SIGecom Job Market Candidate Profiles 2020

Edited by VASILIS GKATZELIS and JASON HARTLINE

This is the fifth annual collection of profiles of the junior faculty job market candidates of the SIGecom community. The seventeen candidates for 2020 are listed alphabetically and indexed by research areas that define the interests of the community. Along with information regarding the candidate's bio, research summary, and representative papers, each profile also contains links to the candidate's homepage, CV, and Google scholar publication profile.

Employers with a relevant job opening can reach all the listed candidates directly by sending the relevant information to the moderated mailing list ecom-candidates2020@acm.org.

–Vasilis Gkatzelis and Jason Hartline

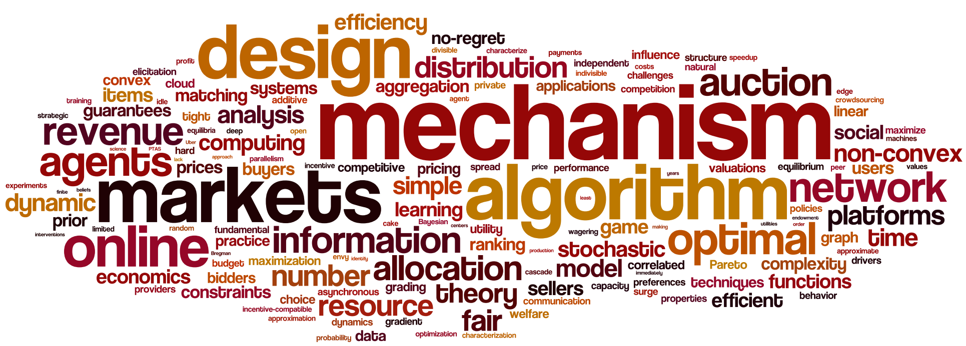

Fig. 1. Generated using the research summaries of the candidates.

Contents

NAVID AZIZAN (Homepage, CV, Scholar)

Thesis: Optimization Algorithms for Non-Convex and Networked Systems

Advisor: Adam Wierman and Babak Hassibi, California Institute of Technology

Brief Biography: Navid Azizan is a PhD candidate in Computing and Mathematical Sciences (CMS) at the California Institute of Technology (Caltech), where he is co-advised by Adam Wierman and Babak Hassibi. During the summer of 2019, he was a research scientist intern at Google DeepMind. He received the B.Sc. degree form Sharif University of Technology and the M.Sc. degree from the University of Southern California, in 2013 and 2015, respectively. His research interests broadly lie in optimization, networks, markets, and machine learning. He is the recipient of several awards, including the 2016 ACM GreenMetrics Best Student Paper Award, the Amazon PhD Fellowship, and the PIMCO PhD Fellowship.

Research Summary: While convexity is a foundational assumption in much of the literature in both computing and economics, most systems are *non-convex* in practice, and not accounting for these non-convexities can lead to significant inefficiencies. The other challenging aspect of today's large-scale systems is that they are *networked*. My research is aimed at developing new tools for solving non-convex and networked problems and applying these tools to real-world systems in various domains such as market design and machine learning.

One of the most important examples of non-convex problems in economics is that of pricing in non-convex markets (which is, for instance, of crucial importance to energy markets). When the costs are non-convex, there may be no linear prices that support a competitive equilibrium, and there is a large literature on how to determine prices in such markets. However, all the existing pricing schemes lack important economic properties. In [1], we provide a new approach to pricing in non-convex markets, which is applicable to general non-convex costs, and guarantees all the economic properties sought in the literature. Further, we provide a polynomial-time algorithm to compute the prices that works for a very general family of non-convex costs. Lastly, we show how to extend the result to networked markets.

Another important, and perhaps the most prevalent, non-convex problem of the modern era in computing is that of training deep neural networks. While these problems are non-convex in general, we prove in [3] that the huge number of parameters in typical deep models makes the training algorithm converge to a *global* minimum, and that in such settings, stochastic gradient descent acts as a particular form of regularizer. We further propose using an alternative algorithm, i.e., stochastic mirror descent, for training deep nets, which we prove induces a different form of regularization, and leads to improving the state of the art by a wide margin.

Representative Papers:

[1] Optimal Pricing in Markets with Non-Convex Costs (*Operations Research*, 2019, and EC'19) with Y. Su, K. Dvijotham, and A. Wierman

[2] Opportunities for Price Manipulation by Aggregators in Electricity Markets (*IEEE Trans. Smart Grid*, 2017, and *Performance Evaluation Review*, 2016) with K. Dvijotham, N. Chen, and A. Wierman

[3] Stochastic Gradient/Mirror Descent: Minimax Optimality and Implicit Regularization (ICLR'19, and NeurIPS'18 Workshop) with B. Hassibi

YUAN DENG (Homepage, CV, Scholar)

Thesis: Dynamic Mechanism Design in Complex Environments

Advisor: Vincent Conitzer, Duke Univeresity

Brief Biography: Yuan is a fifth-year PhD candidate at Duke University in Computer Science, where he is advised by Vincent Conitzer. Before that, he was an undergraduate student at Tsinghua University. He spent two summers interning in the market algorithms group at Google Research New York: with Sébastien Lahaie and Vahab Mirrokni in 2018, and with Balasubramanian Sivan in 2019. He is a recipient of the 2018 Google PhD Fellowship in Algorithms, Optimizations and Markets, and James B. Duke Fellowship. His research is broadly situated at the interface between economics and computer science, and mainly seeks to understand how to design mechanisms in dynamic environments.

Research Summary: As a fundamental problem in mechanism design, pricing in repeated auctions has been extensively studied in recent years, partly motivated by the popularity of selling online ads via auctions. Dynamic auctions open up the possibility of linking auctions across time to achieve higher revenue and/or welfare than static auctions, but the complexity and reliance on exact distributional information limit the deployment of dynamic auctions in practice.

