SIGecom Job Market Candidate Profiles 2017

Edited by VASILIS GKATZELIS and JASON HARTLINE

This is the second annual collection of profiles of the junior faculty job market candidates of the SIGecom community. The twenty four candidates for 2017 are listed alphabetically and indexed by research areas that define the interests of the community. The candidates can be contacted individually or via the moderated mailing list ecom-candidates2017@acm.org.

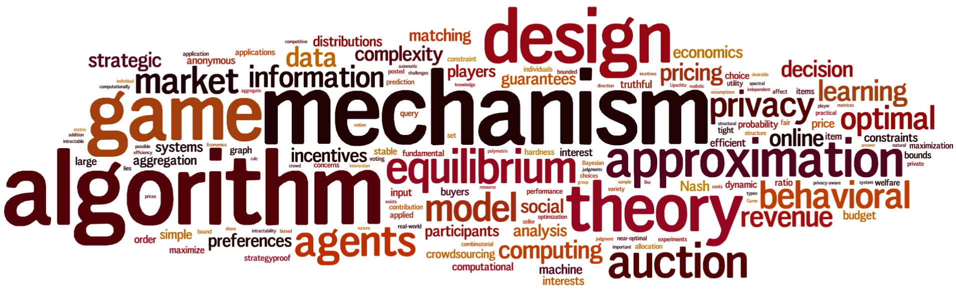

Fig. 1. Generated using the research summaries of the candidates.

Contents

SIMINA BRÂNZEI

Thesis. Computational Fair Division

Advisor. Peter Bro Miltersen, Aarhus University

Brief Biography. I'm a postdoc at the Hebrew University of Jerusalem hosted by Noam Nisan and Michael Schapira. I completed my Ph.D. in 2015 at Aarhus University, Denmark, advised by Peter Bro Miltersen, and undergraduate and Masters at the University of Waterloo, Canada, advised by Kate Larson. Before joining the Hebrew University I was a fellow at the Simons Institute for the Theory of Computing at U.C. Berkeley for the Economics and Computation program, and a (Ph.D.) visiting scholar at Carnegie Mellon University and Tsinghua University. Before grad school I worked in industry for 2 years as an intern at Google New York (AdSense project), IBM Toronto Lab (compiler optimization), and Legg Mason Canada (as a quantitative analyst).

Research Summary. I'm interested in game theory and algorithms, and my work has been investigating from a computational point of view questions inspired from economic systems.

My thesis studied the classical problem of fair division, where a heterogeneous divisible good, known as *cake* —such as land, time, mineral deposits, memory— must be allocated among participants with conflicting preferences. A major challenge in fair division is designing mechanisms that compute fair and efficient allocations for self-interested participants. One of the main results in my thesis is an impossibility: the only truthful protocols in the standard query model for cake cutting are essentially dictatorships. To complement this strong impossibility, we designed an algorithmic framework for fair division, where the cake cutting protocols are computer programs, built with simple instructions and executed jointly by the agents. We show how to construct programs that have fair outcomes in the equilibrium.

I've also been working on analyzing and designing markets for strategic agents, mechanism design, verification, unsupervised learning (How should you select a clustering algorithm for the task at hand?), social networks, voting dynamics, equilibrium computation (e.g. Stackelberg solution concept, applied in recent years to airport security), matchings, cooperative games, academic authorship, and others.

Representative Papers.

[1] A Dictatorship Theorem for Cake Cutting (IJCAI 2015)
with P. Bro Miltersen

[2] Verifiably Truthful Mechanisms (ITCS 2015)
with A. Procaccia

[3] Nash Social Welfare Approximation for Strategic Agents (preprint 2016)
with V. Gkatzelis and R. Mehta

MITHUN CHAKRABORTY

Thesis. On The Aggregation Of Subjective Inputs From Multiple Sources

Advisor. Sanmay Das, Washington University in St. Louis

Brief Biography. I am currently a PhD student in the Department of Computer Science and Engineering at Washington University in St. Louis (WUSTL), after spending three years at Rensselaer Polytechnic Institute and another year at Virginia Tech with my adviser Sanmay Das. I received a WU-CIRTL Practitioner Level recognition from the Teaching Center at WUSTL in May 2016, and the Department Chair Award for Outstanding Teaching in April 2016 for teaching Social Network Analysis (CSE316A) in Fall 2015. I was invited to give a talk on my research at the 20th Conference of the International Federation of Operational Research Societies (IFORS) in July 2014, and was also selected to attend the Summer School on Algorithmic Economics at Carnegie Mellon University in August, 2012. I earned my Bachelor's degree (with Honours) in Electronics and Telecommunication Engineering from Jadavpur University, Kolkata, India, in 2009.

Research Summary. My research interests span algorithmic economics and machine learning, an overarching theme being the aggregation of multiple subjective inputs (e.g. forecasts on uncertain events, noisy observations of hidden truths, etc.) into a single group judgment or estimate. I use tools and concepts from across computer science, economics, and operations research.

My recent work has focused on price properties in prediction markets. In a NIPS 2015 paper, we showed that when a prediction market implemented with a market scoring rule interacts with myopic risk-averse traders, the price process behaves like an opinion pool, a classical family of belief combination rules [2]. An IJCAI 2016 paper proposes a new model for studying manipulation in prediction markets when participants can influence the predicted outcome but some of them have a non-zero probability of being non-strategic, and shows that the equilibrium of this game is one of two types, depending on this probability - either collusive and uninformative or partially revealing [1]. Another major theme of my research is the design and analysis of market-making algorithms. In a AAAI 2015 paper, we proposed a practical adaptation of the logarithmic market scoring rule, which takes market orders only, to a setting with limit order books, and uncovered interesting properties of the resulting market ecosystem. Although my work has been mainly theoretical, I was also part of a team that designed and ran human-subject experiments to empirically evaluate the behavior of two different market making approaches ([3], AAAI 2013).

In addition to the above contributions pertaining to incentivized information aggregation, I have also worked on approaches for learning from differentially informed agents, abstracting away from truth-telling incentives. This includes devising an approximately Bayesian algorithm for learning a real-valued target from a sequence of censored noisy signals, and showing that it performs asymptotically almost as well as if we had uncensored signals (UAI 2011). I am currently involved in developing and analyzing communication protocols aimed at no-regret learning for an ad hoc team exploring and exploiting a multi-armed bandit.

In summary, my goal is to enrich and deepen our understanding of subjective input aggregation techniques used in various fields of study, and identify and address challenges encountered in practical applications of these techniques.

Representative Papers.

[1] Trading On A Rigged Game: Outcome Manipulation In Prediction Markets (IJCAI 2016) with S. Das

[2] Market Scoring Rules Act As Opinion Pools For Risk-Averse Agents
(NIPS 2015 spotlight presentation: top 16.6% of accepted papers) with S. Das

[3] A Bayesian Market Maker (EC 2012)
with A. Brahma, S. Das, A. Lavoie, and M. Magdon-Ismail

YU CHENG

Thesis. Algorithms and Complexity of Optimal Information Revelation

Advisor. Shang-Hua Teng, University of Southern California

Brief Biography. Yu Cheng is a Ph.D. candidate in the Computer Science Department at the University of Southern California (USC), advised by Shang-Hua Teng. Prior to that, he completed his B.S. in Computer Science at Shanghai Jiao Tong University in China. Outside of work, Yu advanced to the World Finals of the ACM-ICPC (International Collegiate Programming Contest) twice, and became USC's student coach for competitive programming subsequently.

Research Summary. My research interests are mainly in algorithmic game theory and spectral graph theory. In algorithmic game theory, my work has focused on algorithms and complexity for optimal information revelation, and query complexity of approximate Nash equilibria.

Strategic interactions often take place in an environment rife with uncertainty, and the act of exploiting informational advantage to affect the decision of other players is ubiquitous. We study the optimization problem faced by an informed principal in Bayesian games, who reveals (partial) information to the players about the state of nature to obtain a desirable equilibrium, a task often referred to as signaling or persuasion. In a paper in FOCS 2015, we present a broadly applicable algorithmic framework and solve a number of open problems in this area using our framework. In a paper in EC 2016, we push the boundary on computational aspects of Bayesian persuasion by proving new hardness results on signaling in zero-sum games and network routing games.

I have also studied the query complexity of Nash equilibria in large games. The problem of computing Nash equilibrium is known to be hard. However, there are still many open questions regarding the complexity of approximate Nash equilibria. We prove that for binary-action, n-player games and constant ϵ, the query complexity of an ϵ-approximate equilibrium is exponential in $(n/\log n)$. Previously, an exponential lower bound on the query complexity was known only for the stronger notion of "well-supported Nash equilibrium".

My research in spectral graph theory focuses on speeding up the matrix form of Newton's method using spectral sparsification. Newton's method provides a powerful framework for designing efficient and parallel algorithms. However, when applied to large-scale matrices, the intermediate matrices quickly become dense, even when the input matrix is sparse. We present an approach using spectral sparsification to keep all matrices sparse while preserving the effectiveness of Newton's method. As applications, we give the first nearly-linear time algorithms for sampling Gaussian graphical models with symmetric diagonally dominant precision matrices, and for solving the matrix roots problem for graph Laplacians.

