Report from Dagstuhl Seminar 17251

Game Theory Meets Computational Learning Theory

Edited by

Paul W. Goldberg[1], Yishay Mansour[2], and Paul Dütting[3]

1 **University of Oxford, GB,** paul.goldberg@cs.ox.ac.uk
2 **Tel Aviv University, IL,** mansour@tau.ac.il
3 **London School of Economics, GB,** p.d.duetting@lse.ac.uk

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17251 "Game Theory Meets Computational Learning Theory". While there have been many Dagstuhl seminars on various aspects of Algorithmic Game Theory, this was the first one to focus on the emerging field of its intersection with computational learning theory.

Seminar June 18–23, 2017 – http://www.dagstuhl.de/17251

1998 ACM Subject Classification I.2 Artificial Intelligence, I.2.6 Learning (K.3.2), F. Theory of Computation, J.4 Social and Behavioral Sciences (economics)

Keywords and phrases Algorithmic Game Theory, Computational Learning Theory, Economics

Digital Object Identifier 10.4230/DagRep.7.6.68

1 Executive Summary

Paul W. Goldberg
Yishay Mansour
Paul Dütting

Algorithmic Game Theory (AGT) has been an identifiable research field for about 20 years by now. It emerged as an important research community in the 1990s, with the ACM-EC conference starting in 1999, and the conferences WINE and SAGT also support this community; in addition, the field is also represented in the main CS theory and AI conferences. Among former Dagstuhl seminars on topics in AGT, there have been a sequence of Dagstuhl seminars on Equilibrium Computation, and another sequence of seminars on Computational Social Choice, and also on Electronic Markets and Auctions.

Machine learning has of course become very pervasive, with the vast accessibility of "big data" and has the motivation to develop new methodologies for harvesting the vast amounts of data, improve our ability to automatically carry out many tasks, from classifying documents and pictures, to identifying normal trends and anomalies.

It is perhaps not surprising that it is timely to investigate the connections between economics and big data, more specifically the interface between game theory and machine learning. Much of econometrics is about handing data and deriving understanding from data.

In the AGT context, this would seem to apply most readily to data emanating from "economic" sources, and of course there are plenty of examples. The most notable of these is learning user preferences from examples.

There have been workshops at the AGT/Machine Learning interface in the ACM-EC conference (the 2017 EC held the 3rd workshop on Algorithmic Game Theory and Data

Science) but there is clearly space for more meetings on this topic, and this is the first one to take place at Dagstuhl. As such, it contributed to the development of the European community in this research area (noting that ACM-EC is usually held in the USA).

It was pleasing that the seminar attracted quite a high proportion of participants who were visiting Dagstuhl for this first time, alongside others who have made multiple visits. There was a good balance amongst representatives of Algorithmic Game Theory, Machine Learning, and Economics.

One can classify the AGT/Machine Learning topics as follows.

- Usage of ML ideas (reinforcement learning, multi-arm bandits, etc.) into decision making under uncertainty (and the search for game-theoretic solution notions such as equilibria)
- usage of game-theoretic tools into machine learning approaches (as in Generative Adversarial Nets).
- A basic test case is learning user valuations from historical data. For example, given the outcome of previous auctions to learn the distribution of the users' valuations and the goal is to define near optimal mechanisms. (This is also an aspect in learning "revealed preferences".)
- Query complexity of solution concepts of games, aspects of which are applicable to learning adversary preferences in the context of security/patrolling games.

The seminar was structured around longer invited/tutorial talks (typically lasting 1.5-2 hours), one or two such talks taking place each day. These were followed by shorter contributed talks.

We thank Argy Deligkas for serving as collector of the abstracts.

Keynote/tutorial talks

Monday Sven Seuken, University of Zurich Design of Machine Learning-Based Mechanisms; Yaron Singer, Harvard University: Learning, Optimization, and Noise

Tuesday Claudio Gentile, Universita dell'Insubria: No Regret and Sequential Prediction

Wednesday Denis Nekipelov, University of Virginia: Robust Inference for Non-Robust Models

Thursday Jamie Morgenstern, University of Pennsylvania: The Sample Complexity of Single-Parameter Auction Design

Friday Yakov Babichenko, Technion: Informational Bounds on Equilibria, and its Relation to Learning

Related topics not covered

There is ongoing work at the intersection of macroeconomics and machine-learning techniques, that was out of scope of this meeting but may be of interest later. For example, ongoing work on Vector Autoregressive models in the context of multivariate time series modelling, which may later lead to interesting problem in computational learning theory. Another topic that was only touched-on is agent-based models of macroeconomics.