My research addresses two aspects of dynamic mechanism design. The first concerns how to design dynamic auctions when only the estimated distributional information is available to the seller. We introduce a framework for the design of mechanisms that are robust to both estimation errors in the distributional information and the buyers' strategic behavior. We combine our framework and techniques from online learning to design auctions that estimate the distributional information on the fly and achieve low regret against the revenue-optimal dynamic auction [1]. The second thread explores how to design dynamic mechanisms with constraints, in particular, when the buyers have budget constraints. We characterize the simple structures of optimal dynamic mechanisms. Applying the characterization, we design a non-clairvoyant dynamic mechanism with budget constraints that achieves a constant approximation of the revenue-optimal dynamic auction [2].

Another line of my research aims to investigate how to strategize against no-regret agents. This line of research is partly motivated by online advertising markets, in which the buyers may use no-regret algorithms in bidding while the seller can exploit the buyers' no-regret strategies to increase revenue. In general two-player normal-form games with one strategic optimizer and one no-regret learner, we present a tight characterization of what various properties of the learner's no-regret algorithms imply for the optimizer's best behavior [3]. We further extend these techniques in designing revenue-optimal auctions against no-regret buyers.

Representative Papers:

[1] A Robust Non-Clairvoyant Dynamic Mechanism for Contextual Auctions (NeurIPS'19) with S. Lahaie and V. Mirrokni

[2] Non-Clairvoyant Dynamic Mechanism Design with Budget Constraints and Beyond (SSRN) with V. Mirrokni and S. Zuo

[3] Strategizing against No-regret Learners (NeurIPS'19) with J. Schneider and B. Sivan

LUDWIG DIERKS (Homepage, CV, Scholar)

Thesis: Market Design for Cloud Computing

Advisor: Sven Seuken, University of Zurich

Brief Biography: Ludwig is a Ph.D. student advised by Prof. Sven Seuken in the Computation and Economics Research Group at the Department of Informatics of the University of Zurich. In Summer 2018, Ludwig was an intern at Microsoft in the Office of the Chief Economist. In 2015, he received a M.Sc. in Mathematics in Operations Research from TU Munich. He wrote his Master's thesis on "Cooperative Games with Payoff Restrictions" under the supervision of Prof. Felix Brandt. During his M.Sc. he visited the NCKU in Tainan, Taiwan (R.O.C.), for one year as an exchange student under the supervision of Prof. Ruey-Lin Sheu.

Research Summary: In my research, I employ a variety of market design and operations research techniques to identify and analyze inefficiencies in existing markets and to develop mechanisms that mitigate or avoid them. During my PhD, my work was mostly focused on the domain of cloud computing.

Large amounts of capacity in cloud computing centers often stand idle because they are reserved for tasks that do not actually use them at all times (e.g., maintenance, or users with long-term contracts). A natural idea to increase a cloud provider's profit is to sell this idle capacity on a secondary market, in the form of preemptible capacity that can be taken back at any time. In a first paper [1], we focus on preemptible spot markets, i.e., markets where users can directly bid for capacity. We model the provider's profit optimization problem by combining queuing theory and game theory to analyze the equilibria of the resulting queuing system. The main result is an easy-to-check condition under which a provider can simultaneously achieve a profit increase and create a Pareto improvement for the users by offering a spot market (using idle resources) alongside a fixed-price market. In an as of yet unpublished second paper, we extend the analysis to secondary fixed-price markets and compare both secondary market types.

In a second research strand, I focused on the cluster admission control problem: Many modern cloud workload are characterized by scaling resource demands. A provider therefore has to continuously decide whether she can add additional workloads to a given compute cluster or if doing so would impact existing workloads' ability to scale. In [2], we formalize the problem as a constrained partially observable Markov decision process (POMDP). As no way to feasibly solve this POMDP is known, we systematically relax it to design quick heuristic admission policies. These policies estimate moments of each workload's distribution of future resource usage. Through simulations we evaluate the performance of our policies compared to current industry standards. We further evaluate by how much utilization can be improved with learned or elicited prior information and show how to incentivize users to provide this information.

Representative Papers:

[1] Cloud Pricing: The Spot Market Strikes Back (EC'19)
with S. Seuken

[2] On the cluster admission problem for cloud computing (NetEcon'19)
with I. Kash and S. Seuken

RUPERT FREEMAN (Homepage, CV, Scholar)

Thesis: Eliciting and Aggregating Information for Better Decision Making

Advisor: Vincent Conitzer, Duke University

Brief Biography: Rupert Freeman is a postdoc at Microsoft Research New York City. Previously, he received his Ph.D. from Duke University under the supervision of Vincent Conitzer. His research focuses on topics such as resource allocation, social choice, and information elicitation. He is the recipient of a Facebook Ph.D. Fellowship and a Duke Computer Science outstanding dissertation award.

Research Summary: My research focuses on problems where the goal is to elicit private information (like probabilistic judgments or subjective preferences) and aggregate it into a single output (like a forecast, action, or resource allocation). I am most excited by problems and solutions that draw on a variety of techniques from computer science and economics such as auctions and market design, game theory, algorithm design, artificial intelligence, and machine learning.

In a recent line of work, we have been exploring connections between wagering mechanisms — a class of elicitation mechanisms grounded in the classic economic theory of proper scoring rules — and other seemingly unrelated problems. For instance, a reinterpretation of outcome-contingent payments as an allocation of Arrow-Debreu securities yields a surprising equivalence between wagering and a natural resource allocation problem where agents have additive preferences over a set of divisible items [1]. The equivalence immediately yields new mechanisms and results for both settings. In other work [2], we use repeated application of a wagering mechanism to design the first incentive-compatible mechanism for selecting the winner of a forecasting competition, such as the Netflix Prize or a Kaggle competition. In an ongoing project [3], we have found that similar ideas can be used to design algorithms for learning from expert advice that are incentive compatible when experts care about the weight assigned to them by the learning algorithm, without suffering any loss in terms of regret guarantees.