Representative Papers.

[1] Well-Supported versus Approximate Nash Equilibria: Query Complexity of Large Games. (to be submitted) with X. Chen and B. Tang

[2] Hardness Results for Signaling in Bayesian Zero-Sum and Network Routing Games (EC 2016) with U. Bhaskar, Y. K. Ko, and C. Swamy

[3] Mixture Selection, Mechanism Design, and Signaling (FOCS 2015) with H.Y. Cheung, S. Dughmi, E. Emamjomeh-Zadeh, L. Han, and S.-H. Teng

ILAN COHEN

Thesis. Online Algorithms and Game Theory

Advisor. Yossi Azar, Tel Aviv University

Brief Biography. Ilan Cohen is a PhD student at the Blavatnik School of computer science in Tel Aviv University under the supervision of Professor Yossi Azar. He holds an M.S. in computer science from Tel Aviv University and a B.S. cum laude in computer science from the Technion Institute in Haifa. His research involves online and approximation algorithms with game theoretical aspects. During the past three years, he has been a teaching assistant in the Algorithms course. Prior to his doctoral program, he worked as an algorithms developer and a programmer at IDF in the intelligence corps.

Research Summary. My research interests lie at the intersection of approximation algorithms, online algorithms and game theory. My work is divided into three parts. The first part adds game theoretical aspects to fundamental online problems. The second part involves oblivious algorithms that are motivated by designing prompt mechanisms for online bounded capacity auctions. The third part covers various subjects in online packing and covering problems.

Online algorithms deal with making irrevocable decisions while handling a sequence of events. In our scenario, the events are strategic in nature and have a private cost function, and seek to maximize their utility, i.e. minimize their private cost incurred by making a decision plus the surcharge posted on the decision by our dynamic pricing scheme. An example of this is the "parking problem" where an online sequence of cars arrive in some metric space and need to park in a vacant parking spot. Online algorithms know the next car's destination and order it where to park, while in our setting the algorithm sets a surcharge for each parking place (without knowing the next car's destination) and defers the decision on where to park to the car itself. This scenario is natural for problems such as: k-server, online metric matching and metrical task systems. We achieve essentially the same approximation ratio (up to a constant) as the best known online algorithms for these problems.

A bounded capacity auction is a single-item periodic auction for bidders that arrive online, where the amount of participating bidders is bounded. The algorithm decides which agents will participate and the allocation and pricing rule. We show a reduction from a simple stochastic balls and bins game to this problem. Although the algorithm for the game is oblivious (i.e., it does not receive input), we devise a non-uniform randomized algorithm. We establish a lower bound of 1.5 and an upper bound of 1.55, which implies a 1.55 competitive ratio mechanism for this auction.

In online packing and covering problems we establish almost tight lower and upper bounds for packing multidimensional vectors into bins. In this work we give almost tight bounds on the number of bins where the competitive ratio depends on the number of dimensions and the ratio between the maximum coordinate to the bin size. Additionally, we have worked on online covering with convex objective functions, including application such as unrelated machine scheduling with startup costs.

Representative Papers.

[1] The Loss of Serving in the Dark (STOC 2013, ICALP 2015)
with Y. Azar and I. Gamzu

[2] Tight Bounds for Online Vector Bin Packing (STOC 2013)
with Y. Azar, S. Kamara, B. Shepherd

[3] Pricing Online Decisions: Beyond Auctions (SODA 2015)
with A. Eden, A. Fiat and L. Jez

RICCARDO COLINI-BALDESCHI

Thesis. Approximation Algorithms in Mechanism Design: an application to sponsored search auctions

Advisor. Stefano Leonardi, Sapienza University of Rome

Brief Biography. I am a postdoctoral researcher at LUISS working with Prof. Marco Scarsini on "Games Under Uncertainty". I obtained my PhD in Computer Science from Sapienza, University of Rome in February 2016, under the supervision of Prof. Stefano Leonardi. My dissertation "Approximation Algorithms in Mechanism Design" consisted of two main topics. The first topic studied social efficiency and revenue properties in auctions where buyers have budget constraints and matching preferences. The second studied the design of simple mechanisms that approximate the social efficiency in a double auction setting. During my PhD, I was very happy to cooperate with Prof. Amos Fiat visiting the University of Tel-Aviv in two distinct periods. At the beginning of my PhD, I had the opportunity to be a visiting student at CWI, Amsterdam with Prof. Guido Schaefer. Previously, I received M.Sc. degree (summa cum laude) in computer science at Sapienza and B.Sc. in computer engineering from the same institute.

Research Summary. My research interests fall in the area of algorithmic game theory. I am interested in the design of approximation algorithms that can be applied in the context of sponsored search auctions. Sponsored search auctions are used every day by billions of users to promote their product and are the main source of income for many large companies. Even though this topic has attracted a lot of attention from different research communities, many problems are still open and a gap between practical implementations and theoretical knowledge exists. My research can be split in two macro areas: prior-free auctions with budgets constraints and double auctions. With regards to prior-free auctions with budgets the first goal is design truthful auctions that are Pareto-optimal (PO) in a setting that is as close as possible to the real sponsored search setting. Thus, I consider buyers that have budgets, matching preferences, and a seller that owns different items with multiple copies and CTRs. In this setting we were able to provide several IC, IR, and PO mechanisms and to derive a complete characterization for PO mechanisms. The second goal concerns revenue maximization in auctions with budgets. Due to previous results (of hardness and uniqueness), I moved to the design of non-truthful algorithms that approximate the optimal envy-free revenue when buyers have budget constraints and matching preferences. We initially proved that the optimal revenue cannot be approximated within $\Omega(\min\{n, m\}^{1/2-\epsilon})$ for n buyers and m items. We then provided several constant-approximation algorithms in different subproblems that circumvent the above hardness result. I also focused on double auctions in a setting without budget constraints and with distributional knowledge over the values of the agents. Due to the impossibility result by Myerson and Satterthwaite, it is known that is not possible to maximize social efficiency with an IR, BIC, and WBB mechanism. Thus the main question is: are there simple, truthful, and strong budget balanced mechanisms that can provide a good approximation to the optimal social welfare? In a recent paper we positively answered this question providing the first mechanisms that achieve all three major design requirements for double auctions (IR, DSIC and SBB) and approximate the optimal social welfare within a constant factor. Moreover, we extended the very natural class of sequential posted price mechanisms to the two-sided market case.

Representative Papers.

[1] Approximately Efficient Double Auctions with Strong Budget Balance (SODA 2016) with B. de Keijzer, S. Leonardi, and S. Turchetta

[2] Revenue Maximizing Envy-Free Fixed-Price Auctions with Budgets (WINE 2014) with S. Leonardi, P. Sankowski, and Q. Zhang

[3] On Multiple Keyword Sponsored Search Auctions with Budgets (ICALP 2012 and ACM TEAC) with S. Leonardi, M. Henzinger, and M. Starnberger

RACHEL CUMMINGS

Thesis. The Implications of Privacy-Aware Choice

Advisor. Katrina Ligett, Caltech and Hebrew University

Brief Biography. Rachel Cummings is a Ph.D. candidate in Computing and Mathematical Sciences at the California Institute of Technology. Her research interests lie in the intersection of computer science and economics, specifically problems surrounding algorithmic game theory, data privacy, and learning theory. Her work focuses on understanding how privacy concerns affect the behavior of players in games, and incorporating privacy-aware behavior into mechanism design. She received her B.A. in Mathematics and Economics from the University of Southern California and her M.S. in Computer Science from Northwestern University. She won the Best Paper Award at the 2014 International Symposium on Distributed Computing, and she is the recipient of a Simons Award for Graduate Students in Theoretical Computer Science.

Research Summary. Each day, people make billions of observable choices: which websites to visit, which groceries to buy, which newspapers to read. In our digital age, these choices can be tracked and recorded en mass, and are frequently used for personalized pricing, targeted advertising, and broader policy decisions. When people know that their choices today affect their utility in the future, they can strategically modify their behavior to ensure that their data reveal less–or perhaps, more favorable–information about themselves. Existing models of decision-making are ill-equipped to describe such privacy-aware choices. My thesis addresses this problem by answering two primary questions: How do privacy concerns affect behavior? And how should we design markets or games differently to account for privacy-aware players?

In [3], we initiate the study of the testable implications of choice data in settings where agents are privacy-aware. We adapt the standard conceptualization of consumer choice theory to a situation where the consumer is aware of, and has preferences over, the information revealed by her choices. We show that little can be inferred about consumers' preferences once we introduce the possibility that she has concerns about privacy. This holds even when consumers' privacy preferences are assumed to satisfy relatively strong structural assumptions.

A designer might be temped to counteract these effects through privacy policy, which dictates how behavioral information is shared. In [1], we study how privacy technologies affect behavior in a simple economic model of targeted advertising. An advertiser would like to use a consumer's past purchases to target the ad he shows, but is given only an imperfect signal about the purchases. We study equilibrium behavior as a function of the privacy level promised to the consumer and show that this behavior can be highly counter-intuitive. The effect of adding noise/uncertainty in equilibrium can be the opposite of what we would expect if we ignored equilibrium incentives.