2 Table of Contents

Executive Summary
Paul W. Goldberg, Yishay Mansour, and Paul Dütting 68

Overview of Talks

Forecast Aggregation
Yakov Babichenko 72

Aggregating Earnings-Per-Share Predictions
Amir Ban 72

The optimal strategy for a linear regression game
Peter L. Bartlett 72

Bayesian Methods for Market Clearing
Gianluca Brero and Sébastien Lahaie 73

Distributed Methods for Computing Approximate Equilibria
Argyrios Deligkas, John Fearnley, and Rahul Savani 73

Optimal Auctions through Deep Learning
Paul Dütting, Zhe Feng, Harikrishna Narasimhan, and David C. Parkes 74

Optimal Auctions for Correlated Bidders with Sampling
Hu Fu, Jason D. Hartline, Nima Haghpanah, and Robert Kleinberg 74

No Regret and Sequential Prediction
Claudio Gentile 75

Non-revelation Mechanism Design
Jason D. Hartline and Samuel Taggart 75

Optimizing Worst-case Benchmarks
Aleck Johnsen and Jason D. Hartline 76

Best-Response Dynamics in Combinatorial Auctions with Item Bidding
Thomas Kesselheim and Paul Dütting 76

Bayesian Methods for Clearing Markets
Sébastien Lahaie and Gianluca Brero 77

Submultiplicative Glivenko-Cantelli and Uniform Convergence of Revenues
Yishay Mansour 77

The Sample Complexity of Single-Parameter Auction Design
Jamie Morgenstern 78

Revenue Optimization with Approximate Bid Predictions
Andrés Muñoz Medina 78

Robust Inference for Games via Theoretical Guarantees
Denis Nekipelov 78

Weighted Voting Via No-Regret Learning
Ariel D. Procaccia 79

Design of Machine Learning-based Mechanisms
Sven Seuken 79

Learning, Optimization, and Noise
Yaron Singer 80

Learning from Untrusted Data
Gregory Valiant 81

Prediction with a Short Memory
Gregory Valiant 82

Sample Complexity of Multi-Item Profit Maximization
Ellen Vitercik 82

Learning Mastermind – Some Ideas and Challenges
Bernhard von Stengel 83

Peer Prediction Mechanisms and their Connections to Machine Learning
Jens Witkowski 83

Learning Cooperative Solution Concepts
Yair Zick 84

Participants 85

3 Overview of Talks

3.1 Forecast Aggregation

Yakov Babichenko (Technion – Haifa, IL)

Joint work of Itai Arieli, Yakov Babichenko, Smorodinsky Rann

Bayesian experts with a common prior that are exposed to different types of evidence possibly make contradicting probabilistic forecasts. A policy maker who receives the forecasts must aggregate them in the best way possible. This is a challenge whenever the policy maker is not familiar with the prior nor the model and evidence available to the experts. We propose a model of non-Bayesian forecast aggregation and adapt the notion of regret as a means for evaluating the policy maker's performance. Whenever experts are Blackwell ordered taking a weighted average of the two forecasts, the weight of which is proportional to its precision (the reciprocal of the variance), is optimal. The resulting minimal regret is equal to $(5^{1.5} - 11)/8 \simeq 0.0225425$, which is 3 to 4 times better than naive approaches such as choosing one expert at random or taking the non-weighted average.

3.2 Aggregating Earnings-Per-Share Predictions

Amir Ban (Tel Aviv University, IL)

Joint work of Amir Ban, Yishay Mansour

We investigate whether stock-market analysts possess differential expertise, and if so, whether such an observation makes it possible to aggregate multiple analyst estimates into a result that is significantly more accurate than their consensus average. We do a retrospective study using historical quarterly earnings-per-share forecasts and actual results for large publicly traded companies, setting a goal of getting the most accurate aggregated forecast possible. Using this goal as criterion, we find that analysts possess both an individual forecast bias as well as individual expertise. The individual bias is significant, while the individual expertise is of lower significance. Together, they enable a $20\% - 30\%$ accuracy improvement over consensus average.

3.3 The optimal strategy for a linear regression game

Peter L. Bartlett (University of California – Berkeley, US)

Joint work of Peter Bartlett, Alan Malek

We consider a linear regression game: at each round, an adversary reveals a covariate vector, the learner predicts a real value, the adversary reveals a label, and the learner incurs the squared prediction error. The aim is to minimize the difference between the cumulative loss and that of the linear predictor that is best in hindsight. We present the minimax optimal strategy for this game, and show that it can be efficiently computed, for natural constraints on the adversarially chosen covariate sequence that prevent the adversary from misrepresenting the scale of the problem. This strategy is horizon-independent, that is, it incurs no more regret than the optimal strategy that knows in advance the number of rounds of the game. We also provide an interpretation of the minimax algorithm as a follow-the-regularized-leader strategy with a data-dependent regularizer.