In an EC 2019 paper [4], we consider a social choice problem where we are required to aggregate a set of individual budget proposals into a single proposal. We present a solution to this problem by drawing a connection between prices in a simple market and generalized median mechanisms. Although markets do not generally yield incentive-compatible mechanisms for *resource allocation*, they provide us with a powerful tool for incentive-compatible *preference aggregation* through their prices.

Representative Papers:

[1] An Equivalence Between Wagering and Fair-Division Mechanisms (AAAI 2019) with D.M. Pennock and J.W. Vaughan

[2] Incentive-Compatible Forecasting Competitions (AAAI 2018) with J. Witkowski, J.W. Vaughan, D.M. Pennock, and A. Krause

[3] No-Regret and Incentive-Compatible Prediction with Expert Advice (In preparation) with D.M. Pennock, C. Podimata, and J.W. Vaughan

[4] Truthful Aggregation of Budget Proposals (EC 2019) with D.M. Pennock, D. Peters, and J.W. Vaughan

OPHIR FRIEDLER (Homepage, CV, Scholar)

Thesis: Simple Mechanisms for Complex Environments

Advisor: Michal Feldman, Tel Aviv University

Brief Biography: Ophir is a fifth-year PhD student at the Computer Science Department in Tel Aviv University, under the supervision of Michal Feldman, where he also obtained his MSc. Prior to his MSc, Ophir obtained his BSc at the Technion - Israel Institute of Technology.

Research Summary: My main research interests are in Algorithmic Game Theory and Mechanism Design. During my PhD, I study various types of simple mechanisms, in complex settings where optimal mechanisms are poorly understood, intractable, or cannot be realistically implemented.

I put an emphasis on less standard settings such as agents with valuations that may admit complementarities across items (and not only substitutes). In [1], we study previously studied (EC 2015) *simple* mechanisms, which were analyzed for agents with complement-free valuations. In [1] we prove such mechanisms have good social welfare guarantees for valuations with limited degree of complementarity. Moreover, we provide different but similar mechanisms with improved guarantees for general valuations. For the revenue objective, in [2] we show that the the better mechanism between selling items separately, or selling the grand bundle of all items, has good revenue guarantees in settings with limited degree of complementarity.

In an EC 2017 paper, we study competition complexity – the number of extra bidders needed for a prior free mechanism, to extract at least as much revenue as the optimal, prior-based mechanism. We show that for additive bidders with feasibility constraints, the prior free VCG mechanism with a linear number of extra bidders, extracts at least as much revenue as the optimal, poorly understood, computationally intractable mechanism. In [3] we slightly relax the objective from the optimal revenue, to 99% of the optimal revenue. We provide a host of results showing that for additive bidders, 99% of the revenue can be achieved, by using substantially less (to no extra) competition.

In an ongoing work [4], we study Walrasian equilibium, subject to the endowment effect, where consumers tend to inflate the value of items they own. Following up on an EC 2018 work, we provide a more flexible formulation of the endowment effect and prove stability results with approximate efficiency for generalized settings, as well as settings where the market designer can pre-pack items to indivisible bundles.

Representative Papers:

[1] Simple Mechanisms for Agents with Complements (EC 2016)
with M. Feldman, J. Morgenstern, and G. Reiner

[2] A Simple and Approximately Optimal Mechanism for a Buyer with Complements (EC 2017)
with A. Eden, M. Feldman, I. Talgam-Cohen, and S.M. Weinberg

[3] 99% Revenue via Enhanced Competition (EC 2018)
with M. Feldman and A. Rubinstein

[4] A General Framework for Endowment Effects in Combinatorial Markets (Arxiv)
with T. Ezra and M. Feldman

NIKHIL GARG (Homepage, CV, Scholar)

Thesis: Learning and Pricing in Human-centric Platforms

Advisors: Ashish Goel & Ramesh Johari, Stanford University

Brief Biography: I am a PhD candidate at Stanford University, where I am part of the Stanford Crowdsourced Democracy Team and the Society and Algorithms Lab. I use tools from across computer science, probability, and economics to study online platforms. I received my MS in Electrical Engineering from Stanford, and a BS in Computer Engineering and a BA in Plan II (Liberal Arts) from the University of Texas at Austin. I have interned at Uber, NASA, Microsoft, the Texas Senate, and IEEE's policy arm, and am a NSF Graduate Research Fellow and McCoy Family Center for Ethics in Society Graduate Fellow.

Research Summary: Human-centric platforms increasingly mediate interactions between people, at their best enabling fair and efficient agreements within large, diverse groups. Principled design of such platforms connects behavioral considerations to more classical Econ-CS challenges such as defining and achieving appropriate objectives; understanding statistical and learning-theoretic limitations; and leveraging experiments and extant data. So far, I have worked on designing online marketplaces [1,2] and civic engagement platforms [3], focusing on mechanisms that efficiently learn heterogeneous participant preferences and price goods accordingly. In all my work I aim to bridge the gap between coarse theoretical insights and the fine-grained questions practitioners must answer, driven both by theory and data. My work has informed deployments at Uber, a large online labor platform, and in participatory budgeting elections in several U.S. cities.

In my work with Hamid Nazerzadeh [1], we consider how to design surge (dynamic) pricing in ride-hailing platforms, which is used to balance the supply of available drivers with the demand for rides. We show that due to the temporal dynamics of surge, trips of different time lengths vary in the opportunity cost they impose on drivers, and so some drivers may strategically reject trip requests to maximize their earnings, to the detriment of other drivers. We develop an incentive compatible pricing scheme to resolve this issue, with a simple enough closed-form expression to enable transparency and communication of surge prices through a heat-map. Finally, with calibrated numerics and empirical analysis, we compare additive vs multiplicative surge in practice.