Another line of my work addresses the challenge of purchasing and aggregating data from strategic individuals with complex incentives and privacy concerns. An analyst wishes to perform a learning task, but first must incentivize each individual to share her data. In a variety of settings, we design mechanisms that incentivize players to truthfully share data, allow the analyst to accurately perform his learning task, and minimize the analyst's costs.

Representative Papers.

[1] The Strange Case of Privacy in Equilibrium Models (EC 2016)
with K. Ligett, M. Pai, and A. Roth

[2] Truthful Linear Regression (COLT 2015)
with S. Ioannidis and K. Ligett

[3] The Empirical Implications of Privacy-Aware Choice (EC 2014 & Operations Research 2016) with F. Echenique and A. Wierman

ARGYRIOS DELIGKAS

Thesis. Algorithms for Computing Approximate Equilibria in Bimatrix, Polymatrix and Lipschitz Games

Advisor. Rahul Savani, University of Liverpool

Brief Biography. Argyrios Deligkas is a fourth-year PhD student at the University of Liverpool, under the supervision of Rahul Savani. Argyrios' defence is in September 2016, and from December 2016 he will start a one-year postdoc at the Technion. His research interests include algorithms and complexity, game theory, in particular equilibrium computation and auction and mechanism design. His PhD was funded by a Microsoft Research Fellowship 2012-2015. During his PhD, he served as a lecturer at the University of Liverpool in the Fall 2015 semester. He has an undergraduate degree in Mathematics and a Masters degree in Mathematics and Computer Science from the University of Patras, Greece.

Research Summary. My primary research interest is Algorithmic Game Theory. More specifically, my work is focused on three fundamental areas in the field: equilibrium computation, combinatorial auctions, and mechanism design. I have studied various problems related to these fields, including approximation algorithms and mechanisms, randomized algorithms and mechanisms, query and communication complexity and hardness of approximation.

Equilibrium is the central solution concept that is studied in Game Theory. However, computing an exact equilibrium is PPAD-hard, so there are unlikely that there exists a polynomial-time algorithm for this problem. The hardness of computing exact equilibria has lead to the study of approximate equilibria. My research has focused on designing efficient algorithms and inapproximability results for the computation of equilibria on classical models, such as bimatrix, polymatrix and concave games, and on more recently defined models, such as biased and Lipschitz games.

Bimatrix games are the class of games that have received a lot attention. In [2] we developed a technique for computing approximate well-supported Nash equilibria. This new technique allowed us to improve upon the best known approximation guarantee that can be achieved in multiple settings, such as computational complexity, communication complexity, and query complexity. Polymatrix games are a natural succinct representation of many-player games, where players are vertices in a graph and edges in the graph correspond to bimatrix games. In STOC 2015 it was shown that is PPAD-hard to compute an ϵ-Nash equilibrium in a polymatrix game for some very small but constant epsilon. In [1] we developed a polynomial-time algorithm that finds an (almost) 0.5-Nash equilibrium of a polymatrix game. Our approximation guarantee does not depend on the number of players, which is a property that was not previously known to be achievable for polymatrix games, and still cannot be achieved for general many-player strategic-form games.

In [3] we study games with non-linear utility functions. We observe that Lipschitz continuity of the players' utility functions allows us to provide algorithms that find approximate equilibria. We prove that the best known technique ("Sampling") for strategic-form games can be applied efficiently on Lipschitz games. We provide a quasi-polynomial time approximation scheme that computes an ϵ-equilibrium for Lipschitz games, for every $\epsilon > 0$.

Representative Papers.

[1] Computing Approximate Nash Equilibria in Polymatrix Games (Algorithmica: in press) with J. Fearnley, R. Savani, and P. Spirakis

[2] Distributed Methods for Computing Approximate Equilibria (submitted) with A. Czumaj, M. Fasoulakis, J. Fearnley, M. Jurdziński, and R. Savani

[3] Lipschitz Continuity and Approximate Equilibria (submitted) with J. Fearnley and P. Spirakis

ALICE GAO

Thesis. Eliciting and Aggregating Truthful and Noisy Information

Advisor. Yiling Chen, Harvard University

Brief Biography. Alice Gao is a postdoctoral research fellow in Computer Science at University of British Columbia, where she is primarily funded by the Canadian NSERC Postdoctoral Fellowship. Her research tackles problems at the intersection of artificial intelligence, game theory, and crowdsourcing, using both theoretical and experimental methods. She is currently working on designing effective peer grading systems for large classes. Alice obtained her PhD in Computer Science from Harvard University in 2014, advised by Yiling Chen. Her PhD dissertation was recognized by the 2014 Victor Lesser Distinguished Dissertation Runner-up Award and an Honorable Mention of the 2015 SIGecom Doctoral Dissertation Award. Her PhD research was supported by a Canadian NSERC Postgraduate Scholarship for Doctoral Students and a 2014 Siebel Scholar Scholarship. She earned her Bachelor's degree in Computer Science and Mathematics from University of British Columbia.

Research Summary. Accurate information is vital for solving many problems, such as grading open-ended assignments in large classes, making product recommendations based on contributed reviews, and making policy decisions based on event predictions. To solve these problems, many systems are designed to collect judgments and beliefs of multiple dispersed individuals about some events of interest. However, such systems operate under the assumption of full cooperation by its participants, assuming that participants will voluntarily make useful and accurate contributions. In reality, participants have many incentives to behave strategically, by not investing time and effort to formulate and submit contributions, or by intentionally reporting biased or false information. Because of such strategic behavior, these systems have low quality data, which does not always lead to sound solutions. The goal of my research is to design systems and mechanisms to motivate participants to contribute high quality information about their judgments and beliefs in order to produce accurate estimates and to recommend good decisions.

My research seeks to address several challenges with building an effective system for collecting and aggregating dispersed information. First, participants should be motivated to contribute sufficient information to the system. Second, if participants do not yet have relevant information, they should be incentivized to expend costly effort to acquire such information. Third, the system must effectively combine many pieces of noisy information to produce the best single estimate.

I address these challenges using techniques from artificial intelligence and game theory, with both theoretical and experimental methods. I use game theoretic analysis to characterize participants' strategic behavior and how such behavior affects the quality of collected information. I also conduct controlled human subject experiments to verify theoretical predictions and to study participants' behavior in practice. In addition, I use machine learning to develop accurate models of participants' behavior based on experimental data.

Currently, I am designing peer grading mechanisms, which achieves the strongest incentive properties while requiring the least amount of trusted evaluations. I am also studying the effect of incentives and social interaction on the wisdom of crowds effect. I have worked on analyzing the effect of strategic behavior on information aggregation in prediction markets.

Representative Papers.

[1] Incentivizing Evaluation via Limited Access to Ground Truth: Peer Prediction Makes Things Worse (in submission) with J.R. Wright and K. Leyton-Brown

[2] Trick or Treat: Putting Peer Prediction to the Test (EC 2014) with A. Mao, Y. Chen, and R.P. Adams

[3] Market Manipulation with Outside Incentives (JAAMAS 2015) with Y. Chen, R. Goldstein, and I.A. Kash

CHIEN-JU HO

Thesis. Design and Analysis of Crowdsourcing Mechanisms

Advisor. Jennifer Wortman Vaughan, Microsoft Research

Brief Biography. Chien-Ju is a postdoctoral associate at Cornell University, working with Arpita Ghosh. He obtained his PhD in Computer Science from UCLA in 2015, advised by Jenn Wortman Vaughan. He also visited Harvard EconCS group from 2012 to 2015, hosted by Yiling Chen. His research interests are in crowdsourcing, algorithmic economics, online learning, optimization, and behavioral experiments. His dissertation was on the design and analysis of crowdsourcing mechanisms. He is the recipient of the Google Outstanding Graduate Research Award at UCLA in 2015. His work was nominated for Best Paper Award at WWW 2015.

Research Summary. The goal of my research is to design mechanisms which combine computer power and human effort together to solve problems that neither humans or computers can solve along. My research draws on ideas from game theory, machine learning, optimization, and behavioral experiments.

My dissertation research explored the incentives and learning problems in crowdsourcing mechanisms, with the goal of improving the quality of crowdwork. For example, I studied how to assign heterogeneous tasks to online arriving workers with unknown, heterogeneous skill sets. I proposed algorithms that are provably near-optimal in minimizing the assignment cost and ensuring high accuracy of aggregating workers' responses. I also studied how to design incentives to encourage high-quality contribution. In particular, I explored the problem of selecting optimal crowdsourcing contracts, i.e., performance-based payments, in which workers' payments depend on the quality of their work. I formulated this problem as a repeated version of the principal-agent problem from economics and solved it using a novel multi-armed bandit approach.

In another line of my research, I explored the problem of online purchasing data held by strategic agents for machine learning tasks. The challenge is to use past data to actively price future data in order to obtain learning guarantees, even when agents' costs can depend arbitrarily on the data itself. I showed how to convert a large class of no-regret algorithms into online posted-price and learning mechanisms. The mechanisms gave robust risk (prediction error) bounds with a budget constraint.

