Tutorial on Algorithmic Game Theory and Data Science

Jamie Morgenstern[1] and Vasilis Syrgkanis[2]

[1]University of Pennsylvania
[2]Microsoft Research, NYC

1 Description

The increasing availability of data-sets from large scale systems and markets, renders necessary the development of data-driven methodologies in fields such as mechanism design and game theoretic analysis. Moreover, many data-sets that we are asked to analyze in applications such as electronic markets, cloud computing, energy marketplaces and crowdsourcing systems, are the outcome of strategic interactions, rather than the outcome of stochastic natural processes. Thereby the use of game-theoretic techniques when addressing learning theoretic problems on such data-sets is mandatory.

Recent work in algorithmic game theory has began to address these two concerns and there is still room for much further development. Being at this crossing point the tutorial will serve two purposes; reviewing recent progress and formulating concrete open questions at the intersection of algorithmic game theory and data science. The tutorial will complement the corresponding workshop on algorithmic game theory and data science, by providing basic techniques and ideas, as well as placing the work presented at the workshop in a bigger scope.

The tutorial will consist of two parts. The first part will cover questions related to learning good mechanisms when having access to sample data-sets of private parameters of participants. This is in line with the direction of introducing data-driven methods in mechanism design and more generally of how data science can affect algorithm game theory. The second part will cover learning theoretic questions when the input data-sets are the outcome of strategic interactions or more broadly are coming from "strategic sources". Therefore, the second part will address topics related to how algorithmic game theory can influence data-science.

The first component of the tutorial will focus on "standard" sample complexity questions within mechanism design. We will first overview what is known about the sample complexity of revenue and welfare maximization, in both single-parameter and multiparameter settings. Many of these practical results show that low-dimensional classes of mechanisms can be learned with fewer samples, giving theoretical justification for trading off the number of parameters for revenue optimality when designing a mechanism. There are a wide variety of theoretical open problems, both in the small and large sample regimes, studying the optimal bicriterial guarantees for a variety of auction design problems. Most of these results hinge on a variety of assumptions about the distributions from which buyers are drawn. Practical work studying whether these assumptions are satisfied in real-world datasets can guide future theory and practice of designing mechanisms which generalize well from training to test.

In the second part of the tutorial, we will review recent progress in econometric theory for game theoretic settings. We will address questions, such as, i) how to perform parametric inference from observed

strategic interactions, ii) how does the structure of a game or mechanism affect the statistical properties of the resulting inference methods, iii) how to perform A/B testing in strategic market environments. A large part of this section will also cover techniques from the Economics literature and present the state of the art from the Econometrics community when addressing inference problems in game theoretic settings. This will help in two respects: i) to build upon existing results, ii) to understand how the computer science and algorithmic perspective can influence or inform existing results and questions already addressed by the econometric community.

2 Related Work

There is a small but growing body of work at the interface of AGT and data science. Below we list some relevant work addressing issues outlined above. The tutorial will touch upon several of the papers listed below, at various extends and depth levels.

Learning Optimal Mechanisms from Valuation Samples.

- R. Cole and T. Roughgarden, The Sample Complexity of Revenue Maximization. STOC'14.
- I. Segal. Optimal pricing mechanisms with unknown demand. AER'03.
- Z. Huang, Y. Mansour, and T. Roughgarden, Making the Most of Your Samples. EC'15.
- Hu Fu, Nima Haghpanah, Jason D. Hartline, and Robert Kleinberg. Optimal auctions for correlated buyers with sampling. EC'14.
- Kareem Amin, Afshin Rostamizadeh, Umar Syed. Learning Prices for Repeated Auctions with Strategic Buyers. NIPS'13.
- Yash Kanoria, Hamid Nazerzadeh. Dynamic Reserve Prices for Repeated Auctions: Learning from Bids. WINE'14.
- Hu Fu, Nicole Immorlica, Brendan Lucier, Philipp Strack. Randomization beats Second Price as a Prior-independent Auction. EC'15.
- Jamie Morgenstern and Tim Roughgarden. The Pseudo-Dimension of Nearly-Optimal Auctions. NIPS'15.
- Nikhil Devanur, Zhiyi Huang and Christos Alexandros Psomas. Sample Complexity of Auctions with Side Information. STOC'16.