Representative Papers:

[1] Driver Surge Pricing (In Submission)
with H. Nazerzadeh

[2] Designing Informative Rating Systems: Evidence from an Online Labor Market (In Submission) with R. Johari

[3] Iterative Local Voting for Collective Decision-making in Continuous Spaces (JAIR 2019 & WWW 2017)
with V. Kamble, A. Goel, D. Marn, and K. Munagala

[4] Word Embeddings Quantify 100 Years of Gender and Ethnic Stereotypes (PNAS 2018) with L. Schiebinger, D. Jurafsky, and J. Zou

ALECK JOHNSEN (Homepage, CV)

Thesis: Bid-Inversion Mechanisms and Robustness Lower Bounds

Advisor: Jason Hartline, Northwestern University

Brief Biography: Working in industry before graduate school, Aleck has accumulated 10 years of experience in trading, covering most roles: floor clerk and trader, algorithmic trader, algorithm designer, data scientist, and programmer; involving equities, fixed income, and options. He has a CFA charter and has done investor-side venture analysis of startups. During early graduate summers, Aleck wrote a C# e-file module (for transmitting to the IRS) for Forte International Tax (Evanston, IL). For 12 weeks (2016), Aleck visited booking.com (Amsterdam) with advisor Jason Hartline and colleague Dr. Denis Nekipelov (Virginia), where they helped booking's ranking team apply principles of auction theory. At Northwestern, Aleck taught EECS 214 Data Science and Data Management. His undergraduate degrees were in civil engineering; and mathematics; at University of Illinois (Urbana-Champaign).

Research Summary: Since visiting booking.com, Aleck's research has focused on auction inversion techniques, covering both "online" implementation theory and "offline" inference. The major questions are as follows. Can we design a predictive interface (e.g., a dashboard) such that we get both agents wanting to best respond to predictions, and we can infer values to run a desired social choice rule?[1] Can we design a mapping from non-truthful agent bids back to private values such that when we implement a desired choice rule, the side-effects (e.g., payments) are consistent with the original agent reports?[forthcoming] And, can we infer agents' private values in auctions from the allocation rule and the prices they pay?[2]

Aleck has further (unpublished) work on prior independent design and benchmark design for worst case analysis. For prior independence; if a correlated distribution (over n variables) can be "decomposed" in two distinct ways- i.e., into distinct distributions over i.i.d. product distributions- then a lower bound on the performance follows as a consequence. A small set of meaningful examples of such "unidentified" correlated distributions (with decompositions) have been discovered. These results can extend to online algorithm settings, not just auctions. Within revenue auctions, if the optimal prior independent mechanism can be discovered- even for $n = 2$- and happens to respect a specific structure, as a corollary it will also solve an open question of Hartline-Roughgarden regarding "optimal" prior free benchmarks.

Aleck's other interests include the interaction of game theory and big data, e.g. the bias-variance game studied in [3]; and results towards optimal expert learning with fixed time horizon.

Representative Papers:

[1] Dashboard Mechanisms for Online Marketplaces (EC 2019; arxiv)
with J. Hartline, D. Nekipelov, O. Zoeter

[2] Inference from Auction Prices (SODA 2020; arxiv)
with J. Hartline, D. Nekipelov, Z. Wang

[3] Bias-Variance Games (arxiv)
with Y. Feng, R. Gradwohl, J. Hartline, D. Nekipelov

BO LI (Homepage, CV)

Thesis: Mechanism Design with Unstructured Information

Advisor: Jing Chen, Stony Brook University

Brief Biography: Bo Li is a postdoctoral researcher in the Department of Computer Science at University of Oxford, hosted by Edith Elkind. He received his PhD in 2019 from the Department of Computer Science at Stony Brook University, under the supervision of Jing Chen. He is broadly interested in algorithms, AI and computational economics, including problems related to fair division, mechanism design, online algorithms, and their applications to Blockchain. During his PhD, he spent a summer as a visiting student at ITCS (SUFE) with Pinyan Lu, a winter as a research assistant at ICT (CAS) with Xiaoming Sun, a summer as a research assistant at CityU with Minming Li, and a summer as a research intern at Algorand with Jing Chen. He is a recipient of the Catacosinos Fellowship, the Special Department Chair Fellowship, and a Sigma Xi Award. He completed his B.S. in Applied Maths and M.S. in Operations Research at Ocean University of China.

Research Summary: Fairly dividing a number of items among agents is an important topic in multi-agent system and AI. It has many applications in practice such as sharing rents and distributing goods/tasks. Part of my research is focused on a fundamental problem: how to characterize the fairness and efficiently compute such fair divisions for various situations. For mechanism design, I study how to design resilient mechanisms that work properly even in less foreseeable environments, such as when the mechanism designer's information is imprecise or less structured, or when the computation/communication ability of the agents are limited.

With the emergence of many online platforms and dynamic data, I also study fair division and mechanism design problems in online settings, which capture many real-word scenarios, like resource sharing in data centers and labor crowdsourcing. We study to what extent the resource can be fairly divided among a number of online agents with different preferences, and how to disperse a number of online facilities to serve the demands uniformly distributed in a metric space. Furthermore, we show how to simultaneously achieve truthfulness and fairness in a 2-sided market with one side being processors and the other being online customers.

Recently, I have become interested in many game-theoretic questions raised in blockchain and bitcoin. For example, as bitcoin has been proved to be vulnerable for several attacks, we show how to refine the protocol so that participants do not launch these attacks in a Nash equilibrium. Another problem faced by blockchain is the lack of motivation for the participants to relay their information in the communication network, and in a working paper, we design a free market where strong incentives are provided to the users to fully propagate their information.

Representative Papers:

[1] Strategyproof and Approximately Maximin Fair Share Allocation of Chores (IJCAI 2019) with H. Aziz and X. Wu

[2] Efficient Approximations for the Online Dispersion Problem (SICOMP 2019) with J. Chen and Y. Li

[3] Bayesian Auctions with Efficient Queries (working paper) with J. Chen, Y. Li, and P. Lu

AMIN RAHIMIAN (Homepage, CV, Scholar)

Thesis: Faster and Further Spreads in Social Networks

Advisor: Elchanan Mossel and Dean Eckles, MIT

Brief Biography: I am a postdoctoral associate at MIT Institute for Data, Systems, and Society (IDSS), co-advised by Elchanan Mossel (MIT Math) and Dean Eckles (MIT Sloan). I did my PhD in Electrical and Systems Engineering at the University of Pennsylvania, advised by Ali Jadbabaie. Broadly speaking my works are at the intersection of networks, data and decision sciences. I borrow tools from applied probability, statistics, algorithms, as well as decision and game theory. I am mostly interested in applications involving social and economic networks.

