Uncertainty quantification for reliable decision-making. I work on advancing the theoretical understanding of model uncertainty [1] and deriving novel methods [2] to endow decision-making systems with performance and uncertainty guarantees (e.g., for important risk measures like value-at-risk (VaR)/ conditional value-at-risk (CVaR)), especially for complex black-box models' (e.g., large language models) deployment.

Building socially responsible digital systems. I also work on identifying ethical and societal concerns in deploying AI in digital systems and develop algorithmic solutions to provably mitigate such concerns (see [3]).

Representative Papers:

[1] Happymap: A Generalized Multi-Calibration Methods (ITCS'23) with C. Dwork, and L. Zhang

[2] Distribution-Free Statistical Dispersion Control for Societal Applicationss (NeurIPS'23, Spotlight) with T. Zollo, J. Snell, T. Pitassi, and R. Zemel

[3] FIFA: Making Fairness More Generalizable in Classifiers Trained on Imbalanced Data (ICLR'23) with J. Zhang, L. Zhang, T. Ye, Y. Coley, W. Su, and J. Zou

KATE DONAHUE (Homepage, CV, Google Scholar)

Thesis: AI as a Resource: Strategy, Uncertainty, and Societal Welfare ('24)

Advisor: Jon Kleinberg, Cornell University

Brief Biography: Kate Donahue is a final year computer science PhD candidate at Cornell advised by Jon Kleinberg. She works on algorithmic problems relating to the societal impact of AI such as fairness, human/AI collaboration and game-theoretic models of federated learning. Her work has been supported by an NSF fellowship and recognized by a FAccT Best Paper award. During her PhD, she has interned at Amazon, Google, and Microsoft Research.

Research Summary: Artifical intelligence (AI) is emerging as an important new step in a long lineage of innovations transforming human society - with effects that can be positive, as well as negative. My research agenda is to study the societal impact of AI, viewing viewing AI as a resource that can be used for social welfare, among multiple self-interested agents operating under uncertainty. My goal is to provide clear, generalizable insights for how to best use AI/algorithms in a wide range of situations.

In the first vein of research, I study data as a resource, through the lens of "model-sharing games". In this setting, each agent has access to data that can be used to create an ML model (e.g. each agent may be a hospital with data on patient outcomes). The agent can choose to build the model using only their data, or collaborate with other hospitals with related data, reflecting scenarios like federated learning and data cooperatives. However, there may be data shift between agents (e.g. different patient distributions), which means the same model performs differently for each agent, and collaborating may not always be optimal. In [KD AAAI '21], we propose a theoretical framework for studying this scenario, called "model-sharing games", which we then build upon to explore questions around stability, optimality, fairness, and incentive-compatibility [1, 2].

Beyond the model itself, I also consider the broader system the AI is embedded in: specifically, models as inputs to human decisions. Here, we view human time as a resource and try to maximize the performance of the combined human-AI system [3]. One natural way for two agents to interact is filtering: one agent (i.e algorithm) narrows the items into a subset, from which the second agent (i.e. human) selects the final item, a process which occurs in settings as varied as product recommendation, route selection, or categorical prediction. In ongoing work with K. Kollias and S. Gollapudi, we show how performance of an AI-human system is affected by accuracy gaps between the human and algorithm, the size of the set the algorithm presents to the human, and common human cognitive biases.

Representative Papers:

[1] Optimality and Stability in Federated Learning (Neurips 2021)
with J. Kleinberg

[2] Models of fairness in federated learning (WWW '23)
with J. Kleinberg

[3] Human-Algorithm Collaboration: Achieving Complementarity and Avoiding Unfairness (FAccT '22)
with A. Chouldechova, K. Kenthapadi

JESSIE FINOCCHIARO (Homepage, CV, Google Scholar)

Thesis: Designing Consistent and Convex Surrogates for General Prediction Tasks ('22)

Advisor: Rafael Frongillo, University of Colorado Boulder

Brief Biography: Jessie Finocchiaro (she/her) is a NSF Mathematical Sciences Postdoctoral Research Fellow and Fellow in the Center for Research on Computation and Society (CRCS) at Harvard University, hosted by Yiling Chen. Previously, she completed her PhD in Computer Science at the University of Colorado Boulder under the supervision of Rafael Frongillo and BS in Computer Science and Mathematics at Florida Southern College.

Research Summary: Algorithmic predictions supplement human and algorithmic decision-making in a variety of domains: predicted risk of disease leads to flagged scans for additional review from medical professionals, predicted hyperparameters change physics models for simulator design in engineering, and predicted market trends change investment strategies in quantitative finance. In practice, such algorithmic predictions are often made by training a model which minimizes some *loss* function measuring error, and predictions are used to guide some downstream *decision* or recommendation. Algorithmic predictions experiencing low error can often lead to poor decision-making and recommendations if these algorithms are not able to make "smarter" errors when necessary.

Subtle challenges emerge when incorporating such structure into loss functions. Loss functions often need to balance several desiderata; aligning the loss with decision task is often at odds with designing losses that are computationally tractable to minimize (e.g., convex). Historically, convex losses have been constructed in an ad-hoc manner, and often do not align with the intended decision task. Conversely, feasibility or equity concerns conceptually require algorithm designers to modify the optimized loss, but little work has characterized *how* various constraint formulations change decision-making. *My research agenda examines the bidirectional relationship between algorithm design and decision-making in machine learning and algorithmic economics.*

(1) *Decision-making to algorithm design.* Given a decision problem, design a "good" loss (part of the machine learning algorithm) that is statistically consistent for the decision problem [1, 3].
(2) *Algorithm design to decision-making.* Given a fixed algorithm, examine how values embedded into the algorithm change decision-making [2].

