EECS 547: Electronic Commerce (Algorithmic Game Theory)

(Tentative Syllabus)

Lecture: Tuesday and Thursday 10:30am-Noon in 1690 BBB
Section: Friday 3pm-4pm in EECS 3427
Email: AGT-staff@umich.edu

Instructor: Grant Schoenebeck
Office Hours: Wednesday 3pm and by appointment; 3636 BBBB

Graduate Student Instructor: Biaoshuai Tao
Office Hours: 4-5pm Friday in Learning Center (between BBB and Dow in basement)

Introduction:

As the internet draws together strategic actors, it becomes increasingly important to understand how strategic agents interact with computational artifacts. Algorithmic game theory studies both how to models strategic agents (game theory) and how to design systems that take agent strategies into account (mechanism design). On-line auction sites (including search word auctions), cyber currencies, and crowd-sourcing are all obvious applications. However, even more exist such as designing non-monetary award systems (reputation systems, badges, leader-boards, Facebook "likes"), recommendation systems, and data-analysis when strategic agents are involved.

In this particular course, we will especially focus on information elicitation mechanisms (crowd-sourcing).

Description:

Modeling and analysis for strategic decision environments, combining computational and economic perspectives. Essential elements of game theory, including solution concepts and equilibrium computation. Design of mechanisms for resource allocation and social choice, for example as motivated by problems in electronic commerce and social computing.

Format:

The two weeks will be background lectures.

The next half of the class will be seminar style and will focus on crowd-sourcing, and especially information elicitation mechanisms. Each class period, the class will read a paper (or set of papers), and a pair of students will briefly present the papers and lead a class discussion. Students are expected to contribute to the in-class discussion, present one or more papers, and to execute a final paper. A brief work sheet/response paper about the day's reading will be required before each lecture to help facilitate discussion. The seminar style portion of this course will focus on crowd-sourcing, and

especially information elicitation mechanisms. Students will execute a final project which contains some portion of original research.

The last portion of the course will be lectures on classical algorithmic game theory results. To the extent that these are necessary for understanding the papers we read, they will be interspersed with the seminar style classes.

Goals:

This course has several goals. First, students should gain a basic grasp of classical algorithmic game theory results. Second, students acquire an understanding of information elicitation mechanisms that enables them to read and understand cutting edge research. Third, students should learn how to read a research paper and analyze its strengths and weaknesses. They should see how research progresses both by reading and by producing original research.

Prerequisites:

No previous knowledge of game theory or algorithms will be assumed, but familiarity with mathematical reasoning and basic probability theory will be essential to getting the most out of this course. Students should expect to learn additional mathematics on their own as necessary. Parts of this course will rely on basic ideas from economics (like Nash equilibrium and related notions), so experience with these ideas will also be helpful, but not assumed. I encourage students from a variety of diverse disciplines to consider this course. If you are interested in the course but aren't sure if you have the necessary training, please contact me and ask!

Additionally, you can try to read some of the technical papers listed below to get a sense for what we will be doing. However, note that the first part of the course will survey some background material designed to help you read these papers.

Grade Components:

Course Leading:

Students will be put in pairs and assigned a day to present. The presenters should prepares a worksheet for the other students to fill out to help them understand important themes of the paper. The leaders should prepare to lead a section of class that includes about 1/3 engaging activing, 1/3 lecture, and 1/3 discussion. This presentation should highlight the main contributions of the works and also provide some context.

Final Project:

A main component of this course will be a final project. The goal of this project is to gain a deep understanding of some specific topic related to algorithmic game theory, and to use that understanding to work on an open research problem. Students should work in small groups of 2-3 in order to complete this project. Members of the group should equally contribute to the project. The final project must be approved, but, in general, students are free to pick any topic related to algorithmic game theory.

We will have a poster sessions where final projects are presented. Additionally, students should turn in a write up of their final projects due **December 12**. The project should be written up in 10 pages (max) but may include appendices for completeness. A project proposal is due **November 9** in class (2 pages max), but students are encouraged to get an earlier start.

Problem Sets:

There will be problem sets covering both the initial lecture material, and the lectures at the end of class (none will be given on the information elicitation readings). Each will be due the evening before a class. Late assignments will receive a 10% grade deduction per day up to 3 days late after which they will not be accepted.

Student may collaborate, but should state with whom they collaborated.