Research Summary:

In [1], we consider the choice of k seeds in a social network to maximize the expected spread size. Most of the previous work on this problem (known as influence maximization) focuses on efficient algorithms to approximate the optimal seed sets with provable guarantees, assuming the knowledge of the entire network graph. However, in practice, obtaining full knowledge of the network structure is very costly. To address this gap, we propose algorithms that make a bounded number of queries to the graph structure and provide almost tight approximation guarantees.

In [2], we study how interventions that change the network structure can increase the speed of spread. For simple models in which contagion spreads through each edge independently at random, interventions that randomly rewire the edges would increase the speed of spread. However, for other contagion models that require multiple exposures before adoption (i.e. threshold-based contagions), recent work has argued for the opposite conclusion: highly clustered, rather than random, networks facilitate spread. In [2], we characterize the conditions under which we can reverse the latter result by allowing a small ($o(1)$) probability of sub-threshold adoptions.

In [3], we study the computations that Bayesian agents undertake when exchanging opinions over a network. The agents act repeatedly on their private information and take myopic actions that maximize their expected utility according to a fully rational posterior belief. We show that distinguishing between posteriors that are concentrated on different states of the world is NP-hard. Therefore, even approximating the Bayesian posterior beliefs is hard. We also describe a natural search algorithm to compute agents' actions and beliefs, which we call elimination of impossible signals. We show that if the network is transitive, this algorithm can be modified to run in polynomial time.

Representative Papers:

[1] Seeding with Costly Network Information (EC'19, preparing for OR)
with D. Eckles, H. Esfandiari, and E. Mossel. arXiv:1905.04325

[2] Long Ties Accelerate Noisy Threshold-Based Contagions
with D. Eckles, E. Mossel, and S. Sen. arXiv:1810.03579

[3] Bayesian Decision Making in Groups is Hard (revised and resubmitted to OR)
with J. Hązła, A. Jadbabaie, and E. Mossel. arXiv:1611.01006

ARIEL SCHVARTZMAN COHENCA (Homepage, CV)

Thesis: Circumventing Impossibility Results in Mechanism Design

Advisor: S. Matthew Weinberg, Princeton University

Brief Biography: Ariel Schvartzman Cohenca is a PhD candidate at Princeton University advised by S. Matthew Weinberg. Ariel's work focuses in understanding the trade-off between optimality and simplicity in the design of multi-dimensional auctions. He was awarded the Department of Computer Science's Graduate Student Teaching Award in 2017, and the School of Engineering and Applied Science's Award for Excellence in 2018. During the summer of 2018, Ariel was a research intern at Google-Mountain View under the supervision of Gagan Aggarwal. He obtained his B.S. in Mathematics with Computer Science from MIT in 2015.

Research Summary:

Optimal mechanism design beyond single-item settings remains a central question at the intersection of economics and computer science. The problem is intricate for a number of reasons: the mechanisms may be bizarre, computationally hard to find or simply too complex to present to a bidder. The community's focus, thus, has shifted from to asking, for instance, how complex must a mechanism be in order to extract 99% of the optimal revenue? My work joins that of others in quantifying this trade-off explicitly. Our results suggest that significantly simpler mechanisms can compete with optimal ones if the seller is willing to lose 1% of the optimal revenue (Kothari et al., FOCS 2019, Saxena et al., SODA 2018).

In settings where buyers have correlated valuations simple (or even finite) mechanisms have no hope of competing with optimal ones, even approximately. In light of this, we begin the study of beyond-worst case approximations for correlated bidders via the smoothed-analysis framework. Our results suggest ways to overcome long-standing impossibility results and shed light on the properties that make correlated distributions inapproximable (Psomas et al., EC 2019).

Finally, I am also interested in mechanism design for tournaments: how should a tournament designer pick a reasonable winner from a set of teams? We give an elegant answer to this question, showing that simple tournament formats are optimal among all fair ones that dissuade collusion (Schneider et al., ITCS 2017).

Representative Papers:

[1] Approximation Schemes for a Buyer with Independent Items via Symmetries (FOCS 2019) with P. Kothari, D. Mohan, S. Singla, and S. M. Weinberg

[2] Smoothed Analysis of Multi-Item Auctions with Correlated Values (EC 2019) with A. Psomas, and S. M. Weinberg

[3] The Menu Complexity of "One-and-a-Half-Dimensional" Mechanism Design (SODA 2018) with R. R. Saxena, and S. M. Weinberg

[4] Condorcet-Consistent and Approximately Strategyproof Tournament Rules (ITCS 2017) with J. Schneider, and S. M. Weinberg

ALI SHAMELI (Homepage, CV)

Thesis: Algorithm Design in Online and Matching Markets

Advisor: Amin Saberi, Stanford University

Brief Biography: Ali is a fifth-year Ph.D. student at Stanford university at the MS&E department where he is advised by Amin Saberi. During his Ph.D. Ali did 3 internships at Adobe Research, Google Research, and Microsoft Research. He is the recipient of Stanford Graduate Fellowship, and the recipient of the best paper award in WINE 2017. He is also a visitor at the Simons Institute during the Fall of 2019. Ali is particularly interested in the study of online and matching markets.

Research Summary: The recent explosion in online and two sided marketplaces such as Uber, Airbnb, and Amazon Mturk, emphasizes the importance of better understanding of such platforms. These platforms have various objectives which are related to maximizing revenue or welfare of the market. For example, some online social networks aim to maximize information diffusion [3] as a means to increase welfare, or some matching markets face complicated challenges in providing high quality matches for users [1]. I am interested in identifying, modeling, analyzing, and providing a better understanding of the various challenges that these platforms face toward achieving their goals. Below I provide some examples.