Representative Papers:

[1] An Embedding Framework for the Design and Analysis of Consistent Polyhedral Surrogates (NeurIPS 2019)
with R. Frongillo and B. Waggoner
[2] Using property elicitation to understand the impacts of fairness constraints (Working paper)
[3] The Structured Abstain Problem and the Lovász Hinge (COLT 2022)
with R. Frongillo and E. Nueve IV

BAILEY FLANIGAN (Homepage, CV)

Thesis: Supporting Better Direct Democracy via Social Choice, Social Science, and Algorithms ('24)

Advisor: Ariel Procaccia, Harvard School of Engineering and Applied Sciences

Brief Biography: I am a 5th-year PhD student in Computer Science at Carnegie Mellon University. My research combines theoretical and applied techniques from Econ-CS, social choice, and political science to design and support democratic processes that *directly* involve the public in policymaking, usually through *deliberation.* My research has been recognized by an honorable mention for the *INFORMS Doing Good with Good OR* prize, a *SIGCSE Best Paper Award*, and a *Siebel Scholarship.* Outside of research, at CMU I led the creation of *CS-JEDI,* a PhD course on diversity, equity, and inclusion. For this I received CMU-level and department-level service awards. My PhD was funded by a *Hertz Fellowship* and an *NSF GRFP.*

Research Summary: The main goal of my research is to design and support democratic processes that produce greater social benefit, work against patterns of marginalization, and cultivate public trust. I focus especially on processes that *directly* involve the public in policymaking, usually via *democratic deliberation* — a reasoning-based discussion between constituents about a political decision. Deliberative processes are a high-impact and rapidly growing application domain, with deliberative processes like *citizens' assemblies* and *deliberative town halls* now being used to inform political decisions around the world, even at the national level.

Despite the accelerating uptake of deliberative processes, many pressing questions about their principled implementation and impact remain unanswered. My research targets questions in this space, such as:

Q1 How should participants of deliberative processes be *selected*?
Q2 What does it mean for a deliberative process to be *representative?*
Q3 Why — and when — does deliberation *improve democratic outcomes*?

My work so far on *Q1* uses tools from algorithms, convex and integer optimization, and game theory to build randomized selection algorithms that guarantee notions of fairness, transparency, representation, and robustness to manipulation. These algorithms have been published in *Nature* and deployed in real-world deliberative processes by major groups of practitioners. My lines of work on *Q2* and *Q3* are rooted in innovations on standard voting models, and both rely on ideas from social choice and social science. I am now working with political scientists and practitioners to deploy my work on Q2 within UC Riverside's deliberation platform *Prytaneum.* Likewise, my work on Q3 is being deployed in ongoing deliberative processes, including a town hall that will guide the reform of Chile's constitution.

Representative Papers:

[1] Fair Algorithms for Selecting Citizens' Assemblies (*Nature* '21)
with P. Gölz, A. Gupta, B. Hennig, and A.D. Procaccia

[2] Toward Accounting for Stakes in Voting (*under submission*)
with S. Wang and A.D. Procaccia

[3] Distortion Under Public-Spirited Voting (*EC* '23)
with S. Wang and A.D. Procaccia

YURI RESENDE FONSECA (Homepage, CV)

Thesis: Learning from Optimal Actions: Theory and Empirical Analysis in Digital Platforms ('24)

Advisor: Omar Besbes, Columbia University, Ilan Lobel, New York University

Brief Biography: I am a Ph.D. candidate in Decision, Risk, and Operations at Columbia University, where I am fortunate to be advised by Omar Besbes and Ilan Lobel from NYU-Stern. I previously obtained an M.Sc. and a D.Sc. in Materials Science from the Military Institute of Engineering in Brazil under the supervision of Carlos Nelson Elias.

Research Summary: My research uses tools from operations research, learning theory, and causal inference to analyze observed decisions (emerging from algorithms or agents) in complex environments and learn from their behavior. I am interested in questions pertaining to structural estimation, mechanism design, and algorithm design.

The first part of my research agenda aims at leveraging structural models and machine learning methods to understand the complex interactions emerging in digital platforms and the implications for their design. In "*Signaling Competition in Two-Sided Markets*," I collaborated with a large service marketplace to elucidate how and how much workers anticipate competition when applying for short-term jobs. We develop and estimate a structural model that allows workers to incorporate expectations of competition levels and react to information released by the platform. Our structural model is anchored around an equilibrium concept to summarize workers' interactions at the platform level. I show that by revealing information about competition levels to workers, congestion is reduced, and the expected number of deals in the platform increases.

My second stream of work focuses on single-agent inverse optimization. In particular, I am interested in designing machine learning and optimization methods for learning structural models from revealed preferences. In "*Contextual Inverse Optimization: Offline and Online learning*", I aim to answer the following fundamental question: What and how fast can one learn based on past actions taken by agents? For the offline setting, we characterize an instance-dependent minimax regret and the corresponding optimal policy. The notion of regret is taken for adversarial new instances and captures the generalization power of the model learned. For the online setting, we improve state-of-the-art results, providing an improvement from $O(\sqrt{T})$ to $O(\log T)$ in the order of the regret. We are able to do so by designing novel online learning algorithms that exploit explicitly the geometry of this class of problems.

Representative Papers:

[1] Signaling Competition in Two-Sided Markets (EC'23)
with O. Besbes, I. Lobel, and F. Zheng

[2] Contextual Inverse Optimization: Offline and Online Learning (COLT'21, Operations Research - articles in advance)
with J. Smith, J. Doe, R. Roe, and J. Roe

[3] "Statistical Learning And Inverse Problems: A Stochastic Gradient Approach (NeurIPS'22) with Y. Saporito

ABHEEK GHOSH (Homepage, CV, Google Scholar)

Thesis: Contests: Equilibrium Analysis, Design, and Learning ('24)

Advisors: Edith Elkind and Paul W. Goldberg, University of Oxford

Brief Biography: I am a fourth-year Ph.D. student at Oxford supported by the Clarendon Fund Scholarship. I spent the summer of 2022 working with Milind Tambe at Google Research. I did my master's at UT-Austin and bachelor's at IIT-Guwahati and was previously advised by Umang Bhaskar (TIFR Mumbai).

Research Summary: My recent research has focused on contest theory. Contests are games where agents compete for valuable rewards by putting in costly and irreversible efforts. Classic examples include the all-pay auction and the Tullock contest. Contests sometimes arise naturally, like competition among firms for drug discovery and patents, among students for college seats, but they may also be explicitly organized, like crowdsourcing, cryptocurrencies, and sports.