In addition to theoretical analysis, I conducted empirical behavioral experiments to understand how workers react to different payment schemes in crowdsourcing markets. I conducted a comprehensive set of experiments on Amazon Mechanical Turk and analyzed results from more than 2,000 workers. I then developed a realistic worker behavior model which resolves some conflicting results in the literature. Moreover, I showed our previous theoretical results on contract design continue to hold under this model. My vision in this line of work is to develop theoretical frameworks based on realistic user models with the support of empirical evidence.

Representative Papers.

[1] Low-Cost Learning via Active Data Procurement (EC 2015)
with J. Abernethy, Y. Chen, and B. Waggoner

[2] Incentivizing High Quality Crowdwork (WWW 2015)
with A. Slivkins, S. Suri, and J. Wortman Vaughan

[3] Adaptive Contract Design for Crowdsourcing Markets: Bandit Algorithms for Repeated Principal-Agent Problems (EC 2014 and JAIR 2016)
with A. Slivkins and J. Wortman Vaughan

OMER LEV

Thesis. Agent Modeling of Human Interaction: Stability, Dynamics and Cooperation

Advisor. Jeffrey S. Rosenschein, Hebrew University of Jerusalem

Brief Biography. Omer Lev is currently a post-doctoral fellow at the University of Toronto, hosted by Prof. Allan Borodin. He completed his PhD at the Hebrew University under the supervision of Prof. Jeff Rosenschein, working in the field of multi-agent systems, with an emphasis on computational social choice, auction theory, and networks. During his doctoral studies he was awarded a Microsoft Research PhD scholarship; he also worked as a part-time researcher for two years at Microsoft Research in Israel. Additionally, he received an MSc in mathematics (supervised by Prof. Sergiu Hart), an MBA (majoring in finance and business strategy), and a BSc in mathematics and computer science (in the Amirim honors program), all at the Hebrew University.

Research Summary. My research examines various real-world phenomena using a game-theoretic toolkit, while keeping the analysis closely related to real-world data and observable behavior. For example, in the area of voting and decision making, instead of limiting the analysis to a restricted group of manipulators, I have worked on models that assume all voters may be strategic. My research explores appropriate solution concepts: they should be both stable (e.g., Nash equilibrium), and reasonable, i.e., they should not result in states that are obviously implausible in real-world settings. In parallel with my theoretical work, I have also examined, using various simulations, whether proposed models indeed give rise to realistic results.

In addition to my research on voting and decision making, I have also worked on various other aspects of algorithmic game theory: pricing items in settings that are appropriate for web-based sellers (e.g., Amazon); information diffusion in social networks; recommendation systems using a social network to derive the recommendations; axiomatic approaches to network-based mechanisms; cooperative game theory, and more.

More recently, I have focused on studying crowd-based mechanisms and activities. These include both crowdfunding and crowdsourcing (umbrella terms encompassing a variety of activities). On the topic of crowdfunding, I am working to improve our fundamental understanding of current markets using real-world data, to provide us with a better starting point for analysis. On the (better-understood) topic of crowdsourcing, I have explored the theoretic models underlying some of these scenarios, such as all-pay auctions; I have also worked on mechanism design for peer evaluation (as is used in MOOC grading or paper reviewing) so as to find better (e.g., strategyproof) mechanisms. The use of various crowd-based mechanisms is increasing, and there is still much to do regarding these topics. We need to improve our understanding of how these mechanisms work today: both to characterize people's actions when using them, as well as to build models that are better able to capture this behavior. Beyond this, we need to build better mechanisms to improve the way crowdsourced activities work, and to achieve various desired properties that will enhance the outcome for participants.

Representative Papers.

[1] Strategyproof Peer Selection: Mechanisms, Analyses, and Experiments (AAAI 2016) with H. Aziz, N. Mattei, J.S. Rosenschein, and T. Walsh

[2] How Robust is the Wisdom of the Crowds? (IJCAI 2015) with N. Alon, M. Feldman, and M. Tennenholtz

[3] A Local-Dominance Theory of Voting Equilibria (EC 2014) with R. Meir and J.S. Rosenschein

JAMIE MORGENSTERN

Thesis. Market Algorithms: Incentives, Learning, and Privacy

Advisor. Avrim Blum, Carnegie Mellon University

Brief Biography. Jamie Morgenstern is a Warren Center Fellow at the University of Pennsylvania, hosted by Michael Kearns, Aaron Roth and Ricky Vohra. She completed her PhD in Computer Science at Carnegie Mellon University with Avrim Blum in 2015. Her research lies at the intersection of machine learning, economics, and privacy, drawing on tools from all of these to design robust algorithms which perform well in the presence of uncertainty and strategic agents. Recently, she has begun designing and characterizing machine learning algorithms which are *fair*, e.g. that may not discriminate against smaller populations. She was a co-organizer of the Tutorial on Algorithmic Game Theory and Data Science at EC 2016. She was awarded the NSF Graduate Research Fellowship (GRFP), a Microsoft Research Women's Scholarship, and a Simons Award for Graduate Students in Theoretical Computer Science. She spent a summer during graduate school as an intern at Microsoft Research Redmond, and another at Alcatel Lucent Technologies. Prior to starting her PhD, she completed her undergraduate work at the University of Chicago, with a BA in math and a BS with Honors in Computer Science.

Research Summary. My research uses tools from machine learning and privacy to design robust mechanisms, as well as using tools from mechanism design and privacy to design learning algorithms which behave well in the presence of data controlled by strategic or privacy-aware individuals.

Machine Learning for Auction Design. My work focuses on designing auctions which can be optimized from limited data about agents participating in the auction. What sorts of auctions should one consider when choosing based upon limited amounts of previous data? Exact optimality is often at odds with learnability from data and simplicity of the auction format, both of which are desirable properties for an auction in practice. In “The Pseudo-Dimension of Nearly Optimal Auctions,” we design revenue-maximizing auctions for selling a single item, and show these auctions can be learned from a polynomial-sized sample of buyers. We extend this work in “Learning Simple Auctions”, where we show approximately optimal auctions for multiple items can also be learned from a small number of samples. This work and “Do Prices Coordinate Markets?” suggest that pricings are sufficiently simple that a small number of samples used to pick prices will suffice to guarantee near-optimal revenue, welfare, and demand of any particular good using those prices.

Privacy as a tool for truthfulness. In “Approximately Stable, School Optimal, and Student-Truthful Many-to-One Matchings (via Differential Privacy)”, we use privacy as a tool to compute approximately stable, school-optimal matchings which are approximately truthful for the student side of the market. The study of stable matching algorithms has been one of intense interest over the last 30 years; algorithms are known which yields stable matchings, which optimizes for one side of the market, and is truthful for agents on that side. However, for the other side of the market, players may have incentive to misreport their preferences. More generally, Roth [Roth82] showed that no algorithm can yield stable matchings and be truthful for both sides of the market. Numerous papers have made distributional assumptions about players' preferences, and show that with high probability, few people will have any incentive to misreport to this algorithm. In this work, we designed a private version of the standard stable matching algorithm which, for *w*orst-case preferences from both hospitals and residents, is approximately resident-truthful, stable, and hospital-optimal.

Representative Papers.

[1] Do Prices Coordinate Markets? (STOC 2016)
with J. Hsu, R. Rogers, A. Roth, and R. Vohra

[2] The Pseudo-Dimension of Nearly Optimal Auctions (NIPS 2015)
with T. Roughgarden

[3] Approximately Stable, School Optimal, and Student-Truthful Many-to-One Matchings (via Differential Privacy) (SODA 2015)
with S. Kannan, A. Roth, and S. Wu

SWAPRAVA NATH

Thesis. Mechanism Design for Strategic Crowdsourcing

Advisor. Yadati Narahari, Indian Institute of Science

Brief Biography. Swaprava is a Fulbright-Nehru Post-doctoral fellow at the Computer Science Department, Carnegie Mellon University. Earlier, he was a Lecturer and Post-doctoral Fellow at the Economics and Planning Unit, Indian Statistical Institute, New Delhi. He completed his PhD from the Dept. of Computer Science and Automation, Indian Institute of Science, Bangalore. His research interest is in the area of strategic multi agent systems, which includes crowdsourcing, resource allocation, online advertising, auctions, matching, social networks. He has completed internships at Xerox Research Centre, Europe (XRCE), in 2010, and at Harvard University, in 2011. He is a recipient of the Honorable Mention Award of Yahoo! Key Scientific Challenges Program, 2012. His PhD research was supported by the Tata Consultancy Services PhD Fellowship, 2010. In 2015, he received the Fulbright-Nehru post doctoral grant to continue his research on Internet Economics at Carnegie Mellon University.

Research Summary. My research is motivated by the strategic interaction between multiple decision making agents (mainly on the Internet), where "efficient" economic decisions are to be taken in a computationally efficient manner, e.g., in online advertising, (spectrum) resource allocation, crowdsourcing etc. I take an axiomatic approach to design mechanisms (primarily with money) that are robust against the conflicting interests of the agents, and blends the analysis techniques of computer science and social choice theory.

My PhD thesis looked at several aspects of crowdsourcing. The first question involves eliciting the "private" skills of the crowd-workers which can be time-varying. Using the assumption of Markovian skill evolution, we provide mechanisms that are truthful in this domain, which blends analysis techniques of Markov decision processes and mechanism design with interdependent valuations. The second approach considers crowdsourcing contests (a la DARPA red balloon challenge 2009) where certain novel strategic behavior is observed and lays down the design desiderata. We show that not all properties can be satisfied simultaneously and provide tight approximation guarantees to them [1].