One of the most well studied problems in mechanism design is the school choice problem. Although variations of this model are very interesting, even making small changes to this problem renders it too hard to solve. This motivated me to study the school choice problem subject to lower and upper bound quotas. This is a very important question since it allows us to enforce diversity in schools. In our OR paper [1], we efficiently solve this problem while approximately satisfying all the distributional constraints. Our algorithm is very general and has applications in various settings such as refugee assignment, or resident matching.

Motivated by the logistical challenges faced by different firms in their production lines, we defined and tackled the production constrained bayesian selection problem [2]. We modeled this problem as an Online Bayesian Selection problem subject to laminar matroid constraints. Our main result was a PTAS which we obtained by characterizing, relaxing, and rounding the LP for the optimal online policy. Our techniques are of independent interest and have applications in other settings such as selling flight itineraries or products with inventory replenishment.

More recently, I have been studying crowdsourcing platforms at Microsoft Research. What peaked my interest about this topic was that these platforms are intrinsically very inefficient from a mechanism design perspective since both sides of the market have an incentive to circumvent the platform to avoid paying any additional fees. This raises problems of great theoretical and practical interest.

Representative Papers:

[1] Assignment Mechanisms under Distributional Constraints (Operations Research, SODA'19) with I. Ashlagi, and A. Saberi

[2] Nearly Optimal Pricing Algorithms for Production Constrained and Laminar Bayesian Selection (EC'19) with N. Anari, R. Niazadeh, and A. Saberi

[3] Information Aggregation in Overlapping Generations (WINE'17 (Best Paper)) with M. Akbarpour, and A. Saberi

BIAOSHUAI TAO (Homepage, CV, Scholar)

Thesis: Complexity, Algorithms, and Heuristics of Influence Maximization

Advisor: Grant Schoenebeck, University of Michigan

Brief Biography: I am currently working toward the Ph.D. degree with the Computer Science and Engineering Division, University of Michigan, Ann Arbor. I received the B.S. degree in mathematical science with a minor in computing from Nanyang Technological University in 2012.

Research Summary: My research focuses on the computational complexity and algorithm analysis aspects of economics problems. While I mainly work on social network analyses and resource allocation problems, I also work on other problems.

My primary research area is the *influence maximization problem*. In summary, my contribution to the *influence maximization problem* literature is two-fold. Firstly, I proved several novel and fundamental results about the algorithmic complexity of submodular influence maximization, some of which have been opened for more than 15 years. For example, I showed that influence maximization for well-studied cascade models, including the independent cascade model and the linear threshold model, is APX-hard, even if the networks are undirected [1]. Secondly, I proposed and studied a few new sociologically founded nonsubmodular cascade models and showed how they give fundamentally different recommendations to the seed-pickers in contrast to the submodular case [2, 3].

I have also been working on resource allocation problems, in particular, the *cake cutting problem*. *Fairness* is the most fundamental solution concept and is proved to be achievable in general, and my research focuses on the possibility and the algorithmic complexity of other solution concepts, such as *efficiency* and *strategy-proofness*, on top of that *fairness* is ensured. I showed that the problem of maximizing the social welfare (efficiency) while guaranteeing fairness is NP-hard to approximate to within a factor of $\Omega(\sqrt{n})$, and this problem admits a PTAS if agents' utilities are linear [5]. This result has already appeared in textbooks. I also studied the cake cutting problem in a game theory setting where agents can misreport their utility functions on the cake. I proved that i) strategy-proofness and fairness cannot be simultaneously obtained when the number of agents is finite and ii) they can be obtained if the size of the market tends to infinity [4].

Representative Papers:

[1] Influence Maximization on Undirected Graphs: Towards Closing the $(1-1/e)$ Gap (EC'19) with G. Schoenebeck

[2] Think Globally, Act Locally: On the Optimal Seeding for Nonsubmodular Influence Maximization (RANDOM'19) with G. Schoenebeck and F.Y. Yu

[3] Beyond Worst-Case (In)approximability of Nonsubmodular Influence Maximization (ACM:TOCT'19 and WINE'17) with G. Schoenebeck

[4] Cake Cutting: Envy and Truth (IJCAI'17) with X. Bei, N. Chen, G. Huzhang, and J. Wu

[5] Optimal Proportional Cake Cutting with Connected Pieces (AAAI'12) with X. Bei, N. Chen, X. Hua, and E. Yang

YIXIN TAO (Homepage, CV)

Thesis: Market Efficiency, Dynamics, and Optimization

Advisor: Richard Cole, New York University

Brief Biography: My BS degree in Computer Science was from the ACM Honor Class at Shanghai Jiao Tong University. Now, I am pursuing a Ph.D. in the Computer Science Department, Courant Institute of Mathematical Sciences, at NYU.

Research Summary: My research focuses on algorithmic game theory (AGT) and optimization. In AGT, my work mainly concerns market efficiency and market dynamics. I am also interested in fair division. In optimization, my work addresses asynchronous implementations of coordinate descent.

Market Dynamics: A major goal in AGT is to justify equilibrium concepts from a complexity perspective. One appealing approach is to identify natural distributed algorithms that converge quickly to an equilibrium. In the Fisher Market setting, we established new convergence results for two generalizations of Proportional Response when buyers have CES utility functions. The starting points are new convex and convex-concave formulations of such markets. The two generalizations correspond to suitable mirror descent algorithms applied to these formulations. Our results follow from new notions of strong Bregman convexity and of strong Bregman convex-concave functions, and associated linear rates of convergence.

Fair Division: Pareto Efficiency and envy-freeness are the foremost notions of, respectively, efficiency and fairness for the allocation problem. The question of whether there exists an allocation that is both Pareto Efficient and envy-free has been studied for a long time. Unfortunately, for general utility functions, in both the divisible and indivisible cases, solutions that are simultaneously Pareto Efficient and envy-free cannot be ensured. Our work focus on the indivisible case. We show that for any cardinal utility functions (including complementary utilities for example) and for any number of items and players, there always exists an ex-ante mixed allocation which is envy-free and Pareto Efficient.