Final Exam:

A final exam will be administered 10:30am-12:30pm on Tuesday, December 19th in BBB 1690 (our usual room). The exam will be closed book, but students will be allowed one sheet of paper (8.5" x 11"), front and back, for any notes. While the focus of the class will be on projects, this exam will ensure that students keep up with the boarder material of the course. A practice exam will be available help students study (and to indicate what the exam will focus on). A review session will be held in BBB 3901 10am-11am, Friday, December 15.

Reading Responses:

You are required to read papers and other listed reading materials before each class. (Materials listed under additional readings in the schedule are optional.) You must upload reading worksheet responses by midnight before class. Your comments should include answers to posted reading questions (if any) and general comments. For research papers, things to think about for general comments include (you don't need to hit all of these...):

- What is the main contribution of this paper? Why is this interesting?
- What was the main insight?
- What are you still confused about?
- Is anything missing?
- Where does this paper fit?
- What are conceptual contributions?
- What are technical contributions?
- What are technical hurdles?
- What are strengths of the paper?
- What are weaknesses?
- What are ideas for future work?
- What is a good idea for a final project relating to this paper?

Professor Michael Mitzenmacher has a blog post on a how to read a paper that you may find helpful if you do not have much experience reading research papers
http://mybiasedcoin.blogspot.com/2010/03/reading-research-paper.html

You response will be marked highly if it clear that you have read and thought about the paper. It is okay to have typos as long as your response is easily understood. We will drop the lowest two marks of the term.

In-class Participation:

Students are expected to actively contribute comments to class discussion. Quality is much more important than quantity. Both thoughtful questions and thoughtful answers are valued.

Grading:

Participation and Comments	15%
Problem Sets	15%
In Class Presentation(s)	20%
Final Project	35%
Final	15%

Rubrics will be posted for different assignments. Scores will be aggregated according to the above formula. Students will be ordered according to their final grades. Grade will be "curved" and assigned subjectively so that:

A+	Demonstrates a strong and in-depth understanding of the material. Will be able to see new applications of tools in the future and apply the material to it. Can expertly combine tools in innovative ways.
A	Demonstrated a strong and in-depth understanding of the material. Will be able to see new applications of tools in the future and apply the material to it.
A -	Demonstrated a solid understanding of the material. Will be able to see some new applications of tools in the future and apply the material to it.
B+	Demonstrated a good grasp of most of the material. Will be able to apply most tools in the future, but typically, the applications must be pointed out.
B	Demonstrated a good grasp of some of the material and a fair understanding of most. Will be able to apply some tools in the future, but typically, the applications must be pointed out.
B-	Demonstrated a fair understanding of most of the material. Can only apply tools in very straightforward manners even after applications are pointed out.

< C+	Failed to demonstrate understanding of most material. Failed to demonstrate the ability to even straight-forwardly apply tools.

Tentative Schedule

Readings will be posted in the class nb site; email the course staff if you do not have access!

Background

September 5: Overview, Class Policies, Game

September 7: Intro to Game Theory I

- Algorithmic Game Theory - Nisan: Chapter 1 Sections 1.1-1.4.

September 12: Intro to Game Theory II

- Algorithmic Game Theory- Nisan: Chapter 1 Sections 1.4-1.8.

September 12: Course Background

- Background Handout

Information Elicitation

September 19: 3 Foundational Results — Presented by Zach and Linh

- Miller, Resnick, Zechhauser. "Eliciting informative feedback: The peer-prediction method." Management Science, 2005. http://pubsonline.informs.org/doi/pdf/10.1287/mnsc.1050.0379
- Prelec. "Baysian Truth Serum for subjective data." https://www.researchgate.net/profile/Drazen_Prelec/publication/8231017_A_Bayesian_Truth_Serum_for_Subjective_Data/links/0deec51bb0d8741001000000/A-Bayesian-Truth-Serum-for-Subjective-Data.pdf
- Dasgupta, Ghosh. "Crowdsourced judgement elicitation with endogenous proficiency." https://arxiv.org/pdf/1303.0799.pdf

September 21: Information Theory Interpretation — Presented by Innocent

- Kong, Schoenebeck. "A Framework For Designing Information Elicitation Mechanisms That Reward Truth-telling." https://arxiv.org/abs/1605.01021 (new version available soon!)