3.4 Bayesian Methods for Market Clearing

Gianluca Brero (Universität Zürich, CH) and Sébastien Lahaie (Google – New York, US)

We cast the problem of combinatorial auction design in a Bayesian framework in order to incorporate prior information into the auction process and minimize the number of rounds. We develop a generative model of agent valuations and market prices such that clearing prices become maximum a posteriori estimates given observed agent valuations. This generative model then forms the basis of an auction process which alternates between refining estimates of agent valuations and computing candidate clearing prices. We provide an implementation of the auction using assumed density filtering to estimate valuations and expectation maximization to compute prices. An empirical evaluation over a range of valuation domains demonstrates that our Bayesian auction mechanism is very competitive against a conventional combinatorial clock auction, even under the most favorable choices of price increment for this baseline.

3.5 Distributed Methods for Computing Approximate Equilibria

Argyris Deligkas (Technion – Haifa, IL), John Fearnley (University of Liverpool, GB), and Rahul Savani (University of Liverpool, GB)

Joint work of Artur Czumaj, Argyrios Deligkas, Michail Fasoulakis, John Fearnley, Marcin Jurdzinski, Rahul Savani
Main reference Artur Czumaj, Argyrios Deligkas, Michail Fasoulakis, John Fearnley, Marcin Jurdzinski, Rahul Savani: "Distributed Methods for Computing Approximate Equilibria", in Proc. of the Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11–14, 2016, Proceedings, LNCS, Vol. 10123, pp. 15–28, Springer, 2016.
URL http://dx.doi.org/10.1007/978-3-662-54110-4_2

We present a new, distributed method to compute approximate Nash equilibria in bimatrix games. In contrast to previous approaches that analyze the two payoff matrices at the same time (for example, by solving a single LP that combines the two players' payoffs), our algorithm first solves two independent LPs, each of which is derived from one of the two payoff matrices, and then computes an approximate Nash equilibrium using only limited communication between the players. Our method gives improved bounds on the complexity of computing approximate Nash equilibria in a number of different settings. Firstly, it gives a polynomial-time algorithm for computing *approximate well supported Nash equilibria (WSNE)* that always finds a 0.6528-WSNE, beating the previous best guarantee of 0.6608. Secondly, since our algorithm solves the two LPs separately, it can be applied to give an improved bound in the limited communication setting, giving a randomized expected-polynomial-time algorithm that uses poly-logarithmic communication and finds a 0.6528-WSNE, which beats the previous best known guarantee of 0.732. It can also be applied to the case of *approximate Nash equilibria*, where we obtain a randomized expected-polynomial-time algorithm that uses poly-logarithmic communication and always finds a 0.382-approximate Nash equilibrium, which improves the previous best guarantee of 0.438. Finally, the method can also be applied in the query complexity setting to give an algorithm that makes $O(n \log n)$ payoff queries and always finds a 0.6528-WSNE, which improves the previous best known guarantee of 2/3.

3.6 Optimal Auctions through Deep Learning

Paul Dütting (London School of Economics, GB), Zhe Feng, Harikrishna Narasimhan, and David C. Parkes

Main reference Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David C. Parkes: "Optimal Auctions through Deep Learning", CoRR, Vol. abs/1706.03459, 2017.
URL http://arxiv.org/abs/1706.03459

Designing an auction that maximizes expected revenue is an intricate task. Indeed, as of today–despite major efforts and impressive progress over the past few years–only the single-item case is fully understood. In this work, we initiate the exploration of the use of tools from deep learning on this topic. The design objective is revenue optimal, dominant-strategy incentive compatible auctions. We show that multi-layer neural networks can learn almost-optimal auctions for settings for which there are analytical solutions, such as Myerson's auction for a single item, Manelli and Vincent's mechanism for a single bidder with additive preferences over two items, or Yao's auction for two additive bidders with binary support distributions and multiple items, even if no prior knowledge about the form of optimal auctions is encoded in the network and the only feedback during training is revenue and regret. We further show how characterization results, even rather implicit ones such as Rochet's characterization through induced utilities and their gradients, can be leveraged to obtain more precise fits to the optimal design. We conclude by demonstrating the potential of deep learning for deriving optimal auctions with high revenue for poorly understood problems.

3.7 Optimal Auctions for Correlated Bidders with Sampling

Hu Fu (University of British Columbia – Vancouver, CA), Jason D. Hartline (Northwestern University – Evanston, US), Nima Haghpanah, and Robert Kleinberg

Main reference Hu Fu, Nima Haghpanah, Jason D. Hartline, Robert Kleinberg: "Optimal Auctions for Correlated Buyers with Sampling", CoRR, Vol. abs/1406.1571, 2014.
URL http://arxiv.org/abs/1406.1571

Correlation among buyers' valuations enables a revenue maximizing seller to fully extract the social surplus, with no money left to the buyers. This was shown by a classical result by Cremer and McLean. The model has been criticized for allowing arbitrary dependence of the mechanism on the prior: any uncertainty on the prior disrupts the mechanism. We examine this criticism from a learning point of view. We allow uncertainty on the prior but grant the seller sample access from the true prior, and study the number of samples that suffice for surplus extraction. We give precise bounds on the number of samples needed, which show that surplus extraction needs much less information than learning the prior itself. In a sense, this is because the buyers "collaborate" in the learning, driven by their incentives. Our upper bound on the number of samples is by an algebraic argument.