Asynchronous Optimization: Coordinate descent is a core tool in machine learning and elsewhere. Large problem instances are common. To help solve them, two orthogonal approaches are known: acceleration and parallelism. Asynchronous parallel algorithms are appealing as they reduce the need for waiting, albeit at the cost of having to cope with out-of-date data. In our work, we give a comprehensive analysis of Asynchronous Stochastic Accelerated Coordinate Descent. We show: A linear speedup for strongly convex functions so long as the parallelism is not too large; a substantial, albeit sublinear, speedup for strongly convex functions for larger parallelism; a substantial, albeit sublinear, speedup for convex functions.

Representative Papers:

[1] An Analysis of Asynchronous Stochastic Accelerated Coordinate Descent (in preparation) with R. Cole

[2] Dynamics of Distributed Updating in Fisher Markets (EC 2018) with Y. K. Cheung, and R. Cole

[3] Large Market Games with Near Optimal Efficiency (EC 2016) with R. Cole

ALEXANDROS VOUDOURIS (Homepage, CV)

Thesis: Design and Analysis of Algorithms for Non-Cooperative Environments

Advisor: Ioannis Caragiannis, University of Patras

Brief Biography: I was awarded my Ph.D. in September 2018 by the University of Patras, where I was advised by Ioannis Caragiannis. Since then, I have been working as a postdoctoral researcher at the Department of Computer Science, University of Oxford, under the supervision of Edith Elkind. My main research interests are on the design of simple algorithms (mechanisms) with provable efficiency guarantees for many fundamental problems in Algorithmic Game Theory and Computational Social Choice. I received my MS.c. in Computer Science and Technology (December 2014), and my Diploma (5-year degree) in Computer Engineering and Informatics (July 2013) from the University of Patras.

Research Summary: My first research direction has focused on the *efficiency and complexity in strategic environments*. In this context, one of my most interesting works is about the efficiency of *allocation mechanisms* for the distribution of divisible resources [1]. This setting captures important scenarios, including the allocation of bandwidth in communication networks and cpu time in cloud computing, and as such it has been extensively studied in the related literature. My work deviates from previous papers, by making the more realistic assumption that the users have hard budget constraints, which limit the payments they can afford. This requires a new characterization of worst-case equilibria for the analysis of resource allocation mechanisms, as well as the design of new ones.

On the other hand, my second research direction has focused on the design and analysis of algorithms for *rank aggregation* with applications in *peer grading* [3] and *crowdsourcing environments* [2]. Peer grading is a method that has been adopted by the most prominent MOOC platforms to address the problem of grading the huge number of students that participate. According to a particularly simple variant of peer grading, known as *ordinal peer grading*, each student is assigned a small number of exam papers and is asked to *order* them in terms of quality from best to worst. Given the *partial* rankings provided by the students, rank aggregation methods can merge them into a global complete one, representing the relative performance of the students. Naturally, this raises many questions related to the efficiency of rank aggregation methods in terms of how well the outcome ranking agrees with the *ground truth* ranking, which a professional grader would come up with. In [3] we proposed a framework for the analysis of simple rank aggregation rules, by combining theory, simulations, as well as field experiments.

Representative Papers:

[1] The Efficiency of Resource Allocation Mechanisms for Budget-Constrained Users (EC 2018) with I. Caragiannis

[2] Optimizing Positional Scoring Rules for Rank Aggregation (AAAI 2017, AIJ 2019) with I. Caragiannis, X. Chatzigeorgiou, and G. A. Krimpas

[3] How Effective Can Ordinal Peer Grading Be? (EC 2016) with I. Caragiannis and G. A. Krimpas

DAVID WAJC (Homepage, CV, Scholar)

Thesis: Matching Algorithms Under Uncertainty, and Matching Lower Bounds

Advisor: Bernhard Haeupler, Carnegie Mellon University

Brief Biography: David is a PhD candidate at Carnegie Mellon University's computer science department. He is broadly interested in algorithms, in particular algorithms under uncertainty, including online, dynamic and distributed algorithms. His research has been published in major theory of computation venues, including FOCS, SODA, EC, PODC and ICALP. Before joining CMU, he was a Research Engineer at Yahoo! Labs, after completing an MSc and BSc (summa cum laude) at the Technion. During his studies, David has spent a semester at EPFL, has gone on numerous visits to the Simons Theory of Computing Institute at Berkeley, and has interned at Google Research and IBM R&D.

Research Summary: The proliferation of user-facing mobile and web-based apps has made online problems rise in prominence. In many such applications an online algorithm must match agents (e.g., riders and passengers) immediately and irrevocably on arrival of an agent or matching opportunity. How can such an algorithm guarantee good performance (*competitiveness*) compared to the hindsight-optimal solution? Much of my research addresses this question.

One example of such online matching-related problems is Internet ad allocation. For this problem a worst-case optimal competitive ratio of $1 - \frac{1}{e}$ was known, but experiments showed this problem is easier in practice. In an EC'15/TEAC'18 paper we gave a possible explanation of this behavior, by studying instances with an imbalanced thicknesses on the advertisers' and ad slots' sides. For such instances we showed that greedy fares well, with a competitive ratio tending to one as this imbalance grows. We then designed optimal online algorithms, whose competitive ratios tend exponentially faster to one in terms of this imbalance.

In another work (SODA'18), we studied online matching in the well-studied class of d-regular graphs, for which we presented tight $1 - \tilde{\Theta}(1/\sqrt{d})$ upper and lower bounds. Underlying our work is an online rounding scheme for bounded fractional matchings, which we later used (in FOCS'19) to optimally edge color graphs online.

Finally, a beautiful algorithm of Karp et al. proved the greedy algorithm is sub-optimal for online matching in bipartite graphs under one-sided vertex arrivals. Whether there exist better-than-greedy algorithms for edge arrivals or general vertex arrivals were vexing open questions. In recent work (also in FOCS'19) we answered these questions, negatively for the former, and positively for the latter.