In my postdoctoral research, I considered several axiomatic problems: the first is of resource allocation where individual preferences are "selfish", which appears in the context of spectrum auctions, cloud computing, and divisible combinatorial auctions. We provide a characterization of the strategyproof mechanisms in this domain [2]. In the second problem, we consider agent preferences that are separable into several components and show that the structure of the truthful decision schemes are decomposable into components. The third problem involves computational voting, where every voter is capable of only presenting a ranking over the alternatives, depending on which the mechanism has to pick a fixed subset of alternatives, and we provide tight approximation to the optimal social welfare [3]. Finally, our working paper on the trade-off between efficiency and budget balance provides the limits of approximability of these two properties (these two properties are known to be unsatisfiable together for strategyproof mechanisms, [Green-Laffont 1979]) in the general quasi-linear domain.

In summary, my research has contributed to the field of social choice with axiomatic and computational characterizations of several classical and contemporary problems.

Representative Papers.

[1] Mechanism Design for Time Critical and Cost Critical Task Execution via Crowdsourcing (WINE 2012)
with P. Dayama, D. Garg, Y. Narahari, and J. Zou

[2] Affine Maximizers in Domains with Selfish Valuations (TEAC 2015)
with A. Sen
[3] Subset Selection Via Implicit Utilitarian Voting (IJCAI 2016)
with I. Caragiannis, A. Procaccia, and N. Shah

ILAN NEHAMA

Thesis. Computational Issues in Judgment Aggregation

Advisor. Noam Nisan, Hebrew University of Jerusalem

Brief Biography. Ilan received his B.A. in Math and B.Sc. in Computer Science at the Technion (Summa Cum Laude), and M.A. in Computer Science with specialization in Rationality at The Hebrew University under the supervision of Prof. Gil Kalai; Thesis: Implementing Social Choice Correspondences using k-Strong Nash Equilibrium (Summa Cum Laude, GPA: 97.77, 2/194). Ilan submitted his Ph.D. at The Hebrew University (Benin School of Computer Science & Federmann Center for the Study of Rationality) on June 2016. During his graduate studies, he served as a lecturer in a programming course, as well as a TA in several MA courses both in the Computer Science department - Mathematical Tools in CS, and the Economics department - Microeconomics A & Microeconomics B: Game Theory and Information Economics. Ilan's works are mainly in theoretical Game Theory, Social Choice, and Judgment Aggregation, and on computational aspects and the usage of methods from Computer Science in these fields.

Research Summary. I am interested in the computer science approach to questions in social choice, game theory, and microeconomics, and specifically to judgment aggregation questions. Judgment aggregation investigates which procedures a group could or should use to form collective judgments on a given set of propositions or issues, based on the judgments of the group members.

One of my interests is shedding light on phenomena in judgment aggregation using approximation and perturbation viewpoints. That is, studying the way phenomena studied in the literature change when perturbing the classic strict properties. E.g., requiring an aggregation rule to satisfy a property with high probability, or generalizing players' rationality constraints to being close to rational (bounded rationality). For example, in "Approximately Classic Judgment Aggregation" I studied the perturbations of the "Doctrinal Paradox" scenarios. In these scenarios, one looks for "consistent" and "independent" aggregation mechanism for a set of permissible opinions. I presented a relaxation where these two desired properties hold only with high probability, and showed that under uniform distribution, for two basic agendas, there is no non-trivial mechanism that satisfies the perturbed constraints. In subsequent works, I show similar results for most of the truth-functional agendas, and for many other distributions, although all having independent representation. A challenging technical point in this research agenda is dealing with non-uniform distributions. The vast majority of analytic works of distributions of interesting social choice phenomena (such as Condorcet cycles) are w.r.t. uniform or almost uniform distributions, and I would like to extend these works to more natural distributions of the opinions (e.g., Polya-Eggenberger models).

During my PhD, I have also worked on analysis of the computational complexity of problems in aggregation scenarios, and studied decision under ambiguity and equilibria when there is ambiguity regarding the types of the players.

While the main stream of my research agenda is the above, I also researched topics outside of the scope of judgment aggregation, both in ad-hoc small research groups for a specific research question (e.g., in mechanism design, or defining power indexes in voting scenarios), or individual works I found myself working on (my work on decision under ambiguity).

Representative Papers.

[1] Approximately Classic Judgment Aggregation (AMAI 2013 & WINE 2011)

[2] Mechanism design on discrete lines and cycles (EC 2012)
with E. Dokow, M. Feldman, and R. Meir

[3] Analyzing games with ambiguous players types using the MINthenMAX decision model (Meeting of the Society for Social Choice and Welfare 2016 & World Congress of the Game Theory Society 2016.)

RAD NIAZADEH

Thesis. Robustness and simplicity of mechanisms and algorithms

Advisor. Robert Kleinberg, Cornell University

Brief Biography. Rad Niazadeh is a Ph.D. candidate in computer science at Cornell University, advised by Robert Kleinberg. He holds his M.Sc. and B.Sc. in electrical engineering from Sharif University of Technology, Iran. His research lies in different areas of theoretical computer science, with a focus on economic aspects of algorithms. In particular, he is interested in algorithmic mechanism design and online algorithms, with applications to electronic commerce and computational advertising. Rad was a long-term visitor at Microsoft Research New England in 2014, where he was also an intern during summer 2015. During Fall 2015, he was a research intern at Yahoo! Research labs, and he is an intern at Microsoft Research Redmond during summer 2016. Rad was awarded the Google PhD Fellowship in 2016, and Jacobs Scholar Fellowship in 2011.

Research Summary. I study theory problems stemming from mechanism design and online algorithms, with an emphasis on the themes of simplicity, practicality, and generality. I gravitate toward problems whose motivation stems from real-world applications and whose analysis leverages techniques from applied probability and high-dimensional optimization. The following paragraphs summarize some of my past contributions and ongoing work.

Simple and near-optimal mechanisms: Contemporary research in algorithmic mechanism design often grapples with the question of whether simple mechanisms can approximate the performance of optimal ones. One of the most well-known instances of this question, due to Hartline and Roughgarden, pertains to anonymous posted pricing in a single-item auction with independent (non-identical) bidders. Our work supplies the best known upper and lower bounds on how well anonymous pricing approximates optimal revenue, and shows that our upper bound is tight with respect to the "ex-ante relaxation". I also worked on affine fee-setting mechanisms in markets with intermediaries; such mechanisms are commonly applied by eBay and Amazon, for example. For single item trade, our work supplies mild conditions on the buyer side under which the best such mechanism can obtain a constant fraction of the intermediary's optimal revenue for all seller distributions.

Robustness in online algorithms: Pessimistic performance guarantees of worst-case analysis are a strong stimulus for assuming random arrival order of the input in many online problems. However, in most applications the input ordering is random, but not uniformly so. When do algorithmic performance guarantees proven under the uniform-random-ordering assumption remain valid under weaker assumptions? My work initiates a systematic study of this question, proving that existing algorithms for the classical secretary problem remain near-optimal under a wide class of distributions, including "almost all" distributions with at least log(log n) bits of entropy.

I have also worked on the open problem of black-box reductions to convert any algorithm into a Bayesian incentive compatible mechanism with bounded welfare loss in expectation. Recently we showed an efficient reduction for general type spaces by utilizing sophisticated tools from applied probability. My other research includes core selecting algorithms for sale of space, competitive equilibria for non-quasilinear agents, and mechanism design for value maximizers.

Representative Papers.

[1] Optimal Auctions vs Anonymous Pricing (FOCS 2015)
with S. Alaei, J. Hartline, Y. Yuan, and M. Pountourakis

[2] Secretary Problems with Non-Uniform Arrival Order (STOC 2015)
with T. Kesselheim and B. Kleinberg

[3] Simple and Near-Optimal Mechanisms For Market Intermediation (WINE 2014)
with Y. Yuan and B. Kleinberg.

EMMANOUIL POUNTOURAKIS

Thesis. Simple Mechanisms in Static and Dynamic Settings

Advisor. Nicole Immorlica, Microsoft Research

Brief Biography. Emmanouil Pountourakis is a PhD candidate in the Theory and Economics group of the Department of Electrical Engineering and Computer Science at Northwestern University, advised by Nicole Immorlica. Since 2014 he has been a long term visitor at Microsoft Research, New England. He holds an undergraduate and Masters degree in Computer Science from the University of Athens. During Summer 2012 he completed a research internship at CWI, Amsterdam. He was a student visitor in the Institute of Advanced Studies at Hebrew University for the Special Semester in Algorithmic Game Theory in 2011. Emmanouil Pountourakis has a broad interest in algorithmic mechanism design. His current research focuses on revenue maximization in static and dynamic environments, and theoretical analysis of behavioral models of time-discounting. In the past he has worked on a variety of topics including cost-sharing, matching, and mechanism design without money.

