

Research Statement

Bo Waggoner

Warren Center for Network and Data Sciences, University of Pennsylvania

Fall 2017

Overview

Information permeates our world, and in many forms: data on disk, knowledge in an expert's mind, statistical observations of the natural world. To make decisions or predictions – treatments in health care, climate and public policy, user behavior and suggestions in e-commerce – we seek to understand and design *systems for acquiring, aggregating, and acting on that information*. Such systems can involve computers (e.g. machine-learning algorithms); humans (e.g. voting systems, prediction markets); or a mix (e.g. crowdsourcing platforms).

However, there is a growing recognition of the consequences these systems have for the people they interact with, directly or indirectly. Statistics on health care data can help design lifesaving treatments – but can also have huge privacy impacts on that data's owners. Platforms for prediction or decisionmaking can encounter strategic misreporting in order to manipulate the outcomes. E-commerce systems built on seemingly-neutral principles can quickly become discriminatory if designed without fairness protections in mind.

To overcome these challenges, we must understand how social constraints such as privacy, fairness, and strategic behavior impact these systems; and how these systems impact the people they touch. Then, we can hope to build better systems with societal impact as a primary goal.

My research agenda takes a theoretical approach to this problem both from the top down and bottom up. The latter involves work on subproblems or components of such systems; or on tools in related fields including machine learning, privacy, algorithms, and game theory or “EconCS”.

Meanwhile, the top-down approach involves design or analysis of an entire end-to-end system for acquiring and aggregating information. Some examples are systems that conduct machine-learning or make decisions in provably good ways when individuals hold **data points**. However, the system must be designed with goals such as privacy [1, 17], fairness [5], and resistance to strategic manipulation of the data [7]. Other examples are systems that aggregate **beliefs** or knowledge. A foremost such system in practice is prediction markets; we would like to know how to redesign these systems to prevent manipulation and misleading behavior [4], preserve participants' privacy [17]; or elicit different *kinds* of information [8].

Below, I outline below two of the directions I am most excited about. The first involves a more concrete and technically-focused area; the second touches several fields but centers around a conceptually unifying question. I hope they illustrate what I find compelling about this research agenda: It tends to bring together techniques and concepts from many fields of research – game theory, algorithms, machine learning, privacy, etc. – yet can also revolve around surprisingly technically focused questions. Most important, these technical questions translate into practical consequences.

When working on these problems, I envision a world where everyone is in control of their own data and is fairly compensated for its use; and one where users of online systems can act honestly and simply without being taken advantage of. To achieve this, we need systems that are designed from the ground up to respect principles of privacy, fairness, and incentives. I hope this research can help us move at least a little closer to this goal.

Information Elicitation and Machine Learning

Information elicitation refers to the following setting: Given some “features” x , such as a patient’s medical data, one makes a prediction about some “outcome” y , such as the patient’s blood pressure one year later. The prediction r will be evaluated by a *loss function* $\ell(r, y)$. In microeconomics, such functions (particularly *proper scoring rules*) are used to evaluate expert predictions. Meanwhile, however, this also captures the classic problem in *machine learning*, which is to produce hypotheses h that make predictions $h(x)$ such that, on average, $\ell(h(x), y)$ is small.

The key question of information elicitation, which is fundamental to both the economics and machine-learning settings, is this: **What is the relationship between the loss function chosen and the kinds of predictions collected?** For example, proper scoring rules are known to “elicit” probability distributions over the outcome; the squared loss $\ell(r, y) = (r - y)^2$ is known to elicit the expectation of y ; and absolute loss $\ell(r, y) = |r - y|$ is known to elicit the median.

In Casalaina-Martin et al. [3], Frongillo et al. [9], we investigate a new twist: the loss function $\ell(r, y_1, y_2)$ can use two observations y_1, y_2 (or in general, k observations) to evaluate the prediction. We show that such losses can greatly improve *elicitation complexity* (dimensionality of one’s hypothesis) and *sample complexity* (amount of data required) when attempting to learn higher-order properties of a distribution, such as variances, norms of the distribution, confidence intervals, or so on. In Frongillo and Waggoner [8], we investigate *prediction markets*, which are used to forecast political elections, sporting events, and internally within companies for various purposes. We ask how such markets might look like if a company wants to know, for instance, a 95% percentile prediction of the release date rather than the “expected” release date of a product. Using theory of information elicitation, we are able to characterize the structure of viable prediction markets for a variety of such statistics.