September 26: Review — Presented by Professor

September 28: Cheap Signals — Presented by Zhiya and Chris G.

- Gao, Wright, Leyton-Brown. "Incentivizing Evaluation via Limited Access to Ground Truth: Peer-Prediction Makes Things Worse." https://arxiv.org/abs/1606.07042
- Kong, Schoenebeck. "Eliciting Expertise without Verification." Available soon.

October 3: Pricing (and Privacy) — Presented by Zach and Kristopher S.

- Arpita Ghosh, Katrina Ligett, Aaron Roth, and Grant Schoenebeck. “Buying Private Data without Verification“ EC ‘14. https://arxiv.org/abs/1404.6003
- [Acemoglu, Daron, Mohamed Mostagir, and Asuman Ozdaglar. Managing innovation in a crowd. No. w19852. National Bureau of Economic Research, 2014. https://dspace.mit.edu/bitstream/handle/1721.1/84477/Acemoglu14-04.pdf?sequence=1]

October 5: Aggregating Information Presented by Hunter and Innocent

- Dražen Prelec, H. Sebastian Seung, John McCoy. “A solution to the single-question crowd wisdom problem.” Nature 2017. http://www.nature.com/nature/journal/v541/n7638/full/nature21054.html?foxtrotcallback=true
- [Arpita Ghosh, Satyen Kale, and Preston McAfee. “Who moderates the moderators?: Crowdsourcing abuse detection in user-generated content.” EC ’11. http://vita.mcafee.cc/PDF/UGC2.pdf]

October 10: Peer Grading Presented by Ziwei and Chen

- James R. Wright, Chris Thornton, Kevin Leyton-Brown. “Mechanical TA: Partially Automated High-Stakes Peer Grading.” SICCSE ’15. http://www.cs.ubc.ca/~kevinlb/papers/2015-SIGCSE-MechTA.pdf
- de Alfaro, Luca, Michael Shavlovsky, and Vassilis Polychronopoulos. "Incentives for Truthful Peer Grading." arXiv preprint arXiv:1604.03178 (2016). https://arxiv.org/pdf/1604.03178.pdf

October 12: Eliciting Functions Presented by Hunter and Alicia

- Salganik, Matthew J., and Karen EC Levy. "Wiki surveys: Open and quantifiable social data collection." *PloS one* 10, no. 5 (2015): e0123483. http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0123483
- [Khatib, Firas, Seth Cooper, Michael D. Tyka, Kefan Xu, Ilya Makedon, Zoran Popović, and David Baker. "Algorithm discovery by protein folding game players." PNAS 2011. http://www.pnas.org/content/108/47/18949.full.pdf]

October 17: <Fall break>

October 19: Information Over Time Presented by Linh and Aditya

- Aaronson, Scott. "The complexity of agreement." STOC ’05. https://www.scottaaronson.com/papers/agree-econ.pdf
- Ely, Jeffrey, Alexander Frankel, and Emir Kamenica. "Suspense and surprise." Journal of Political Economy 123, no. 1 (2015): 215-260. http://fbe.usc.edu/seminars/papers/AE_3-8-13_KAMENICAsuspense.pdf

October 24: Predict Markets Presented by Dong and Aditya

- Hanson, R., 2012. Logarithmic markets coring rules for modular combinatorial information aggregation. *The Journal of Prediction Markets*, *1*(1), pp.3-15. http://www.ubplj.org/index.php/jpm/article/download/417/448
- Hanson, R., 2006. Decision markets for policy advice. *Promoting the general welfare: American democracy and the political economy of government performance*, pp.151-173. https://www.researchgate.net/profile/Robin_Hanson2/publication/46454134_Decision_Markets_for_Policy_Advice/links/00b4951facaf4d176d000000.pdf

October 26: Prediction markets and Outcome manipulation Presented by Nate and Alicia

- Freeman, Rupert, Sébastien Lahaie, and David M. Pennock. "Crowdsourced Outcome Determination in Prediction Markets." In AAAI, pp. 523-529. 2017. http://www.aaai.org/ocs/index.php/AAAI/AAAI17/paper/download/14675/13807
- Chen, Yiling, Xi Alice Gao, Rick Goldstein, and Ian A. Kash. "Market Manipulation with Outside Incentives." In AAAI. 2011. http://www.aaai.org/ocs/index.php/AAAI/AAAI11/paper/download/3747/3948/