Research Summary. My research interests lie in the intersection of theoretical computer science, economics, and game theory. The main focus of my work has been in algorithmic mechanism design, the theory of algorithm design in strategic environments. In my earlier research, I studied the performance loss of optimal algorithms in the presence of incentive and computational constraints in the areas of group-strategyproof cost-sharing, stable marriage, and mechanism design without money.

In these explorations I have been intrigued by the fact that despite how well several aspects of mechanism design are understood, theoretically optimal mechanisms are rarely observed in practice. This phenomenon, the misalignment of predicted and observed behavior, also occurs in the agents participating in the mechanism. My current research focuses on understanding the observed practices and providing theoretical justification for their prevalence.

One such example is revenue maximization using simple auctions for the single item setting. Recently, I provided a tight analysis of the approximation ratio of anonymous pricing to the standard upper bound benchmark of ex-ante relaxation. It implied tightened upper bounds for the approximation of optimal auction from four to $e \approx 2.718$ for both anonymous posted pricing and second price auction with anonymous reserve, providing the first improvement to this open question in the last half decade.

Furthermore, one of my on-going works studies anonymous pricing in the dynamic environment of selling a single item over several periods. Our goal is to study the revenue of the seller in a perfect Bayesian equilibrium. Unfortunately, there are several negative results which show that the seller's revenue can be arbitrarily low due to the lack of commitment power. Surprisingly, my current results indicate that restricting the seller to use an anonymous pricing scheme alleviates these negative results and allows good revenue guarantees.

Lastly, I am interested in the interaction of revenue maximization and different behavioral models. My recent work studies optimal contract design with a present-biased agent. When making a decision, the present-biased agent overestimates her present utility by a multiplicative factor. The contract designer exploits this behavior to maximize their revenue. My work introduces regulations and studies optimal contracts under them with the aim of reducing the exploitative power of the mechanism.

Representative Papers.

[1] Procrastination with Variable Present Bias (EC 2016)
with N. Gravin, N. Immorlica, and B. Lucier

[2] Optimal Auctions vs Anonymous Pricing (FOCS 2015)
with S. Alaei, J. Hartline, R. Niazadeh, and Y. Yuan.

[3] Mechanisms for Hiring a Matroid Base without Money (SAGT 2014)
with G. Schäfer

CHRISTOS ALEXANDROS PSOMAS

Thesis. Algorithmic Mechanism Design in Dynamic Environments

Advisor. Christos Papadimitriou, University of California, Berkeley

Brief Biography. Christos-Alexandros (Alex) Psomas is a PhD candidate in the Department of Electrical Engineering and Computer Sciences at UC Berkeley, advised by Christos Papadimitriou. Prior to that, Alex received his MSc in Logic, Algorithms and Computation from the University of Athens,Greece, and his BSc from the Department of Informatics of Athens University of Economics and Business, in

Athens Greece. During his PhD Alex has spent two summers as a research intern at the International Computer Science Institute with Eric Friedman, a summer at Microsoft Research Redmond with Nikhil Devanur, and a summer as an instructor for the Discrete Mathematics and Probability Theory course at UC Berkeley.

Research Summary. My research to date has been in the field of Algorithmic Mechanism Design. I have worked both on fair division problems, as well as revenue maximization. A key topic of my recent work has been dynamic environments.

In my work on revenue maximization I have explored the complexity of optimal dynamic auctions. In mechanism design we typically study one-shot mechanisms. However, multi-stage environments are more common in the real world. In recent work, we show that restricting to static mechanisms results in very significant (super constant) losses in revenue. Why is it then that dynamic mechanisms are not more prevalent in practice? Surprisingly, we prove that the problem of maximizing revenue from selling two items across two rounds, arguably the simplest meaningful dynamic mechanism design problem imaginable, is computationally intractable. This result provides a partial answer to the above question.

In another line of research, I study dynamic environments where monetary transfers cannot be used, for example resource allocation. When a system designer wants to allocate a unit of shared memory between agents arriving and departing over time, when N players are present the solution is clear: everyone gets an N-th. If agents are arriving and departing constantly, the obvious caveat of this solution is its "disruptiveness": for a new agent to get her fair share, everyone has to give a little piece. What if, at every arrival we could only take back resources from a fixed number of agents k? We showed that there exists a non wasteful such algorithm that is almost optimal with respect to fairness.

More recently, I have been interested in the sample complexity of revenue maximization in various settings. We study a new model of a priori identical agents, that can be distinguished based only on some side information. We show almost matching upper and lower sample complexity bounds, and en route, improve bounds on previous models.

Representative Papers.

[1] On the Complexity of Dynamic Mechanism Design (SODA 2016)
with C. Papadimitriou, G. Pierrakos, and A. Rubinstein

[2] Dynamic Fair Allocation with Minimal Disruptions (EC 2015)
with E. Friedman and S. Vardi

[3] Sample Complexity of Auctions with Side Information (STOC 2016)
with N. Devanur and Z. Huang

RYAN ROGERS

Thesis. Private and Robust Mediators in Large Games

Advisor. Aaron Roth & Michael Kearns, University of Pennsylvania

Brief Biography. Ryan Rogers is currently a PhD candidate in the Applied Mathematics department at the University of Pennsylvania, where he is co-advised by Michael Kearns and Aaron Roth. He is president and co-founder of the SIAM student chapter at Penn. Ryan was a research intern with the Privacy Tools Project at

Harvard where he mentored an undergraduate student. He will be spending summer 2016 as a research intern at Microsoft Research NYC. After graduating Summa Cum Laude from Stetson University with a B.S. in mathematics, he worked as a programmer on the International Space Station Training Facility at the Johnson Space Center. He then went on to continue his studies by enrolling in Part III Math at Cambridge University, where he graduated with Distinction and worked with Richard Weber on his Part III essay in Algorithmic Game Theory. Ryan is also an avid rower, having competed in college and throughout graduate school. He currently rows and coaches on the Schuylkill River in Philadelphia.

Research Summary. My research involves applying differential privacy to mechanism design. Differential privacy provides an algorithmic stability condition which limits the sensitivity of the outcome from a mechanism to the decisions of a single agent. This stability guarantee turns out to be a powerful tool in mechanism design. The results from this fusion can be split roughly into two areas: enforcing privacy as a constraint in otherwise classical game theory problems, and using differential privacy as a tool to solve new problems in game theory where privacy may not be a first order concern.

Many crucial applications of mechanism design deal with sensitive information, e.g. medical data. It is then important to add privacy as a constraint on solutions of fundamental problems in game theory while still managing agents' incentives. I then study the impact that privacy has on existing problems in mechanism design and what may or may not be achievable subject to privacy. In particular, my work on barter-exchange economies shows that mechanisms which are differentially private, or even a more relaxed version of it, cannot obtain Pareto optimal allocations. However, I provide a mechanism with a less restrictive, but still strong guarantee of privacy that can obtain an asymptotical exact Pareto optimal allocation as the number of agents grows. A key application of this is to markets for kidney exchanges, which deals with highly sensitive data.

Differential privacy also serves as a powerful tool in coordinating agents to an equilibrium strategy in settings of incomplete information. Such scenarios are more realistic than assuming agents have complete information, especially if the number of agents is large. However, games in incomplete information settings lose many of the nice properties of complete information games, including the quality of equilibria becoming worse and coordinating to an equilibrium becoming even more difficult. It is then desirable in incomplete information games to be able to compute equilibria as if the agents knew the types of everyone else. I then introduce a mediator that takes as input the utilities of the agents and outputs a suggested action for each agent. I show that it is possible to construct a mediator in congestion games where agents cannot significantly improve their utility by misreporting their types or not following the suggestion. Further, the resulting action profile will lead to a Nash equilibrium of the complete information game.

Representative Papers.

[1] Asymptotically Truthful Equilibrium Selection in Large Congestion Games (EC 2014) with A. Roth

[2] Private Pareto Optimal Exchange (EC 2015)

with S. Kannan, J. Morgenstern, and A. Roth

[3] Do Prices Coordinate Markets, (STOC 2016)
with J. Hsu, J. Morgenstern, A. Roth, and R. Vohra

ALAN ROYTMAN

Thesis. Making Decisions Under Uncertainty for Large Data Domains

Advisors. Rafail Ostrovsky, University of California, Los Angeles

Brief Biography. Alan Roytman is a postdoctoral student of computer science at Tel Aviv University, where he holds the I-CORE Postdoctoral Fellowship. He is broadly interested in theoretical computer science, and more specifically in game theory, online algorithms, and streaming algorithms. Before coming to Tel Aviv University, he received his Ph.D. in computer science from the University of California, Los Angeles, where he was advised by Adam Meyerson and Professor Rafail Ostrovsky. His Ph.D. thesis focused on applications dealing with large data that arrives on the fly, particularly online algorithms for load balancing problems and streaming algorithms for clustering problems. He received his B.A. in mathematics and computer science from the University of California, Berkeley.

Research Summary. My research mostly lies within theoretical computer science, and more specifically the area of algorithms, including algorithmic game theory, online algorithms, and streaming algorithms. My research interests in game theory concern measuring the inefficiency of equilibria for various solution concepts, along with mechanism design.