3.8 No Regret and Sequential Prediction

Claudio Gentile (University of Insubria – Varese, IT)

A 90 minute tutorial on classical and recent research in the context of sequential prediction algorithms (expert, bandit, and variants thereof), as well as in online convex optimization and online learning of Lipschitz policies.

3.9 Non-revelation Mechanism Design

Jason D. Hartline (Northwestern University – Evanston, US) and Samuel Taggart

Main reference Jason D. Hartline, Samuel Taggart: "Non-Revelation Mechanism Design", CoRR, Vol. abs/1608.01875, 2016.
URL http://arxiv.org/abs/1608.01875

We consider mechanism design and redesign for markets like Internet advertising where many frequent, small transactions are organized by a principal. Mechanisms for these markets rarely have truthtelling equilibria. We identify a family of winner-pays-bid mechanisms for such markets that exhibit three properties. First, equilibria in these mechanisms are simple. Second, the mechanisms' parameters are easily reoptimized from the bid data that the mechanism generates. Third, the performance of mechanisms in the family is near the optimal performance possible by any mechanism (not necessarily within the family). Our mechanisms are based on batching across multiple iterations of an auction environment, and our approximation bound is asymptotically optimal, with loss inversely proportional to the cube root of the number of iterations batched. Our analysis methods are of broader interest in mechanism design and, for example, we also use them to give new sample complexity bounds for mechanism design in general single-dimensional agent environments.

3.10 Optimizing Worst-case Benchmarks

Aleck Johnsen (Northwestern University – Evanston, US) and Jason D. Hartline (Northwestern University – Evanston, US)

Joint work of Sam Taggart, Aleck Johnsen, Jason D. Hartline

Three common information settings under which algorithms with incomplete information are studied are Bayesian- known distribution, Bayesian- unknown distribution, and worst-case. Whereas Bayesian settings automatically embed optimization problems, worst-case requires benchmarks as a freely chosen parameter, with algorithms subsequently designed to compete pointwise with the benchmark. This talk studies a property of worst-case benchmarks (normalization) that would necessarily allow algorithms approximating the benchmark to extend to give guarantees in the setting of an unknown Bayesian distribution; and another property (resolution) to allow comparison of the efficacy of "normalized" benchmarks, turning benchmark design into an optimization question. Two disparate algorithmic settings of incomplete information are used for the analysis, (online) Expert Learning, and (private values) Auctions.

3.11 Best-Response Dynamics in Combinatorial Auctions with Item Bidding

Thomas Kesselheim (TU Dortmund, DE) and Paul Dütting (London School of Economics, GB)

Main reference Paul Dütting, Thomas Kesselheim: "Best-Response Dynamics in Combinatorial Auctions with Item Bidding", in Proc. of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16–19, pp. 521–533, SIAM, 2017.
URL http://dx.doi.org/10.1137/1.9781611974782.33

In a combinatorial auction with item bidding, agents participate in multiple single-item second-price auctions at once. As some items might be substitutes, agents need to strategize in order to maximize their utilities. A number of results indicate that high welfare can be achieved this way, giving bounds on the welfare at equilibrium. Recently, however, criticism has been raised that equilibria are hard to compute and therefore unlikely to be attained.

In this paper, we take a different perspective. We study simple best-response dynamics. That is, agents are activated one after the other and each activated agent updates his strategy myopically to a best response against the other agents' current strategies. Often these dynamics may take exponentially long before they converge or they may not converge at all. However, as we show, convergence is not even necessary for good welfare guarantees. Given that agents' bid updates are aggressive enough but not too aggressive, the game will remain in states of good welfare after each agent has updated his bid at least once.

In more detail, we show that if agents have fractionally subadditive valuations, natural dynamics reach and remain in a state that provides a $1/3$ approximation to the optimal welfare after each agent has updated his bid at least once. For subadditive valuations, we can guarantee an $\Omega(1/\log m)$ approximation in case of m items that applies after each agent has updated his bid at least once and at any point after that. The latter bound is complemented by a negative result, showing that no kind of best-response dynamics can guarantee more than an $o(\log\log m/\log m)$ fraction of the optimal social welfare.

3.12 Bayesian Methods for Clearing Markets

Sébastien Lahaie (Google – New York, US) and Gianluca Brero (Universität Zürich, CH)

We cast the problem of combinatorial auction design in a Bayesian framework in order to incorporate prior information into the auction process and minimize the number of rounds. We develop a generative model of agent valuations and market prices such that clearing prices become maximum a posteriori estimates given observed agent valuations. This generative model then forms the basis of an auction process which alternates between refining estimates of agent valuations and computing candidate clearing prices. We provide an implementation of the auction using assumed density filtering to estimate valuations and expectation maximization to compute prices. An empirical evaluation over a range of valuation domains demonstrates that our Bayesian auction mechanism is very competitive against a conventional combinatorial clock auction, even under the most favorable choices of price increment for this baseline.