The assumptions like the availability of information and full rationality of agents, which are necessary for equilibrium analysis, may not hold in practice for many applications of contests. My ongoing research with Paul Goldberg studies learning dynamics in contests. We show the convergence of best-response dynamics in Tullock contests with homogeneous agents. We also provide almost tight rate of convergence bounds using techniques from convex optimization and randomized algorithms [2] and a novel analysis of a general class of stochastic processes [1]. These papers also show that best-response dynamics may not converge for non-homogeneous agents; for this case, our ongoing work proves convergence for more general dynamics. Another line of research studies the convergence of learning dynamics when the agents only receive bandit feedback, which is the case in many real-life occurrences of Tullock contests. Related future work: I plan to study correlated and course-correlated equilibria and convergence of no-regret dynamics in contests like rank-order allocation (which includes all-pay auction) and Tullock.

My other research in contest theory includes: (i) Designing contests to improve diversity. In one work, we design contests to get higher participation from many agents rather than a very high effort from a few. In another work, our goal is to incentivize a target group to put in more effort. (ii) Equilibrium complexity. We show that computing an equilibrium in a certain class of contests is CLS-complete.

Other topics. In collaboration with a healthcare charity and researchers at Google, we solve a sequential limited resource allocation problem. We give near-optimal algorithms for restless multi-armed bandits using mean-field methods and show improved performance for two healthcare applications [3]. My other projects have focused on social choice theory, in particular, deliberation [4] and voting.

Representative Papers:

[1] Best-Response Dynamics in Tullock Contests with Convex Costs (WINE 2023)

[2] Best-Response Dynamics in Lottery Contests (EC 2023) with P. W. Goldberg

[3] Indexability is Not Enough for Whittle: Improved, Near-Optimal Algorithms for Restless Bandits (AAMAS 2023) with D. Nagaraj, M. Jain, M. Tambe

[4] Complexity of Deliberative Coalition Formation (AAAI 2022) with E. Elkind, P. W. Goldberg

DENIZALP (DENI) GOKTAS (Homepage, CV, Google Scholar)

Thesis: A Generative Adversarial Theory of Games ('24)

Advisor: Amy Greenwald, Brown University

Brief Biography: Deni is a fifth year Computer Science Ph.D. student at Brown University. His research builds and analyzes multiagent learning algorithms in games and markets with the ultimate goal of building welfare improving technology based on these algorithms. Deni has previously worked as a research scientist intern at JP Morgan's AI research lab as well as Google DeepMind's Game Theory and Multiagent Team, and was a visiting scholar at UC Berkeley's Simons institute.

Research Summary: Over the last two centuries, mathematical economists and game theorists alike, have dedicated a great deal of effort to constructing models of choice that predict the outcome of multiagent interactions. These models can be roughly categorized as general equilibrium models, i.e., market models, and game models. Although both concern preference maximizing agents, the algorithmic literature on general equilibrium and games has evolved mostly independently. I have shown that both classes of models are special cases of two-player zero-sum Stackelberg (i.e., sequential) games, and use this insight to devise efficient algorithms to solve games, and general equilibrium models.

Solving Games: we have recently introduced a class of zero-sum Stackelberg games, along with first-order algorithms to solve such games in polynomial time [1]. We have then shown that any any game is an instance of zero-sum Stackelberg games, and provided polynomial-time first-order methods to approximate equilibria in all Lipschitz-smooth games [2]. In even more recent work, we generalized our results to stochastic games [3].

Solving General Equilibrium Models: we have discovered that a large class of general equilibrium models are equivalent to zero-sum Stackelberg games [1], and have shown an interesting connection between first-order methods to solve these sequential games and age-old auction algorithms [3]. We then generalized this result to stochastic settings, showing that a large class of stochastic general equilibrium models can be seen as zero-sum stochastic Stackelberg games, providing a polynomial-time solution to these games [2].

As our world becomes increasingly more automated, our societies will face an increasing number of multiagent decision making problems. My research aims to solve these problems via the flexible and broad framework of zero-sum Stackelberg games, i.e., generative adversarial optimization. Indeed, the framework and algorithms I introduced have been used to solve problems in automated test generation for autonomous systems, resource allocation, and human-robot interaction [1,2,3].

Representative Papers:

[1] Convex-Concave Min-Max Stackelberg Games. (NeurIPS'21)
with D. Goktas, and A. Greenwald

[2] Exploitability Minimization in Games and Beyond. (NeurIPS'22)
with D. Goktas, and A. Greenwald

[3] Tâtonnement in Homothetic Fisher Markets (EC'23)
with D. Goktas, J. Zhao, and A. Greenwald

YINGKAI LI (Homepage, CV)

Thesis: Approximate Optimality of Simple Mechanisms ('22)

Advisor: Jason Hartline, Northwestern University

Brief Biography: Yingkai is a postdoc at the Cowles Foundation for Research in Economics at Yale University working with Prof. Dirk Bergemann and Prof. Yang Cai. He obtained his PhD from the Department of Computer Science at Northwestern University, where he was advised by Prof. Jason Hartline. He has also interned at Microsoft Research during the summers of 2020 and 2021.

Research Summary: My interests lie broadly at the intersection of computer science and economics, with a focus on mechanism design and information design.

Information elicitation: In many environments, firms or organizations depend on experts to obtain valuable information at a cost to facilitate more informed decision-making. We provide characterizations and compelling economic insights of the optimal mechanisms for incentivizing both the acquisition of costly information and truthful reporting by the agent. We address this in both static [1] and dynamic [2] models. Additionally, we developed polynomial time algorithms for computing these optimal mechanisms in several canonical environments.

Endogenous principal learning: In online marketplaces, platforms can accumulate vast amounts of user data. However, in many cases, these platforms lack the ability to commit to how they will use this collected data in future interactions. This strategic uncertainty can undermine the trust that users place in the platform, which is generally an undesirable outcome for the platforms. In [3], we outline conditions under which the platform consistently favors gathering more user information and conditions under which the platform significantly benefits from abstaining from collecting certain user data.