Learning from Strategic Data.

- S. Chawla, J. Hartline, D. Nekipelov, Mechanism Design for Data Science. EC'14.
- S. Chawla, J. Hartline, D. Nekipelov, A/B testing of Auctions. Manuscript under submission.
- Kareem Amin, Michael Kearns, Peter Key and Anton Schwaighofer. Budget Optimization for Sponsored Search: Censored Learning in MDPs. UAI'12.
- D. Coey, B. Larsen and K. Sweeney. The Bidder Exclusion Effect. NBER Working Paper 20523.

- Kareem Amin, Rachel Cummings, Lili Dworkin, Michael Kearns, and Aaron Roth. Online Learning and Profit Maximization from Revealed Preferences. In the proceedings of AAAI 2015.
- Eyal Beigman, Rakesh Vohra. Learning from Revealed Preferences. In the proceedings of EC 2006.
- Morteza Zadimoghaddam, Aaron Roth. Efficiently Learning from Revealed Preference. WINE '12.
- Avrim Blum, Yishay Mansour, Jamie Morgenstern. Learning Valuation Distributions from Partial Observation. In the proceedings of AAAI 2015.
- Chien-Ju Ho, Aleksandrs Slivkins and Jennifer Wortman Vaughan. Adaptive Contract Design for Crowdsourcing Markets: Bandit Algorithms for Repeated Principal-Agent Problems. EC'14.
- Denis Nekipelov, Vasilis Syrgkanis and Éva Tardos. Econometrics for Learning Agents. EC'15.
- Darrell Hoy, Denis Nekipelov and Vasilis Syrgkanis. Robust Data-Driven Guarantees in Auctions. 1st workshop on Algorithmic Game Theory and Data Science, EC'15.

Empirical Work.

- E Bax, A Kuratti, P Mcafee. Comparing predicted prices in auctions for online advertising. IJIO'12.
- T Blake, C Nosko, S Tadelis. Consumer heterogeneity and paid search effectiveness: A large scale field experiment, Econometrica, 2015.
- M. Ostrovsky and M. Schwarz. Reserve prices in internet advertising auctions: A field experiment. EC'11.
- L Einav, T Kuchler, JD Levin, N Sundaresan. Learning from seller experiments in online markets. Working paper, 2014.
- Renato Paes Leme, Martin Pal and Sergei Vassilvitskii. A Field Guide to Personalized Reserve Prices. WWW'16.

3 Organization

The tutorial will consist of two primary parts, as well as a last smaller section with open questions in the area. The first part will cover questions related to learning good mechanisms when having access to sample data-sets of private parameters of participants. The second part will cover learning theoretic questions when the input data-sets are the outcome of strategic interactions or more broadly are coming from "strategic sources", rather than "non-strategic" samples.

4 Organizers

Jamie Morgenstern is a postdoctoral researcher in the Warren Center at the University of Pennsylvania, hosted by Michael Kearns, Aaron Roth, and Rakesh Vohra. She received her PhD in Computer Science from Carnegie Mellon University. Her research focuses on the application of tools from computer science (such as machine learning and privacy) to mechanism design and other problems within economics. She was the recipient of the Simons Fellowship for graduate students in theoretical computer science 2014-2016, an NSF GFRP fellowship, as well as the Microsoft Research Graduate Women's Scholarship.

Vasilis Syrgkanis is a postdoctoral researcher at Microsoft Research New York, where he is a member of the algorithmic economics and machine learning groups. He received his PhD in Computer Science from Cornell University. His research addresses problems at the intersection of theoretical computer science, machine learning and economics. His work received best paper awards at the 2015 ACM Conference on Economics and Computation (EC) and at the 2015 Annual Conference on Neural Information Processing Systems (NIPS). He was the recipient of the Simons Fellowship for graduate students in theoretical computer science 2012-2014.