3.13 Submultiplicative Glivenko-Cantelli and Uniform Convergence of Revenues

Yishay Mansour (Tel Aviv University, IL)

Joint work of Noga Alon, Moshe Babaioff, Yannai A. Gonczarowski, Yishay Mansour, Shay Moran, Amir Yehudayoff
Main reference Noga Alon, Moshe Babaioff, Yannai A. Gonczarowski, Yishay Mansour, Shay Moran, Amir Yehudayoff: "Submultiplicative Glivenko-Cantelli and Uniform Convergence of Revenues", CoRR, Vol. abs/1705.08430, 2017.
URL http://arxiv.org/abs/1705.08430

In this work we derive a variant of the classic Glivenko-Cantelli Theorem, which asserts uniform convergence of the empirical Cumulative Distribution Function (CDF) to the CDF of the underlying distribution. Our variant allows for tighter convergence bounds for extreme values of the CDF.

We apply our bound in the context of revenue learning, which is a well-studied problem in economics and algorithmic game theory. We derive sample-complexity bounds on the uniform convergence rate of the empirical revenues to the true revenues, assuming a bound on the k-th moment of the valuations, for any (possibly fractional) $k > 1$.

For uniform convergence in the limit, we give a complete characterization and a zero-one law: if the first moment of the valuations is finite, then uniform convergence almost surely occurs; conversely, if the first moment is infinite, then uniform convergence almost never occurs.

3.14 The Sample Complexity of Single-Parameter Auction Design

Jamie Morgenstern (University of Pennsylvania – Philadelphia, US)

Joint work of Jamie Morgenstern, Tim Roughgarden

This tutorial will overview recent literature on the sample complexity of learning (nearly) revenue-optimal auctions for selling to buyers in the single item (and more generally single parameter) setting.

3.15 Revenue Optimization with Approximate Bid Predictions

Andrés Muñoz Medina (Google – New York, US)

Joint work of Andrés Muñoz Medina, Sergei Vassilvitskii
Main reference Andrés Muñoz Medina, Sergei Vassilvitskii: "Revenue Optimization with Approximate Bid Predictions", CoRR, Vol. abs/1706.04732, 2017.
URL http://arxiv.org/abs/1706.04732

In the context of advertising auctions, finding good reserve prices is a notoriously challenging learning problem. This is due to the heterogeneity of ad opportunity types and the non-convexity of the objective function. In this work, we show how to reduce reserve price optimization to the standard setting of prediction under squared loss, a well understood problem in the learning community. We further bound the gap between the expected bid and revenue in terms of the average loss of the predictor. This is the first result that formally relates the revenue gained to the quality of a standard machine learned model.

3.16 Robust Inference for Games via Theoretical Guarantees

Denis Nekipelov (University of Virginia – Charlottesville, US)

Joint work of Darrell Hoy, Denis Nekipelov, Vasilis Syrgkanis

In the past decade the Economics literature has developed a unified approach to inference and prediction of strategic environments. This approach starts with a full theoretical model that characterizes the preferences of agents and the mechanism of interaction between them.

The Econometrician then infers the components of the theoretical model from the data and provides prediction for the new settings by computing a new equilibrium of the strategic model. The issue with this approach is that the initial step where the Econometrician recovers the preferences of the agents from the data boils down to an inversion of a nonlinear mapping that can be discontinuous and even set-valued. In the talk I will discuss the properties of this mapping for classes of simple games and demonstrate that even in cases where this mapping is invertible, the recovered pre-image is very sensitive to the specification of the model. I will also discuss that such poor behavior of the solution will be preserved even when one uses strong conditions to "regularize" this mapping. All these problems will be further amplified in the prediction.

The approach of set inference provides a robust alternative to traditional inference. In this approach the Econometrician recovers an entire set of preferences of the agents in the interactions that are compatible with many possible specifications of the theoretical model. However, computation of such sets is difficult even in simple games. In my talk I discuss a new approach to inference that is based on the idea of the price of anarchy in Koutsoupias and Papadimitriou (1999). The idea of the approach is to bypass the set inference for the primitives of the model and instead directly infer the outcomes of interest such as welfare or revenue in the game. However, unlike the standard price of anarchy which is based on the "worst case scenario"-based prediction for the outcomes, we propose to consider the bounds that are informed by the distribution of the data. I talk about the new notion of the empirical price of anarchy that yields the price of anarchy over all preferences of agents that could have generated the observable distribution of the data. I then discuss some connections between our notion of the empirical price of anarchy and Economic literature on set inference.