Much of my research in game theory concerns studying notions related to the Price of Anarchy in a variety of settings. In particular, we are the first to define the notion of the Price of Mediation, which is the ratio of the worst correlated equilibrium to the worst Nash equilibrium. This ratio aims to quantify how much harm a selfish and inept mediator can cause to society. We give various upper and lower bounds on this ratio for general matrix games and load balancing games. My work also focuses on understanding how the presence of budgets affects the inefficiency of equilibria in simple auction formats, such as simultaneous first price and second price auctions. Moreover, in the context of mechanism design, my work includes approximating social welfare for auction formats that are motivated by crowdsourcing applications.

I also have research interests in the area of online algorithms (where the input to the algorithm arrives on the fly). Many of my interests concern problems motivated by energy efficiency in the cloud and data centers. For instance, I have results for the setting where tasks in data centers are modeled via multidimensional vectors (where each coordinate corresponds to various components on a machine). In particular, I have studied load balancing problems and bin packing problems within this framework. In the load balancing line of work, we have designed algorithms that are simultaneously competitive against two benchmarks that are at odds with each other: energy efficiency and quality of service. In the context of bin packing, we give tight results for the setting where vectors are small relative to each bin's capacity.

Representative Papers.

[1] On the Price of Mediation (EC 2009)
with M. Bradonjić, G. Ercal-Ozkaya, and A. Meyerson

[2] Streaming k-means on Well-Clusterable Data (SODA 2011)
with V. Braverman, A. Meyerson, R. Ostrovsky, M. Shindler, and B. Tagiku

[3] Packing Small Vectors (SODA 2016)
with Y. Azar, I.R. Cohen, and A. Fiat

AVIAD RUBINSTEIN

Thesis. Hardness of Approximation between P and NP

Advisor. Christos Papadimitriou, University of California, Berkeley

Brief Biography. Aviad is a PhD candidate at UC Berkeley, advised by Christos Papadimitriou. Prior to coming to Berekely, he completed an M.Sc. at Tel-Aviv University (2013, advised by Muli Safra), and a B.S. at the Technion (2011). During his PhD he spent 3.5 summers interning at Microsoft Research labs in Hertzeliyah, Beijing, and Boston. Aviad won the Best Student Paper Awards at STOC'15 and ITCS'16.

Research Summary. My research has looked, through the Lens of Theoretical Computer Science, at a variety of problems from the studies of Evolution, Statistics, Stopping Theory, and Game Theory. Two lines of works that I am particularly thrilled about are (1) understanding the terrain of computational complexity between P and NP, and (2) understanding mechanism design subject to simplicity constraints.

A common approach in Computer Science to intractability barriers is approximation. Over the past few decades, many tools have been developed to show that even approximation is NP-hard: the PCP Theorem, for example. My research has focused on problems whose complexity lies between P and NP: they are still intractable, but the standard techniques for showing NP-hardness also break. For example, a heuristic algorithm for Approximate Competitive Equilibrium from Equal Income is used in practice for fair allocation of classes to students; we showed that it is PPAD-complete: an algorithm with provable guarantees is out of reach, yet this problem is unlikely to be NP-hard because an approximate equilibrium always exists. A different obstacle to NP-hardness are approximation algorithms that run in quasi-polynomial time: not quite fast enough to run in practice, yet also not NP-hard unless the Exponential Time Hypothesis (ETH) breaks. Assuming ETH, we can show matching hardness of approximation for fundamental problems such as Densest-k-Subgraph and Community Detection. The most exciting hardness of approximation result is for Approximate 2-Player Nash Equilibrium: it lies in the intersection of PPAD and quasi-polynomial algorithms, and the proof of intractability has to deal with both simultaneously.

Classical work in Economics imposes two constraints on mechanism design: Individual Rationality (IR) and Incentive Compatibility (IC). In recent years it has been shown that even in very simple settings, the revenue-maximizing mechanism is extremely complicated: it requires randomization, has infinite menu-size complexity, and is computationally intractable. As an important step toward bridging theory and practice, we propose adding a third constraint: Simplicity. We have explored the possibilities of Simple Mechanism Design in challenging settings where buyer valuations may have complex combinatorial structure, or evolve over time. Another, more open-ended question I have been interested in is: How should we define "simple"?

Representative Papers.

[1] Inapproximability of Nash equilibrium (STOC 2015)

[2] On the Complexity of Dynamic Mechanism Design (SODA 2016)
with C. Papadimitriou, G. Pierrakos, C.-A. Psomas

[3] Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity (EC 2015) with S.M. Weinberg

SHREYAS SEKAR

Thesis. Incentivizing Self-Interested Agents to form Socially Desirable Solutions

Advisor. Elliot Anshelevich, Rensselaer Polytechnic Institute

Brief Biography. Shreyas Sekar is heading into the final year of his Ph.D at Rensselaer Polytechnic Institute, where he is advised by Elliot Anshelevich. Broadly speaking, Shreyas is interested in research questions falling under the umbrella of "incentivizing self-interested agents in practical scenarios". More specifically, his research tackles problems pertaining to algorithmic pricing, coalition formation and approximation algorithms in social choice domains. He also briefly interned in Microsoft Research, Cambridge, UK, where he worked on the economics of the cloud under Peter Key. Shreyas holds a bachelors degree in Electronics and Communication Engineering from the Indian Institute of Technology at Roorkee

Research Summary. My research attempts to answer questions of the form "How do we design algorithms and incentives in large systems involving self-interested agents?". The broad nature of this question drives my research into seemingly disparate fields such as market pricing, cooperative game theory and computational social choice. The structural thread underlying most of my work involves a combinatorial optimization problem with novel constraints on the input and output arising due to the selfish nature of the agents. Against this backdrop, I am interested in algorithmic solutions that benefit both the designer as well as the participating agents. I will briefly unravel this thread to discuss my research agenda in different domains where such problems arise.

A significant portion of my research involves developing pricing algorithms for markets consisting of a large number of buyers. The scale at which many real-life markets operate preclude all but the most simple of pricing mechanisms: namely posted prices. Moreover, such markets are often characterized by repeated engagement and so, a disproportionate focus on revenue could result in disastrous long-term effects. Bearing this in mind, I developed pricing mechanisms with good revenue that are simultaneously fair, non-discriminatory to the buyers, culminating in constant factor bi-criteria approximations for (revenue, social welfare). In parallel work, I have sought to model and understand how many of the constraints arising in actual markets affect the quality of equilibria: these include mechanisms for markets with competing sellers, and convex production costs.

Another type of constraint that arises in large systems concerns the limited ability of the agents to express their utilities for various solutions. At the same time, it is more convenient for individuals to express their preferences over the set of solutions in terms of rankings. Motivated by this, my research looks at the design of ordinal algorithms, and lies at the intersection of approximation algorithms and social choice mechanisms. Ordinal approximation algorithms act as a bridge between the two worlds, and demand solutions that only use preference rankings as input but provide approximation guarantees with respect to the true utilities. In recent work, I was able to provide the first non-trivial ordinal mechanisms for crucial graph problems such as Matching and Clustering. Exploring the power of such ordinal algorithms in other domains appears to be a promising research avenue.

Representative Papers.

[1] Envy-Free Pricing in Large Markets: Approximating Revenue and Welfare (ICALP 2015) with E. Anshelevich and K. Kar

[2] Computing Stable Coalitions: Approximation Algorithms for Reward Sharing (WINE 2015) with E. Anshelevich

[3] Blindy, Greedy, and Random: Algorithms for Matching and Clustering using only Ordinal Information (AAAI 2016) with E. Anshelevich

CHRISTOS TZAMOS

Thesis. Revenue Optimal Mechanism Design

Advisor. Costis Daskalakis, Massachusetts Institute of Technology

Brief Biography. Christos is a PhD Student with Costis Daskalakis in the Theory of Computation group at MIT. Prior to joining MIT, he completed his undergraduate studies in EECS at the National Technical University of Athens, where he did research with Prof. Dimitris Fotakis. His research interests include algorithmic game theory and learning theory, and more broadly the design, analysis and theory of algorithms. During his time at MIT, Christos spent two summers interning at Google and Yahoo Labs and has mentored two high school students in AGT research through the MIT PRIMES program. Christos is the recipient of the Simons Graduate Award in Theoretical Computer Science and his research has received the best paper award at the 2013 ACM Conference on Economics and Computation.

Research Summary. My research focuses on the intersection of Theoretical Computer Science and Economics. I draw tools and perspective from Computer Science to understand complex economic systems and develop efficient algorithms for their design. Some highlights of my research include the following.

Revenue-Optimal Auctions: The design of auctions that maximize the seller's revenue has been a central problem in Economics for decades. Despite intense research, our knowledge beyond single-item settings is quite limited. To understand the challenges underlying this problem, we studied its combinatorial structure and established a deep connection to a well-studied problem in Computer Science, that of min-cost matching in high dimensional spaces. Exploiting this connection, we showed that designing revenue-optimal auctions is computationally intractable, even in simple settings involving a single buyer with discrete and independently distributed values for the items. To overcome the intractability, we study settings with continuously distributed values and develop an analytical framework for identifying optimal multi-item mechanisms. Our framework is based on optimal transport theory, which is a continuous analog to min-cost matching. Our intractability result appeared in SODA'14, while our characterization framework based on optimal transport first appeared in EC'13 (best paper award) and its strengthening was presented in EC'15 and has been accepted to Econometrica.