Value of Information in Prediction, Learning, and Decisionmaking

Suppose we take a real-world system capable of learning and optimizing – be that system biological, economic, or computational – and freeze it at a moment in time. The environment may have a multitude of information potentially available and useful. However, collecting it may incur costs: privacy losses, monetary compensation, time, computational power. Meanwhile, each piece of information may have uncertain value. **How does the system determine what information to collect and how to make use of it?**

In Abernethy et al. [1], Zheng et al. [18] we consider this problem from an algorithmic perspective. In (respectively) machine-learning and linear optimization settings, the goal is design systems that utilize limited resources to collect the most valuable information and solve the problem at hand. We prove guarantees on their performance in terms of resources such as a monetary budget or limited number of queries to the environment. The key idea is to relate the cost of available information to its potential value to the system at the current point in time. I believe this principle has potential for much broader interesting study and applications.

In Chen and Waggoner [4], Waggoner et al. [17], we consider generalized forms of prediction markets and ask some of the same questions, but in settings where information consists of beliefs or knowledge in the minds of (strategic) experts. Waggoner et al. [17] uses market-like mechanisms to collect data from individuals for solving a machine-learning task, while preserving the privacy of that data (formalized with *differential privacy*). Chen and Waggoner [4] considers the game-theoretic equilibria of prediction markets: how do strategic agents choose to reveal and aggregate private information? We introduce definitions of *substitutes* and *complements* for pieces of information and give evidence that these are natural and fundamental definitions. Then, we show that they characterize “good” and “bad” equilibria respectively, in the sense of information being aggregated quickly versus slowly or not at all. These definitions build on classic information theory (Shannon 1948) and value of information (Howard 1966) literature, but also have connections to such modern problems as submodular optimization. Because the definitions apply to agents in arbitrary decision problems, they connect closely to the theory of information elicitation discussed in the previous section.

Other Work and Future Directions

Other work. Two general areas of interest to me, not yet discussed, are: online learning and randomized or online algorithms [1, 10, 13, 14, 18]; and particularly “EconCS” and algorithmic game theory, which I define as “the (algorithmic) study of systems of goal-directed agents.” I have interests in auctions and “mechanism design” [2, 11], social choice, voting, and fairness [5, 16], and crowdsourcing/“peer prediction” [15].

Future directions: information elicitation. I am excited about information elicitation because of its mathematical depth and elegance; fundamental importance in both decision theory and machine learning; and promise of future applicability. One long-term program for the area is to evaluate the impact of different losses that elicit the same statistic (such as the choice of Bregman divergence when predicting the mean), particularly when one has a constrained hypothesis class – here the role of the loss is to trade off mistakes on some data points versus others. Another problem I expect to grow in importance is the (automated) choice of good *surrogate* loss functions. What makes a good loss function for a given problem (e.g. features such as convexity, dimensionality of the hypothesis, and so on) and how can we find good surrogates when the loss or statistic we care about is computationally intractable?

Future directions: value of information. This is a direction I am particularly excited about for the future as there are many possible environments and models in which to formalize this

problem, and many real-world “big data” settings where there is a new and genuine need to solve such problems. I hope to establish simple-yet-general models that do not just solve one-off problems, but encourage series of works that improve bounds or extend the models. I believe the online learning model of [1, 4] are steps in this direction.

I have a number of open problems regarding substitutes and complements of information and hope to apply these definitions to better understand a variety of problems, such as mechanisms for buying and selling information. Some problems are computational, e.g. the complexity of optimally releasing a subset of information or of finding optimal strategies in a prediction market; while others are game-theoretic, including extensions of our equilibrium results to financial markets models of interest to economists. Finally, some questions are what I call “structural”: understanding the nature of substitutes and complements themselves, for instance, what signals will be substitutes for an agent who may face one of several possible decision problems? Can we define a hierarchy or partial ordering of signal classes by “substitutability”, such that in any scenario that one set of signals are substitutes, a set below it is sure to be substitutes as well?

Future directions: learning, fairness, and privacy. There are many exciting research directions involving the interplay between machine learning with the societal goals of fairness or privacy for holders of data. One direction for fairness involves learning systems with partial feedback. Imagine an automated banking system that gives loans based on people’s financial information. However, it only observes and “learns” from data about the people it gave loans to; so if it develops an unfair bias or discriminatory aspect, the system may never learn to correct this. How can we ensure fairness in such settings?