3.17 Weighted Voting Via No-Regret Learning

Ariel D. Procaccia (Carnegie Mellon University – Pittsburgh, US)

Joint work of Nika Haghtalab, Ritesh Noothigattu, Ariel D. Procaccia
Main reference Nika Haghtalab, Ritesh Noothigattu, Ariel D. Procaccia: "Weighted Voting Via No-Regret Learning", CoRR, Vol. abs/1703.04756, 2017.
URL http://arxiv.org/abs/1703.04756

Voting systems typically treat all voters equally. We argue that perhaps they should not: Voters who have supported good choices in the past should be given higher weight than voters who have supported bad ones. To develop a formal framework for desirable weighting schemes, we draw on no-regret learning. Specifically, given a voting rule, we wish to design a weighting scheme such that applying the voting rule, with voters weighted by the scheme, leads to choices that are almost as good as those endorsed by the best voter in hindsight. We derive possibility and impossibility results for the existence of such weighting schemes, depending on whether the voting rule and the weighting scheme are deterministic or randomized, as well as on the social choice axioms satisfied by the voting rule.

3.18 Design of Machine Learning-based Mechanisms

Sven Seuken (Universität Zürich, CH)

Joint work of Gianluca Brero, Benjamin Lubin, Sven Seuken
Main reference Gianluca Brero, Benjamin Lubin, Sven Seuken: "Probably Approximately Efficient Combinatorial Auctions via Machine Learning", in Proc. of the Thirty-First AAAI Conference on Artificial Intelligence, February 4–9, 2017, San Francisco, California, USA., pp. 397–405, AAAI Press, 2017.
URL http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/15040

In this talk, we present a new paradigm we call "designing machine learning-based mechanisms." In contrast to most prior work, our paradigm uses machine learning (ML) directly on the agents' reports, not to optimize some future mechanism, but to immediately make use of the learning outcome (for the current instance). We instantiate this new idea via combinatorial auctions (CAs), and show how using ML inside CAs can substantially simplify the interaction with the bidders. In our CAs, the bidders report their values (bids) to a proxy agent by answering a small number of value queries. The proxy agent then uses an ML algorithm to generalize from those bids to the whole value space, and the efficient allocation is computed based on the generalized valuations. We discuss that this new design leads to new challenges regarding allocative efficiency, individual rationality, and incentives. However, we show that an iterative auction design and an epsilon-expressive ML algorithm address these challenges. To instantiate our design, we use support vector regression (SVR) as the ML algorithm, which enables us to formulate the winner determination problem as a succinct integer program. Finally, we present some experimental results for two stylized spectrum auction domains. Our results demonstrate that even with a small number of bids, our ML-based auctions achieve high allocative efficiency.

3.19 Learning, Optimization, and Noise

Yaron Singer (Harvard University – Cambridge, US)

Joint work of Eric Balkanski, Avinatan Hassidim, Thibaut Horel, Nicole Immorlica, Harikrishna Narasimhan, David, Parkes, Aviad Rubinstein, Jan Vondrak, Yaron Singer

We will discuss a body of work that revolves around a simple question: what can we optimize from data? This question is relevant to several areas of interest of the workshop including mechanism design and auctions, social networks, decision theory, and online learning. In many cases we wish to optimize a function that is not known, but rather learned from data. In combinatorial auctions, for example, an agent may not know her valuation but rather learns it by observing data. In influence maximization in social networks we do not know the influence model we optimize over but rather learn it from data. In a recent line of work we began exploring these questions, as discussed below.

Optimization under Noise

A natural approach to optimization from data is to first learn a surrogate function that approximates the true function generating the data, and then optimize the surrogate. The problem however, is that even when the function is convex or submodular and the learned function well approximates the true function, the optimization problem may be inapproximable, as shown in the papers below. On the positive side, in work currently under submission we show that when noise is stochastic, the optimal guarantees are achievable for canonical submodular optimization problems.

- Optimization from Noisy Preferences with Avinatan Hassidim Working paper
- Optimal Guarantees for Maximizing Noisy Submodular Functions with Avinatan Hassidim In Submission
- Submodular Optimization under Noise with Avinatan Hassidim Conference on Learning Theory (COLT) 2017
- Robust Guarantees for Stochastic Greedy Algorithms with Avinatan Hassidim International Conference on Machine Learning (ICML) 2017
- Maximizing Approximately Submodular Functions with Thibaut Horel Annual Conference on Neural Information Processing Systems (NIPS) 2016
- Information-theoretic Lower Bounds for Convex Optimization with Erroneous Oracles with Jan Vondrak Annual Conference on Neural Information Processing Systems (NIPS) 2015

Optimization from Samples

Recently, my group has been focused on the question of optimization from data: how much training data do we need in order to optimize a function? In a paper published at NIPS 2015 we show that various submodular functions that arise in the context of diffusion in networks are statistically learnable from data. In recent work we show a sharp impossibility result. We show that various important submodular and convex functions that are statistically learnable such as diffusion functions cannot be optimized from sampled data. The moral is that there are models (functions) that may be statistically learnable and amenable to optimization (submodular or convex), though would still require exponentially-many samples for any algorithm to optimize. On the positive side, in work published at NIPS 2016 we give optimal algorithms for functions with bounded curvature and show they can be optimized well from samples.