Simple mechanisms: In many cases, optimal mechanisms tend to be highly complicated and, as a result, impractical for real-world applications. However, in classic auction settings, we show that simple mechanisms can often achieve approximately optimal performance [4], and are robustly optimal when the designer faces uncertainty about the environment [5]. As part of future research, I am interested in the development of simple mechanisms and the evaluation of their performance in environments where the designer can simultaneously design the mechanisms and specify the information to be gathered.

Representative Papers:

[1] Optimization of Scoring Rules (EC 2022)
with Jason Hartline, Liren Shan, and Yifan Wu

[2] Optimal Scoring for Dynamic Information Acquisition (working paper)
with Jonathan Libgober

[3] Mechanism Design with Endogenous Principal Learning (working paper)
with Daniel Clark

[4] Simple Mechanisms for Non-linear Agents (SODA 2023)
with Yiding Feng and Jason Hartline

[5] Benchmark Design and Prior-independent Optimization (FOCS 2020)
with Jason Hartline and Aleck Johnsen

ZUN LI (Homepage, CV)

Thesis: Artificial Intelligence Methods for Economic and Computer Games ('23)

Advisor: Michael P. Wellman, University of Michigan, Ann Arbor

Brief Biography: Zun Li is a Ph.D. candidate at University of Michigan, Ann Arbor, advised by Prof. Michael P. Wellman. His research lies on the interface between artificial intelligence and computational economics. He had worked as a research scientist intern at DeepMind Alberta in 2022, working with Dr. Marc Lanctot on game-tree search in general-sum imperfect information games. He also had worked as a software engineer intern at Google display ad auction team.

Research Summary: My Ph.D. research followed precisely the chronological order that most game theory textbooks are organized: the most basic normal-form games are first studied, then are games with incomplete information, and then are dynamical games with imperfect information. The only difference here, though, is that my approaches were more from a computational perspective using practical AI methods, instead of deriving the exact mathematical solutions.

My first work [1] adopted a model-based learning approach to solve normal-form games with many players. By using supervised or unsupervised learning techniques, we can learn a succinct representation (such as clusters or a graph) of the true game using payoff data under some structural hypothesis, instead of storing an $O(NM^N)$ tensor. The computation within the learned game can be much more efficient, and the solutions were experimentally shown well in the true games. In my second work [2], I formulated the equilibrium computation in Bayesian games in a similar way as in Deep RL, where each pure strategy is represented as a neural net, and the utilities come in the form of black-box simulation data. Using natural evolution strategies, I proposed algorithms to compute pure equilibria and mixed equilibria. The methods exploit the symmetry structure of the game and scale well in high-dimensonal games. We found that deep neural nets can recover classical analytical solutions in simple games like first- and second-price auctions.

My latest work [3] extended AlphaZero-styled search method to general-sum imperfect information games by replacing MCTS with information-set MCTS, and learning a deep belief network to represent belief states at the root of the search tree. Furthermore, we combine this new search method with policy space response oracle and construct a decision-time AI bot that can conduct test-time search and online Bayesian opponent modeling. We evaluate this bot against humans in a class of negotiation games and found our bot gave comparable social welfare with humans.

Representative Papers:

[1] Structure Learning for Approximate Solutions of Many-Player Games (AAAI'20) with M. Wellman

[2] Evolution Strategies for Approximate Solution of Bayesian Games (AAAI'21) with M. Wellman

[3] Search-Improved Game-Theoretic Multiagent Reinforcement Learning in General and Negotiation Games (Extended Abstract) (AAMAS'23) with M. Lanctot, K. McKee, L. Marris, I. Gemp, D. Hennes, P. Muller, K. Larson, Y. Bachrach and M. Wellman

FRANCISCO MARMOLEJO-COSSÍO (Homepage, CV)

Thesis: Equilibrium Computation in Games and Strategic Aspects of Bitcoin Mining ('20)

Advisor: Paul Goldberg, University of Oxford

Brief Biography: I am a Lecturer and Postdoctoral Fellow at Harvard University hosted by David Parkes. I am also a Research Fellow at Input Output Global (IOG), as well as a Co-organizer of the Mechanism Design for Social Good (MD4SG) research initiative. Previously, I was a Career Development Fellow and Senior Tutor in Computer Science at Balliol College at the University of Oxford, from where I also received a D.Phil. in Computer Science under the supervision of Paul Goldberg, and an M.Sc. in Mathematics and Foundations of Computer Science. I also hold a B.A. in Mathematics (minor Neuroscience) from Harvard University. My current work was recognized as an Exemplary Applied Modeling Track Paper at EC '23, and previous work was nominated for Best Paper at WINE '20.

Research Summary: I work on societally facing challenges, including problems in public health and access to financial services, with the goal of improving access to resources within underserved communities. Beyond methodological advances, I am committed to fostering collaborations between academics and practitioners, in line with the mission of MD4SG, an interdisciplinary research initiative that brings together practitioners and researchers around the world to work towards the aforementioned goal.

My work specifically spans three key problem areas. The first is the issue of resource allocation in resource-constrained communities. I have built, deployed, and piloted algorithmic tools to help decision-makers in Mexican universities allocate scarce testing supplies during the COVID-19 pandemic [1]. In this work, I leveraged techniques from optimization and algorithm design alongside novel modeling paradigms I built through sustained conversations with students, faculty, administration, and testing personnel. Beyond its utility in immediate pandemic response, this work also has potential applications in other general resource-allocation settings. My second specialization area has to do with the disruptive potential of blockchain technology for underserved communities. Given the nascent nature of the technology, my current work has focused on understanding the broader blockchain ecosystem as a precursor to identifying use-cases and adoption barriers. This has involved modeling and studying strategic behavior in decentralized consensus protocols and decentralized finance (DeFi) [2]. In this work, I have partnered with industry partners who deploy and maintain products for a large variety of stakeholders. The third problem area is more theoretical and focuses on informational resources required to compute equilibrium concepts in game theory [3].