Nash-Equilibria in Anonymous Games: Anonymous games are games where a large number of players share the same strategies and the expected payoff of each player only depends on his chosen strategy and the distribution of how many other players choose each of the available strategies. Such distributions, aggregating independent randomizations over a set, are known as Poisson-Multinomial distributions, and are one of the most widely studied families of multi-dimensional distributions. In recent work, we use tools from probability theory to understand their structure, characterizing their similarity to rounded multi-dimensional Gaussian distributions, strengthening the guarantees of the Central Limit Theorem. Exploiting our structural understanding we obtain efficient algorithms for computing approximate Nash equilibria in anonymous games, significantly improving the state of the art. Our main structural result about Poisson Multinomials appeared in FOCS'15, and its application to anonymous games appeared in STOC'16.

Representative Papers.

[1] A Size-Free Central Limit Theorem for Poisson Multinomials and its Applications (STOC 2016) with C. Daskalakis, A. De, and G. Kamath

[2] Strong Duality for a Multiple-Good Monopolist (EC 2015 and Econometrica) with C. Daskalakis and A. Deckelbaum

[3] The Complexity of Optimal Mechanism Design (SODA 2014) with C. Daskalakis and A. Deckelbaum

JAMES R. WRIGHT

Thesis. Modeling Human Behavior in Strategic Settings

Advisor. Kevin Leyton-Brown, University of British Columbia

Brief Biography. James Wright is a Ph.D. candidate in computer science at the University of British Columbia, advised by Kevin Leyton-Brown. He holds an M.Sc. from the University of British Columbia (2010) and a B.Sc. from Simon Fraser University (2000). He studies problems at the intersection of behavioral game theory and computer science, with a focus on applying both machine learning techniques and models derived from experimental and behavioral economics to the prediction of human behavior in strategic settings. He also studies the implications of behavioral game theoretic models on multiagent systems and mechanisms. James's expected graduation date is June 2016.

Research Summary. A wealth of experimental evidence demonstrates that human behavior in strategic situations is often poorly predicted by classical economic models. Behavioral game theory studies deviations of human behavior from the standard assumptions, and provides many models of these deviations. These models typically focus on explaining a single anomaly. Although understanding individual anomalies is valuable, the resulting models are not always well-suited to predicting how people will behave in generic settings, which limits their application to questions of interest in algorithmic game theory, such as "What is the optimal mechanism for implementing a particular objective?".

I am interested applying machine learning techniques to construct behavioral game theoretic models that have high predictive accuracy, and in applying these models to problems in algorithmic game theory. As an example of the first direction, I previously analyzed and evaluated behavioral models in simultaneous-move games, eventually identifying a specific class of models (iterative models) as the state of the art. I then proposed and evaluated an extension that improves the prediction performance of any iterative model by better incorporating the behavior of nonstrategic agents.

Despite growing interest in behavioral game theory over the past decade, many important questions about its application to areas such as mechanism design remain open. For example, foundational analytic techniques such as the revelation principle may not be straightforwardly applicable under some classes of behavioral model. One direction of my current research aims to determine to which classes of behavioral model do such principles apply, and how to handle the cases where they don't apply.

Representative Papers.

[1] Beyond Equilibrium: Predicting Human Behavior in Normal-Form Games (AAAI 2010) with K. Leyton-Brown

[2] Behavioral Game-Theoretic Models: A Bayesian Framework for Parameter Analysis (AAMAS 2012) with K. Leyton-Brown

[3] Level-0 Meta-Models for Predicting Human Behavior in Games (EC 2014) with K. Leyton-Brown

STEVEN WU

Thesis. Privacy and Learning in Strategic Environments

Advisors. Aaron Roth & Michael Kearns, University of Pennsylvania

Brief Biography. Steven Wu is currently a PhD student in the Computer Science department at the University of Pennsylvania, co-advised by Aaron Roth and Michael Kearns. His research focuses on the areas of data privacy, algorithmic economics and learning theory, with an emphasis on algorithm design under privacy, incentive and robustness considerations. His works appear in a varierty of venues including PNAS, STOC, ICML, SODA, EC, COLT, ITCS and WINE. During his PhD, he was a research intern at Microsoft Research, once at the New England lab, and twice at the New York City lab.

Research Summary. My research focuses on data privacy, algorithmic economics and machine learning theory, and the fundamental connections between these areas. I study algorithmic questions in these areas by drawing techniques from differential privacy, game theory, and convex optimization. Some of the key aspects are:

Privacy in Social and Economic Settings. While I broadly work on fundamental problems in differential privacy, I especially like tackling privacy problems in social and economic settings where data are held by strategic agents that are connected with each other. In our recent PNAS'16 paper, we consider the problem of using network meta-data to guide a search for some class of targeted individuals (for example, terrorists, carriers of an infectious disease, or organized criminals). We proposed the notion of "protected differential privacy" and a family of graph search algorithms that discloses individuals in the targeted class, but promises privacy to (innocent) people not in the class. Another direction I have been pursuing is the solving a family of mechanism design problems under the constraint of joint differential privacy, a strong variant of differential privacy in a game theoretic environment. In this line of works with my coauthors, we also establish an interesting connection between joint differential privacy and incentive-compatibility – privacy as stability notion could be used as a tool to obtain truthfulness!

Learning in Strategic Environments. My research also focuses several machine learning problems in economic settings, in which I study the interesting interplay between learning and strategic behavior. For example, in our recent STOC'16 paper, we study the problem of profit maximization from revealed preferences feedback – given the ability to set prices and observe buyers' purchased bundle, how can a seller learn to price the items optimally? We give an efficient algorithm which is also applicable to a wide range of Stackelberg games. In another paper that appears in EC'16, we study the problem of incentivizing agents to do bandit exploration for a social planner. For example, in order to for a navigation platform such as Waze to optimize its route recommendations, it needs incentivize drivers to explore new routes. This problem presents a three-way tradeoff between exploration, exploitation and incentives, and our solution nicely connects online learning theory with techniques in information design.

Representative Papers.

[1] Private Algorithms for the Protected in Social Network Search (PNAS 2016) with M. Kearns, A. Roth, and G. Yaroslavtsev

[2] Watch and Learn: Optimizing from Revealed Preferences Feedback (STOC 2016) with A. Roth and J. Ullman

[3] Bayesian Exploration: Incentivizing Exploration in Bayesian Games (EC 2016) with Y. Mansour, A. Slivkins and V. Syrgkanis

QIANG ZHANG

Thesis. Algorithmic Mechanism Design on Distributed Computing Environment and Facility Location

Advisors. Xiaotie Deng, Shanghai Jiaotong University & Minming Li, City University of Hong Kong

Brief Biography. Qiang is currently a postdoc in the Department of Computer, Control and Management Engineering Antonio Ruberti at Sapienza University of Rome, hosted by Stefano Leonardi. From 2013-2015 he was a postdoc in the Institute of Informatics at University of Warsaw, hosted by Piotr Sankowski. Prior to that, Qiang obtained his Ph.D. in Computer Science from City University of Hong Kong in 2013, advised by Xiaotie Deng and Minming Li. During this Ph.D. study, he spent six months as a visiting student at Northwestern University hosted by Jason Hartline in 2011 and three months at Aarhus University hosted by Peter Bro Miltersen in 2013.

Research Summary. My research lies in the interaction between economics and computer science. My works can be categorized into two following aspects.

First, I investigated the design of efficient pricing algorithms in both static and dynamic settings. The pricing algorithms aim to maximize the revenue of sellers and guarantee some nice properties at the same time. For example, envy-free pricing and allocations are commonly used to model fair equilibrium in a variety of economic settings. Fairness is an important aspect in practical applications as the loss of fairness could cause users to unsubscribe the services. With my collaborators, we studied the revenue-maximizing envy-free pricing problems in various marketplaces. In most of the practical settings, we provided the optimal (or near-optimal) solutions. Besides the static settings, to incorporate the dynamics in markets, in a recent work we studied the pricing algorithms in the setting where the algorithms decide the price of the item at each moment and the buyers in the markets come and leave. However, the challenge is that the algorithms have no knowledge on when the buyers will leave the market. We were able to give a complete answer for this question by providing the best possible algorithms.

Second, I am interested in the design of strategyproof mechanisms for resource allocations in multi-agent systems. It is well-known that it is impossible to design symmetric, Pareto optimal and strategyproof mechanisms for general preferences. However, such mechanisms do exist if the preferences possess a certain underlying structure. Following this direction, my works provide several strategyproof mechanisms with different desirable properties in settings such as cake cutting and facility location games. The results shed light on providing a unify framework to design incentive compatible protocols in large-scale distributed systems. Recently, I am studying the design strategyproof mechanisms for crowd-sourcing and procurement auctions.

Representative Papers.

[1] Online Pricing with Impatient Bidders (SODA 2016)
with M. Cygan, M. Mucha, and P. Sankowski

[2] Revenue Maximization Envy-free Pricing for Homogeneous Resources
(IJCAI 2015) with G. Monaco and P. Sankowski

[3] Facility Location with Double-peaked Preferences (AAAI 2015)
with A. Filos-Ratsikas, M. Li, and J. Zhang