- Maximizing the Spread of Influence From Training Data with Eric Balkanski and Nicole Immorlica In Submission
- Minimizing a Submodular Function from Samples with Eric Balkanski In Submission
- The Sample Complexity of Optimizing a Convex Function with Eric Balkanski Conference on Learning Theory (COLT) 2017
- The Limitations of Optimization from Samples with Eric Balkanski and Aviad Rubinstein Symposium on Theory of Computation (STOC) 2017
- The Power of Optimization from Samples with Eric Balkanski and Aviad Rubinstein Annual Conference on Neural Information Processing Systems (NIPS) 2016
- Learnability of Influence in Networks with Harikrishna Narasimhan and David Parkes Annual Conference on Neural Information Processing Systems (NIPS) 2015

3.20 Learning from Untrusted Data

Gregory Valiant (Stanford University, US)

We consider the problems of estimation, learning, and optimization over a large dataset of which a subset consists of points drawn from the distribution of interest, and we make no assumptions on the remaining points–they could be well-behaved, extremely biased, or adversarially generated. We investigate this question via two models for studying robust estimation, learning, and optimization. One of these models, which we term the "semi-verified" model, assumes access to a second much smaller (typically constant-sized) set of "verified" or trusted data points that have been drawn from the distribution of interest. The key question in this model is how a tiny, but trusted dataset can allow for the accurate extraction of the information contained in the large, but untrusted dataset. The second model, "list-decodable learning", considers the task of returning a small list of proposed answers. Underlying this model is the question of whether the structure of "good" datapoints can be overwhelmed by the remaining data–surprisingly, the answer is often "no". We present several strong algorithmic results for these models, for a large class of mathematically clean and practically relevant robust estimation and learning tasks.

The talk is based on several joint works with Jacob Steinhardt and Moses Charikar, and with Michela Meister.

3.21 Prediction with a Short Memory

Gregory Valiant (Stanford University, US)

We consider the problem of predicting the next observation given a sequence of past observations. We show that for any distribution over sequences of observations, if the mutual information between past observations and future observations is upper bounded by J, then a simple Markov model over the most recent J/ϵ observations obtains expected KL error ϵ–and hence L_1 error $\sqrt{\epsilon}$–with respect to the optimal predictor that has access to the entire past. For a Hidden Markov Model with n states, J is bounded by $\log n$, a quantity that does not depend on the mixing time. We also establish that this result cannot be improved upon, in the following senses: First, a window length of J/ϵ is information-theoretically necessary for expected KL error epsilon or L_1 error $\sqrt{\epsilon}$. Second, the $d^{(J/\epsilon)}$ samples required to accurately estimate the Markov model when observations are drawn from an alphabet of size d is necessary for any computationally tractable learning/prediction algorithm, assuming the hardness of strongly refuting a certain class of CSPs.

3.22 Sample Complexity of Multi-Item Profit Maximization

Ellen Vitercik (Carnegie Mellon University – Pittsburgh, US)

We study the design of pricing mechanisms and auctions when the mechanism designer does not know the distribution of buyers' values. Instead, the mechanism designer receives a set of samples from this distribution and his goal is to use the sample to design a pricing mechanism or auction with high expected profit. We provide generalization guarantees which bound the difference between average profit on the sample and expected profit over the distribution. These bounds are directly proportional to the intrinsic complexity of the mechanism class the designer is optimizing over. We present a single, general theorem that uses empirical Rademacher complexity to measure the intrinsic complexity of a variety of widely-studied single- and multi-item auction classes, including affine maximizer auctions, mixed-bundling auctions, and second-price item auctions. This theorem also applies to multi- and single-item pricing mechanisms in both multi- and single-unit settings, such as linear and non-linear pricing mechanisms. Despite the applicability of our main theorem, we match or improve over the best-known generalization guarantees for many mechanism classes. Finally, our central theorem allows us to easily derive generalization guarantees for every class in several finely grained hierarchies of auction and pricing mechanism classes. We demonstrate how to determine the precise level in a hierarchy with the optimal tradeoff between profit and generalization using structural profit maximization. The mechanism classes we study are significantly different from well-understood function classes typically found in machine learning, so bounding their complexity requires a sharp understanding of the interplay between mechanism parameters and buyer valuations.