Representative Papers:

[1] Welfare-Maximizing Pooled Testing (EC'23, Exemplary Paper Award) with S. Finster, M. González-Amador, E. Lock, E. Micha, and A. Procaccia

[2] Strategic Liquidity Provision in Uniswap v3 (AFT '23) with Z. Fan, D. Moroz, M. Neuder, R. Rao and D. Parkes

[3] Learning Convex Partitions and Computing Game-theoretic Equilibria from Best Response Queries (ACM TEAC and WINE '18) with P. Goldberg

VISHNU V. NARAYAN (Homepage, CV)

Thesis: Multi-Item Auctions and Fair Division ('22)

Advisor: Adrian Vetta, McGill University

Brief Biography: Vishnu V. Narayan is a postdoctoral fellow at Tel Aviv University hosted by Michal Feldman. He received his Ph.D. from McGill University in 2022, where he was advised by Adrian Vetta and studied the structure of equilibria in sequential auctions. At present, his main research focus is on fair division; specifically, he is interested in exploring the power of payments in achieving fairness, and in extending fair division results beyond additive valuations and beyond the division of goods. He is also more broadly interested in other research areas within CS theory, including discrete optimization, graph theory and online algorithms, and has published work on a variety of topics. During his Ph.D., he spent a semester as a visiting fellow at Harvard (hosted by Ariel Procaccia). He has a Best Paper award from SAGT 2019 and a Teaching Assistant Award from McGill University.

Research Summary: I am drawn to problems that tie my previous expertise in combinatorics and optimization together with the Econ-CS domain. The main theme of my research is fair division, which asks how to divide a collection of items amongst agents with different preferences in a manner that everyone agrees is fair. The current decade witnessed an explosion of research in this area, but the majority of this activity is in the division of goods amongst agents with additive valuations. One focus of my research is to push past the frontier of additive valuations, and we have recently made significant progress in this area, both for goods [1,2] and chores [3]. Additionally, I am interested in exploring the effects of transfer payments in fair division. In our EC'20 paper [4], we studied the problem of achieving envy-freeness in indivisible-item instances through the use of a subsidy, and gave a tight upper bound on the amount of subsidy sufficient to always eliminate envy, resolving two conjectures in the process. My ongoing research studies the effects of payments in other settings (such as in chore division and scheduling).

One long-term goal is to complete a comprehensive analysis of the power of transfers in fair division. I also hope to make progress on the remaining big open problems in the area (such as the EFX problem); and intend to apply my expertise in fair division beyond item-allocation to other settings (such as two-sided markets, voting, budgeting, and fair algorithms in machine learning).

Representative Papers:

[1] Fair Division via Quantile Shares [Working Paper]
with Y. Babichenko, M. Feldman, and R. Holzman

[2] Breaking the Envy Cycle: Best-of-Both-Worlds Guarantees for Subadditive Valuations [Working Paper]
with M. Feldman, S. Mauras, and T. Ponitka

[3] Fair Chore Division under Binary Supermodular Costs [AAMAS'23 ext.abs.]
with S. Barman and P. Verma

[4] One Dollar Each Eliminates Envy [EC'20]
with J. Brustle, J. Dippel, M. Suzuki, and A. Vetta

ORESTIS PAPADIGENOPOULOS (Homepage, CV, Google Scholar)

Thesis: Online Decision-Making and Learning under Structured Non-Stationarity ('22)

Advisor: Constantine Caramanis, The University of Texas at Austin

Brief Biography: Orestis is a Postdoctoral Research Scientist at the Data Science Institute of Columbia University, hosted by Vineet Goyal and Assaf Zeevi. His research interests fall into the broad area of sequential decision-making under uncertainty, with applications in revenue management and operations. Before joining Columbia, Orestis completed his PhD in Computer Science at The University of Texas at Austin, under the supervision of Constantine Caramanis.

Research Summary: Through my research, my objective is to design efficient and practical algorithms that produce interpretable and (nearly) optimal solutions to sequential decision-making problems. I am particularly interested in providing the necessary framework for developing these algorithms in a systematic way, together with the analytical tools for proving theoretical guarantees on their performance.

My thesis work focuses on modeling and algorithmically leveraging the effect of human behavior in recommendation systems. In order to capture the interplay between chosen actions and altering users' preferences (e.g., satiation or deprivation), I studied non-stationary generalizations of the multi-armed bandit framework, where the reward distribution of each action is a function of the history (see, for example, my work on "recharging" bandits [1]).

In a recent direction, I turned my focus on assortment optimization, namely, the problem of deciding on an collection of goods (e.g., products, services, etc.) to offer to a customer in order to maximize the expected revenue of the seller. In joint work with Goyal, Humair, and Zeevi [2], we studied an online variation of the problem where the potential goods are observed sequentially and assortment decisions are made instantaneously and irrevocably. Assuming prior distributional knowledge on the features of each good, we developed (nearly) optimal threshold-based online policies for standard demand models and feasibility constraints.

In addition to the above, I have worked on various topics of (sequential) decision-making, including prophet inequalities under limited information [3], recurrent optimal-stopping, online learning under complex feedback structures, and resource allocation [4].

Representative Papers:

[1] Non-Stationary Bandits under Recharging Payoffs: Improved Planning with Sublinear Regret (NeurIPS '22) with C. Caramanis and S. Shakkottai

[2] MNL-Prophet: Sequential Assortment Selection under Uncertainty (WINE '23) with V. Goyal, S. Humair, and A. Zeevi

[3] Single-Sample Prophet Inequalities via Greedy-Ordered Selection (SODA '22) with C. Caramanis, P. Dütting, M. Faw, F. Fusco, P. Lazos, S. Leonardi, E. Pountourakis, and R. Reiffenhäuser

[4] A Constant-Factor Approximation for Generalized Malleable Scheduling under $M^\natural$-Concave Processing Speeds (IPCO '22), with D. Fotakis and J. Matuschke

JUSTIN PAYAN (Homepage, CV, Google Scholar)

Thesis: Data-Driven Optimization for Social Good ('24)

Advisor: Yair Zick, University of Massachusetts Amherst

Brief Biography: I am a sixth-year PhD candidate at UMass Amherst, where I work on resource allocation, data science, and peer review systems. During my PhD, I have enjoyed interning at Amazon Alexa as a research intern in 2020 and 2021, Microsoft as a data science intern in 2022, and Adobe Research as a research intern in 2023.

Research Summary: I apply techniques from combinatorial optimization, market design, and data science to high-stakes social problems. My main application is peer review, but I have applied my skills to other problems.