For privacy, an important question touching theory and practice is *telemetry*, the practice of software systems such as operating systems and web browsers collecting information on users in order to improve performance or detect bugs. Although differential privacy has begun to be deployed for this problem, we still have much to do to ensure users’ privacy because the information collection spans years of a users’ life.

These kinds of questions and many more arise naturally observing how machine learning is applied in practice today in environments with people, who act strategically and have concerns for fairness and privacy. Those environments motivate theoretically deep and interesting questions with a diverse and exciting set of future directions.

References

- [1] Jacob D. Abernethy, Yiling Chen, Chien-Ju Ho, and Bo Waggoner. Low-cost learning via active data procurement. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 619–636, 2015. doi: 10.1145/2764468.2764519.
- [2] Yang Cai, Mohammad Mahdian, Aranyak Mehta, and Bo Waggoner. Designing markets for daily deals. In *Ninth International Conference on Web and Internet Economics*, WINE '13, pages 82–95, 2013. doi: 10.1007/978-3-642-45046-4_8.
- [3] Sebastian Casalaina-Martin, Rafael M. Frongillo, Tom Morgan, and Bo Waggoner. Multi-observation elicitation. In *Proceedings of the 30th Conference on Learning Theory*, COLT '17, pages 449–464, 2017.
- [4] Yiling Chen and Bo Waggoner. Informational substitutes. In *IEEE 57th Annual Symposium on Foundations of Computer Science*, FOCS '16, pages 239–247, 2016. doi: 10.1109/FOCS.2016.33.
- [5] Yiling Chen, Kobbi Nissim, and Bo Waggoner. Fair information sharing for treasure hunting. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI '15, pages 851–857, 2015.
- [6] Yuan Deng, Debmalya Panigrahi, and Bo Waggoner. The complexity of stable matchings under substitutable preferences. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, AAAI '17, pages 480–486, 2017.
- [7] Jinshuo Dong, Aaron Roth, Zachary Schutzman, Bo Waggoner, and Z. Steven Wu. Strategic classification from revealed preferences. 2017.
- [8] Rafael Frongillo and Bo Waggoner. An axiomatic study of scoring rule markets. In *Proceedings of the Ninth Innovations in Theoretical Computer Science Conference*, ITCS '18, 2018. URL <https://arxiv.org/abs/1709.10065>.
- [9] Rafael Frongillo, Nishant Mehta, Tom Morgan, and Bo Waggoner. Multi-observation regression. 2017.
- [10] Sampath Kannan, Jamie Morgenstern, Aaron Roth, Bo Waggoner, and Z. Steven Wu. A smoothed analysis of the greedy algorithm for the linear contextual bandit problem. 2017.
- [11] Robert D. Kleinberg, Bo Waggoner, and E. Glen Weyl. Descending price optimally coordinates search. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 23–24, 2016. doi: 10.1145/2940716.2940760.
- [12] Katrina Ligett, Seth Neel, Aaron Roth, Bo Waggoner, and Z. Steven Wu. Accuracy first: Selecting a differential privacy level for accuracy-constrained ERM. In *Advances in Neural Information Processing Systems 30*, NIPS '17, 2017.
- [13] Aranyak Mehta, Bo Waggoner, and Morteza Zadimoghaddam. Online stochastic matching with unequal probabilities. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pages 1388–1404, 2015. doi: 10.1137/1.9781611973730.92.

- [14] Bo Waggoner. ℓ_p testing and learning of discrete distributions. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, ITCS '15, pages 347–356, 2015. doi: 10.1145/2688073.2688095.
- [15] Bo Waggoner and Yiling Chen. Output agreement mechanisms and common knowledge. In *Proceedings of the Second AAI Conference on Human Computation and Crowdsourcing*, HCOMP '14, 2014.
- [16] Bo Waggoner, Lirong Xia, and Vincent Conitzer. Evaluating resistance to false-name manipulations in elections. In *Proceedings of the Twenty-Sixth AAI Conference on Artificial Intelligence*, AAAI '12, 2012.
- [17] Bo Waggoner, Rafael M. Frongillo, and Jacob D. Abernethy. A market framework for eliciting private data. In *Advances in Neural Information Processing Systems 28*, NIPS '15, pages 3510–3518, 2015.
- [18] Shuran Zheng, Bo Waggoner, Yang Liu, and Yiling Chen. Active information acquisition for linear optimization. 2017. URL <https://arxiv.org/abs/1709.10061>.