3.23 Learning Mastermind – Some Ideas and Challenges

Bernhard von Stengel (London School of Economics, GB)

In the puzzle game "Mastermind" (available under the name "Guess" at http://hewgill.com/puzzles/), a hidden random code has to be guessed, which is a sequence of 4 pegs, each of which has 6 possible colors. Partial feedback is given on each guess in the form of black pegs (one for each correct colour and place, but never where) and white pegs (one for each other correct colour in the wrong place). The 14 possible feedback combinations (all are possible except 3 black 1 white) partition the remaining codes. Minimizing the maximal partition class allows to find the code with at most 5 guesses (Knuth 1976), and 4.34 expected guesses can be achieved optimally. These may require inconsistent guesses which cannot win immediately.

However, these strategies are not suitable for humans. The main cognitive problem is to anticipate the possible feedback information in order to find a good guess. Instead, a human will use simple rules such as "break symmetry", "use two colors", "stay consistent from the third guess onwards". This puzzle should be a suitable challenge for learning algorithms in order to identify complexity-reducing rules that are obvious to humans, such as symmetry of a guessing strategy in colors and position, and how the rules of the game are defined. Methodologically, one needs to be clear if these rules will be assumed or are to be learned. Another issue, of possible general interest, is the use of introspection (of how a human plays) to validate a learning method.

3.24 Peer Prediction Mechanisms and their Connections to Machine Learning

Jens Witkowski (ETH Zürich, CH)

Joint work of David C. Parkes, Rafael Frongillo, Jens Witkowski
Main reference Rafael M. Frongillo, Jens Witkowski: "A Geometric Perspective on Minimal Peer Prediction", ACM Trans. Economics and Comput., Vol. 5(3), pp. 17:1–17:27, 2017.
URL http://dx.doi.org/10.1145/3070903

Peer prediction mechanisms truthfully elicit private information, such as opinions, experiences, or ratings, from self-interested participants. For example, peer prediction can be used to elicit truthful responses to questions such as "Does this blog contain offensive content?", "Would you recommend this hotel to a friend?", or "Would you consider this article fake news?" Importantly, peer prediction mechanisms elicit truthful responses to these questions without ever observing ground truth at any point, e.g., whether a blog does indeed contain offensive content. In this talk, I will give a brief introduction to the peer prediction problem followed by some of the recent work that is exploring connections between peer prediction mechanisms and machine learning.

3.25 Learning Cooperative Solution Concepts

Yair Zick (National University of Singapore, SG)

Joint work of Maria-Florina Balcan, Ariel D. Procaccia, Yair Zick

Solution concepts in cooperative games have usually been considered in the setting where agent preferences are fully known. In recent work, we consider the setting where agent preferences are unknown; we lay the theoretical foundations for studying the interplay between coalitional stability and (PAC) learning in hedonic games and cooperative games. We introduce the notion of PAC stability – the equivalent of core stability under uncertainty – and examine the PAC stabilizability and learnability of several popular classes of hedonic and cooperative games. We show that all classic cooperative games with transferable utility can be PAC stabilized; however, not all hedonic games can be PAC stabilized.

Participants

- Yakov Babichenko
Technion – Haifa, IL
- Amir Ban
Tel Aviv University, IL
- Siddharth Barman
Indian Institute of Science – Bangalore, IN
- Peter L. Bartlett
University of California – Berkeley, US
- Gianluca Brero
Universität Zürich, CH
- Rachel Cummings
California Institute of Technology – Pasadena, US
- Argyris Deligkas
Technion – Haifa, IL
- Paul Dütting
London School of Economics, GB
- John Fearnley
University of Liverpool, GB
- Felix Fischer
University of Glasgow, GB
- Hu Fu
University of British Columbia – Vancouver, CA
- Claudio Gentile
University of Insubria – Varese, IT
- Paul W. Goldberg
University of Oxford, GB
- Jason D. Hartline
Northwestern University – Evanston, US
- Martin Hoefer
Goethe-Universität – Frankfurt a. M., DE
- Aleck Johnsen
Northwestern University – Evanston, US
- Thomas Kesselheim
TU Dortmund, DE
- Max Klimm
HU Berlin, DE
- Sébastien Lahaie
Google – New York, US
- Yishay Mansour
Tel Aviv University, IL
- Jamie Morgenstern
University of Pennsylvania – Philadelphia, US
- Andrés Muñoz Medina
Google – New York, US
- Paresh Nakhe
Goethe-Universität – Frankfurt a. M., DE
- Denis Nekipelov
University of Virginia – Charlottesville, US
- Ariel D. Procaccia
Carnegie Mellon University – Pittsburgh, US
- Rahul Savani
University of Liverpool, GB
- Sven Seuken
Universität Zürich, CH
- Yaron Singer
Harvard University – Cambridge, US
- Alexander Skopalik
Universität Paderborn, DE
- Gregory Valiant
Stanford University, US
- Ellen Vitercik
Carnegie Mellon University – Pittsburgh, US
- Bernhard von Stengel
London School of Economics, GB
- Jens Witkowski
ETH Zürich, CH
- Yair Zick
National University of Singapore, SG