My work ensures that market outcomes are socially desirable, especially when the connections between resource distribution and outcomes of interest are not fully observed. I incorporate techniques from natural language processing, optimization, and machine learning to estimate and optimize over predicted outcomes. Our framework Robust Reviewer Assignment (RRA) assigns reviewers to conference papers by building an uncertainty-aware, predictive model of review quality and ensuring high quality reviews are likely [1]. We are also building algorithms for repeated constrained matching problems, minimizing regret for resource allocation problems over time. A major motivating problem setting is reviewer assignment in ACL Rolling Review.

Because I focus on problems of social importance, I am very interested in translating my algorithms into practice. Our algorithm FairSequence is deployed in the major conference management platform OpenReview. FairSequence assigns reviewers to papers fairly, efficiently, and quickly while remaining flexible to the constantly evolving constraint set required by conference organizers [2]. I am working with the Chief Administrator of IJCAI to build NLP and ML models of review quality, and we are in talks with AAMAS to create a database of reviewer "badges" that will track and reward the best reviewers.

I also apply combinatorial optimization tools to structured NLP and computer vision problems. Transformer models have revolutionized tasks that require natural perception, and can serve as an intuitive user interface. However, many domains have hard structural requirements that cannot be satisfied easily by transformers. At Adobe, we merged computer vision and combinatorial optimization for document structure recognition (ongoing), and at Microsoft we built NLP models for code generation in Excel [3].

Representative Papers:

[1] Into the Unknown: Assigning Reviewers to Papers with Unknown Affinities (SAGT 2023) with C. Cousins and Y. Zick

[2] I Will Have Order! Optimizing Orders for Fair Reviewer Assignment (IJCAI 2022) with Y. Zick

[3] InstructExcel: A Benchmark for Natural Language Instruction in Excel (Findings of EMNLP 2023) with S. Mishra, M. Singh, C. Negreanu, C. Poelitz, C. Baral, S. Roy, R. Chakravarthy, B. van Durme, and E. Nouri

FARZAD POURBABAEE (Homepage, CV)

Thesis: Essays in Venture Capital, Reputation and Learning ('21)

Advisor: Robert M. Anderson and Federico Echenique, UC Berkeley

Brief Biography: Farzad is a postdoctoral fellow at Caltech HSS, working with Federico Echenique and Omer Tamuz. He received his Ph.D. in Economics and M.A. in Statistics from UC Berkeley, advised by Robert M. Anderson.

Research Summary: My research primarily focuses on important economic topics, including mechanism design, experimentation, information economics and social learning. I leverage mathematical tools such as stochastic analysis, high-dimensional statistics, and probability theory to address these subjects.

Mechanism Design: In [1] we study the public-good provision with privacy-aware agents. We propose adding calibrated noise to the individuals' preference messages sent to the planner, thereby protecting their privacy. Using Isoperimetric inequalities on the Boolean hypercube, we analyze the implementability, revenue, social surplus and noise robustness of the optimal provision rules.

Information Economics and Social Learning: In a *misspecified* social learning setting [2], where agents sequentially make decisions by observing a private signal as well as the actions of their predecessors, we show learning outcomes improve if and only if agent' misspecification is *moderate*. In [4], I study a dynamic model of learning and random meetings between a long-lived agent with unknown ability and heterogeneous projects. By applying tools from optimal stopping theory, I find the optimal project selection policy of the agent as a function of her reputation.

Experimentation: My other line of work is related to bandits and strategic experimentation. Specifically, in [3] we investigate the strategic decision between exploring a risky project and exploiting a safe option in a network setting, where individuals observe the past experimentation outcomes of their neighbors in the graph of social connections. In [5], I study the experimentation dynamics of a decision maker in a two-armed bandit setup, where she faces *Knightian* uncertainty regarding the payoff distribution of one arm and thus entertains *Multiplier preferences*. I frame the decision making environment as a two-player differential game against nature in continuous time, and find the optimal experimentation strategy.

Representative Papers:

[1] Binary Mechanisms under Privacy-Preserving Noise (WINE '23)
with Federico Echenique

[2] The Hazards and Benefits of Condescension in Social Learning (EC '23; Revise and Resubmit at *Theoretical Economics*)
with Itai Arieli, Yakov Babichenko, Stephan Müller and Omer Tamuz

[3] The Impact of Connectivity on the Production and Diffusion of Knowledge (working paper; Informs ADA '22) with Gustavo Manso

[4] Reputation, Learning and Project Choice in Frictional Economies (Revise and Resubmit at Economic Theory)

[5] Robust Experimentation in the Continuous Time Bandit Problem (*Economic Theory*, 2020)

ROJIN REZVAN (Homepage, CV, Google Scholar)

Thesis: Simple vs. Optimal in Multi-dimensional Mechanisms ('24)

Advisor: Shuchi Chawla, University of Texas at Austin

Brief Biography: Rojin Rezvan is a fifth-year PhD student at the University of Texas at Austin, advised by Shuchi Chawla. She received her masters degree from the University of Wisconsin-Madison. She is broadly interested in algorithmic game theory and mechanism design. More specifically, she has done research in multi-dimensional mechanism design, fairness in auctions and fair allocation. She is generally interested in the intersection of mechanism design and other fields such as fairness and decentralized systems.

Research Summary: One of the main focuses of my PhD is on the paradigm of "Simple vs. Optimal" in mechanism design for multi-dimensional settings. Multi-item mechanisms can have undesirable properties such as unbounded revenue, *lottery* options in the menu and super-additive pricing function. To circumvent these issues, there are two paths to take: 1) Make some assumptions, such as independence over item value distributions and the buyers' value functions, 2) Examine the validity of the benchmark. The approach we took in [1] and [2] was the latter.

Our proposal is to compare any *simple* mechanism we design to a more realistic benchmark, called "Buy-many". In this setting, it is assumed that each buyer can interact with the menu multiple times. This ensures that super-additive pricing will not happen. The main difference now is while optimal revenue may be unbounded, the gap between revenue of optimal simple mechanisms such as item pricing and optimal buy-many mechanisms is logarithmic in the number of items. In [1], we propose a structure necessary over the item values, with which we will get fine-grained results in terms of approximation and computation of the buy-many revenue via item pricing. In [2], we extend these results and definitions to multi-buyer setting.