Representative Papers:

[1] Online Matching with General Arrivals (FOCS'19)
with B. Gamlath, M. Kapralov, A. Maggiori, O. Svensson

[2] Tight Bounds for Online Edge Coloring (FOCS'19)
with I. R. Cohen and B. Peng

[3] Randomized Online Matching in Regular Graphs (SODA'18)
with I. R. Cohen

[4] Near-Optimum Online Ad Allocation for Targeted Advertising (EC'15, TEAC'18)
with J. Naor

FANG-YI YU (Homepage, CV, Scholar)

Thesis: Dynamics on Social Networks

Advisor: Grant Schoenebeck, University of Michigan

Brief Biography: I am currently a post-doctoral research fellow working with Grant Schoenebeck. I obtained my Ph.D. degree in Computer Science in August 2019 from the Computer Science and Engineering Division at the University of Michigan. I received the B.S. degree in Electrical Engineering with double major in Mathematics from the National Taiwan University in 2013.

Research Summary: The majority of my work involves studying the long-term behavior of dynamical systems with applications, including contagions and opinion formation on social networks, local search algorithms (stochastic gradient descent), and equilibria of no-regret learners [1, 2, 3, 4].

In a recent work [4], I study a large family of stochastic processes which contains stochastic approximation algorithm and stochastic gradient descent with a uniform step size. A key question is how this family of stochastic processes is approximated by their mean-field approximations. We provide a tight analysis: for any non-attracting fixed point in any stochastic process in this family, we show that system can escape the fixed point in $O(n \log n)$ time with high probability. We also show that it takes time $\Omega(n \log n)$ to escape such a fixed point with any constant probability. This result improves previous analysis of stochastic gradient descent escaping saddle points, and provide new insight on evolutionary stable strategies in evolutionary game theory.

In another work [3], I propose a family of binary opinion formation models, including several previous dynamics. I prove the tight bound on the consensus rate on the dense Erdos-Renyi random graphs when the dynamics are "majority-like." Technically, I propose a general framework that upper bounds the hitting time of homogeneous irreversible Markov chain, which is robust against small perturbation.

Recently, I also worked on information elicitation mechanisms which incentivize agents to report their signals truthfully even in the absence of verification. I made connections to variational methods in statistics which enables new information elicitation mechanisms in the continuous setting and a deeper understanding of information elicitation.

Representative Papers:

[1] General Threshold Model for Social Cascades: Analysis and Simulations (EC 16) with Gao, J., Ghasemiesfeh, G., and Schoenebeck, G.

[2] Sybil Detection Using Latent Network Structure (EC 16) with A. Snook, and G. Schoenebeck

[3] Consensus of Interacting Particle Systems on Erdos-Renyi Graphs (SODA 18) with Schoenebeck, G.

[4] Escaping Saddle Points: from Agent-based Models to Stochastic Gradient Descent (in submission) with Schoenebeck, G.

MANOLIS ZAMPETAKIS (Homepage, CV)

Thesis: Efficient Algorithms for Truncated and Mixture Models

Advisor: Constantinos Daskalakis, Massachusetts Institute of Technology

Brief Biography: Manolis Zampetakis is a Ph.D. student at MIT advised by Constantinos Daskalakis. His research interests include statistics, theoretical machine learning, complexity theory and mechanism design. He is a recipient of the 2018 Google PhD Fellowship on "Algorithms, Optimizations, and Market". He has interned at Google Research, NY (2017), Yahoo! Research, NY (2018) and Microsoft Research, New England (2019). He has organized workshops on complexity of total problems (FOCS'18) and on algorithms for learning and economics (WALE 2019).

Research Summary: My research is focused on how the foundations of algorithms and complexity influence other fields like Statistics, Machine Learning and Economics. I believe that these fields significantly benefit from the wide toolkit of TCS and at the same time they transform TCS by introducing new concepts and posing relevant questions that lie beyond the limit of our current understanding.

Truncated Statistics. A classical challenge in Statistics is estimation from truncated samples. Truncation occurs when samples falling outside of a subset of the support of the distribution are not observed. Truncation has myriad manifestations in all areas of the economical and physical sciences. As a simple illustration, the values that insurance adjusters observe are truncated. Indeed, clients usually only report losses that are over their deductible. In our FOCS'18, COLT'19, FOCS'19 papers, we provide the first efficient algorithms to perform statistical estimation from truncated samples in the cases of Gaussian estimation and linear regression.

Expectation - Maximization. The Expectation-Maximization (EM) algorithm, from 1977, is one of the most widely used heuristics for statistical estimation under mixture models with application to medical and economical sciences. Nevertheless, little is known about its theoretical convergence, even in the paradigmatic case of mixture of two multi-normal distributions. In our COLT'17 paper we show that in this case the EM algorithm globally and efficiently converges to the true parameters.

Complexity of Total Problems and Cryptography. An unfulfilled goal of cryptography is to build cryptographic primitives whose security is based on the hardness of a whole complexity class. The notion of NP-hardness has been proven inadequate for this purpose and hence other complexity classes need to be explored. In our FOCS'18 paper we achieve the first necessary step in this direction by identifying the first "natural" PPP-complete problem that is related to lattice-based crypto. This way we answer a longstanding open question from the seminal paper of Papadimitriou 1994 on total search problems. Our work reveals a research direction with open problems that relate complexity theory, cryptography and game theory.

Representative Papers:

[1] Efficient Statistics, in High Dimensions, from Truncated Samples (FOCS 2018) with C. Daskalakis, T. Gouleakis, and C. Tzamos

[2] Ten Steps of EM Suffice for Mixtures of Two Gaussians (COLT 2017) with C. Daskalakis and C. Tzamos

[3] PPP-completeness with Connections to Cryptography (FOCS 2018) with K. Sotiraki and G. Zirdelis

Index