I am also interested in algorithmic fairness. In [4], we ask: is it possible that certain allocation algorithms in ad auctions introduce unfairness to the allocations in addition to the data? The answer is yes: an algorithm that always allocates to the highest bidder, such as FPA, could potentially turn minor differences in bids to large differences in allocation. To circumvent the issue, we propose two different algorithms that ensure fairness, while losing a fraction of the optimal social welfare, or consequently revenue. Currently, I am working to extend this work to cases where the advertisers have budgets.

Representative Papers:

[1] Pricing Ordered Items (STOC 22) with S. Chawla, Y. Teng, C. Tzamos

[2] Buy-many Mechanisms for Many Unit-demand Buyers (WINE 23) with S. Chawla, Y. Teng, C. Tzamos

[3] Prophet Secretary Against the Online Optimal (EC 23) with P. Duetting, E. Gergatsouli, Y. Teng, and A. Tsigonias-Dimitriadis

[4] Individually Fair Auctions for Mutli-Slot Sponsored Search (Best student paper at FORC 22) with C. Chawla, N. Sauerberg

ABHIN SHAH (Homepage, CV, Google Scholar)

Thesis: Data Rich Causal Inference ('24)

Advisors: Devavrat Shah, Gregory W. Wornell, and Alberto Abadie, MIT

Brief Biography: Abhin Shah is a final-year Ph.D. student in the department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology (MIT), where he works with Prof. Devavrat Shah, Prof. Greg Wornell, and Prof. Alberto Abadie. He is a recipient of MIT's Jacobs Presidential Fellowship.

Research Summary: My research develops methods for causal inference from observational data with the goal of personalized, data-driven decision-making. In particular, I address individual-related and distributional questions in causal inference, e.g., what will be the distribution of a consumer's behavior if we expose them to a product? These levels of analysis offer a more nuanced understanding compared to traditional approaches focused on population-related mean questions.

To tackle confounding in observational data, I leverage tools and frameworks from machine learning and statistics. Below, I provide a summary of my research across three threads, each exploiting a different structure to account for confounding.

(1) *Structure in distribution*: I develop computationally efficient estimators for learning exponential family distributions and use this framework to model unobserved confounding. My work estimates individual-level means of outcomes with just one sample per individual. Integrating this framework with matrix completion, we infer individual-level distributions of outcomes.

(2) *Structure in equation model*: I use matrix completion to exploit low-dimensional relationship between outcomes and unobserved factors, as well as interventions and unobserved factors, to provide doubly robust estimates of individual-level means of outcomes. These estimates remain accurate even if either outcomes-factors or interventions-factors relationship is mis-specified.

(3) *Structure in causal graph*: I design data-driven conditional independence tests to infer population-level distribution of the outcome. These tests identify subsets of the observed data that account for unobserved factors, with limited knowledge of the causal generative graph.

My long-term mission is to construct robust and equitable models. While causality is key to robustness, algorithmic fairness plays a pivotal role in ensuring equity. In pursuit of this mission, my research has eliminated disparities across protected groups in applications such as healthcare and criminal-justice by (i) introducing a new notion of fairness for models that predict selectively and (ii) ensuring fairness when protected attribute data is uncertain.

Representative Papers:

[1] On Counterfactual Inference with Unobserved Confounding (under review at Operations Research) with R. Dwivedi, D. Shah, and G. W. Wornell

[2] Doubly Robust Causal Inference with Latent Factor Models (working paper) with A. Abadie, A. Agarwal, and R. Dwivedi

[3] Front-door Adjustment Beyond Markov Equivalence with Limited Graph Knowledge (NeurIPS 2023) with K. Shanmugam and M. Kocaoglu

KANGNING WANG (Homepage, CV, Google Scholar)

Thesis: Approximations for Economic Efficiency and Fairness ('22)

Advisor: Kamesh Munagala, Duke University

Brief Biography: Kangning is a Motwani Postdoctoral Fellow at Stanford University, hosted by Moses Charikar and Aviad Rubinstein. He earned his Ph.D. in Computer Science from Duke University, advised by Kamesh Munagala. He was a J.P. Morgan Research Fellow of the program *Data-Driven Decision Processes* at the Simons Institute, UC Berkeley. He interned twice at Google Research, hosted by Jieming Mao, Renato Paes Leme, and Aranyak Mehta. His work has been recognized by an ACM SIGecom Doctoral Dissertation Award Honorable Mention, a Duke CS Best Dissertation Award, and the WINE 2018 Best Paper Award.

Research Summary: I work broadly in Economics and Computation, with a focus on mechanism design, social choice, information design, and algorithmic fairness. I am particularly interested in developing economic solutions with approximation guarantees for economic objectives such as utility, revenue, efficiency, and fairness.

Bilateral trade is a common economic scenario with a rich literature. The celebrated Myerson-Satterthwaite impossibility theorem shows that bilateral trade generally cannot be efficient under incentives. However, in [1], we reveal that simple mechanisms can always be approximately efficient, answering a prominent open question. Our other works explore the impact of budget constraints, valuation correlation, lack of priors, and impatience on pricing, auctions, and bilateral trade.

In [2], we use a classical price discrimination model to explore how an entity with information about consumer valuations can increase consumer welfare by *persuading* the seller to set personalized prices. We show that, surprisingly, there exists a solution that approximately maximizes all "reasonable" welfare functions simultaneously. Our other works extend this classical price discrimination model to auctions and bilateral trade.

Metric distortion in social choice is a well-studied framework that measures the efficiency of voting rules. In our works including [3], we design new simple voting rules that broke long-standing efficiency barriers (for deterministic rules and for possibly randomized ones), answering frequently asked open questions.

Committee selection and participatory budgeting are common democratic scenarios that require fair solutions. The *core* is often considered to provide the ultimate form of proportionality guarantees. In many of our settings including [4], even though core solutions may not exist, we show that natural relaxations always do.

Representative Papers:

[1] Approximately Efficient Bilateral Trade (STOC 2022)
with Y. Deng, J. Mao, and B. Sivan

[2] Fair Price Discrimination (SODA 2024)
with S. Banerjee, K. Munagala, and Y. Shen

[3] Breaking the Metric Voting Distortion Barrier (SODA 2024)
with M. Charikar, P. Ramakrishnan, and H. Wu

[4] Approximately Stable Committee Selection (STOC 2020)
with Z. Jiang and K. Munagala

LILY XU (Homepage, CV, Google Scholar)

Thesis: High-stakes decisions from low-quality data: AI decision-making for planetary health ('24)

Advisor: Milind Tambe, Harvard

Brief Biography: I am a CS PhD student at Harvard developing methods in machine learning, sequential planning, and causal inference for planetary health challenges, particularly biodiversity conservation and public health. My research has been recognized with AAAI best paper runner-up, the INFORMS Doing Good with Good OR award, a Google PhD Fellowship, and a Siebel Scholarship. I also co-organize the Mechanism Design for Social Good (MD4SG) research initiative.

Research Summary: I develop and deploy machine learning methods to make reliable decisions in high-stakes settings when data are incomplete. Guided by research questions that emerge from my close collaboration with the public sector, my work enables practitioners to take efficient, robust actions necessary for planetary health. I have worked closely with NGOs and government stakeholders to deploy AI for on-the-ground conservation and maternal healthcare. My research spans:

learning under uncertainty. Online learning enables us to proactively collect more data to improve our models, but may lead to unnecessary exploration. I draw inspiration from immersion in the domains with which I work to integrate problem structure into algorithm design, reducing exploration to achieve higher reward more quickly for multi-armed bandits [1] and reinforcement learning [3].

robust, efficient sequential planning. Resource allocation problems involve multistep sequential decisions and combinatorial actions, introducing exponentially large action spaces and NP-hard optimization. I develop algorithms with provably strong guarantees to make these challenging problems more tractable, integrating advances from game theory and mixed-integer programming [2].

causal inference for impact evaluation. Causal inference is hard when we cannot conduct RCTs and the available data is messy. I show that machine learning can help overcome these missing data challenges, studying ranger patrol data from a national park in Uganda with one of the highest levels of poaching in the world. Our results provide the first causal evidence for poaching deterrence, showing that ranger patrols reduce poaching by 46%.

Key to my work is my commitment to **deployment**. I partner closely with stakeholders to identify bottlenecks in existing algorithmic solutions and deploy AI to achieve measurable impact on the ground. My work on predicting poaching hotspots has been deployed in multiple countries and is being scaled to 1,200 protected areas worldwide through integration with SMART conservation software.

Representative Papers:

[1] Dual-Mandate Patrols: Multi-Armed Bandits for Green Security (AAAI 2021) Lily Xu, Elizabeth Bondi, Fei Fang, Andrew Perrault, Kai Wang, Milind Tambe

[2] Robust Reinforcement Learning Under Minimax Regret for Green Security (UAI 2021) Lily Xu, Andrew Perrault, Fei Fang, Haipeng Chen, Milind Tambe

[3] Optimistic Whittle Index Policy: Online Learning for Restless Bandits (AAAI 2023) Kai Wang*, Lily Xu*, Aparna Taneja, Milind Tambe

AVIV YAISH (Homepage, CV, Google Scholar)

Thesis: The Security and Economics of Cryptocurrencies ('24)

Advisor: Aviv Zohar, The Hebrew University

Brief Biography: Aviv is a Ph.D. candidate in the Computer Science department at the Hebrew University. His research delves into the intricate relationship between the economics and security of distributed systems. Among other honors, Aviv received the four-year merit-based Ze'ev Jabotinsky Fellowship for Ph.D. students, the Hebrew University's rector award for first-in-class computer science M.Sc. students, and the Austria-Israel Academic Network Innsbruck visiting researcher fellowship. Between the second year of his M.Sc. and until '23, Aviv served as a lecturer for two undergraduate courses, and has received an award for his teaching. Aviv was also a research associate at Matter Labs, and a visiting researcher at the University of Innsbruck, having visited twice during the spring and autumn of '23.

Research Summary: Cryptocurrencies use decentralized mechanisms to facilitate transactions amounting to a daily volume of billions of USD. Consequently, the actors involved in operating cryptocurrencies (i.e., *miners*) may have large monetary incentives to manipulate these mechanisms. Thus, the fundamental question driving my research is: *how can one design mechanisms that align the incentives of cryptocurrency actors, and ensure the secure and smooth operation of the system?*

To answer this question, I adopt a blend of theoretical and practical approaches. The latter, in particular, allow my work to produce a positive real-world impact on the design of popular cryptocurrency mechanisms.

In [1,3], my colleagues and I formalize and analyze novel strategic deviations for miners. In [1], we analyze the "clash" of incentives created when financial applications operate on top of a cryptocurrency: profits extracted from applications can incentivize miners to manipulate cryptocurrency mechanisms. In [3] we unveil two major results: alongside a novel class of miner deviations, we also present the *first* evidence of miners of a large cryptocurrency deviating strategically in the wild. Specifically, we show that Ethereum's mechanism was under attack for two years.

In [2], we show that the transaction fee mechanism used by Ethereum is flawed: it only compensates miners for processing transactions that are included in the blockchain. Thus, adversarial actors can lower transaction throughput by creating transactions that are recognized as ineligible for inclusion only at the end of their execution, without fully compensating miners for their work. We also show that heuristics used to avoid this issue can be circumvented.

In [1,2,3], we analyze modifications to the discussed mechanisms that diminish the profitability of our manipulations. These were disclosed to the relevant cryptocurrency foundations and companies, leading to changes to their mechanisms.

Representative Papers:

[1] Blockchain Stretching & Squeezing: Manipulating Time for Your Best Interest (EC'22) with S. Tochner, and A. Zohar

[2] Speculative Denial-of-Service Attacks in Ethereum with K. Qin, L. Zhou, A. Zohar, and A. Gervais

[3] Uncle Maker: (Time)Stamping Out the Competition in Ethereum (CCS'23) with G. Stern, and A. Zohar

Index