October 31: Empirical Presented by Ziwei and Chen

- Xi Alice Gao, Andrew Mao, Yiling Chen, and Ryan Prescott Adams. Trick or treat: putting peer prediction to the test. In Proceedings of the 15th ACM Conference on Economics and Computation (EC 2014), pages 507–524. ACM, 2014. http://scholar.harvard.edu/files/hips/files/gao-trick-ec-2014.pdf
- Dreber, A., Pfeiffer, T., Almenberg, J., Isaksson, S., Wilson, B., Chen, Y., Nosek, B.A. and Johannesson, M., 2015. Using prediction markets to estimate the reproducibility of scientific research. *Proceedings of the National Academy of Sciences*, *112*(50), pp.15343-15347. http://www.pnas.org/content/112/50/15343.full.pdf

November 2: Agent-based modeling Presented by Kristopher S. and Zhiya

- Shnayder, Victor, Rafael Frongillo, and David C. Parkes. "Measuring performance of peer prediction mechanisms using replicator dynamics." (2016). November 7: Assigning Players Questions (little pricing) [look at these] / Contest Systems
- [Gilbert, N., 2008. *Agent-based models* (No. 153). Sage. http://www2.econ.iastate.edu/tesfatsi/AgentBasedModels.Chapter1.NGilbert2008.pdf]

November 7: Behavioral Presented by Nate and Dong

- Wright, J.R. and Leyton-Brown, K., 2010, July. Beyond Equilibrium: Predicting Human Behavior in Normal-Form Games. In *AAAI*. http://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/download/1946/2116
- [Rothschild, D.M. and Wolfers, J., 2011. Forecasting elections: Voter intentions versus expectations. https://www.brookings.edu/wp-content/uploads/2016/06/01-voter-expectations-wolfers.pdf]

Classical Results

November 9: Reputation systems; Chapter 27 of AGT

November 14: Reputation systems/social choice; AGT text Chapter 9.1-9.4;

November 16: social choice/Vickrey-Clarke-Groves; AGT text Chapter 9.1-9.4; 10.1-10.2;

November 21: Revenue Optimal Auctions I: Hartline 2.4-3.3

November 23: <Thanks giving>

November 28: Revenue Optimal Auctions II / All Pay Auctions

November 30: Bitcoin I

December 5: Bitcoin II

December 7: Big Data

December 12: Fairness in Machine Learning

Other

Accommodation for students with disabilities

If you think you need an accommodation for a disability, please let me know at your earliest convenience. Some aspects of this course, the assignments, the in-class activities, and the way we teach may be modified to facilitate your participation and progress. As soon as you make me aware of your needs, we can work with the Office of Services for Students with Disabilities (SSD) to help us determine appropriate accommodations. SSD (734-763-3000; **http://www.umich.edu/~sswd/**) typically recommends accommodations through a Verified Individualized Services and Accommodations (VISA) form. I will treat any information you provide as private and confidential.

Academic integrity

All submitted work must be your own, original work unless you clearly mark it as being otherwise. If you are directly quoting, or building on others' writing, provide a citation. See **the Rackham Graduate policy on Academic and Professional Integrity** for the definition of plagiarism, and associated consequences.

Laptops

Students are welcome to use Laptops in class for the purposes of the class: taking notes, referring to papers, referring to student responses. Laptops should not be used for email, Facebook, Youtube etc. This will affect class discussion and distract your neighbor.

Additional Readings:

Peer prediction Like Mechanisms:

Boi Faltings, Radu Jurca, Pearl Pu, and Bao Duy Tran. Incentives to counter bias in human computation. In Second AAAI Conference on Human Computation and Crowdsourcing, 2014

Sharad Goel, Daniel M. Reeves, and David M. Pennock. Collective revelation: A mechanism for self-verified, weighted, and truthful predictions. In Proceedings of the 10th ACM Conference on Electronic Commerce (EC 2009), 2009.

Radu Jurca and Boi Faltings. Enforcing truthful strategies in incentive compatible reputation mechanisms. In Internet and Network Economics, pages 268–277. Springer, 2005.

Radu Jurca and Boi Faltings. Collusion-resistant, incentive-compatible feedback payments. In Proceedings of the 8th ACM Conference on Electronic Commerce (EC 2007), pages 200–209. . ACM, 2007.

Radu Jurca and Boi Faltings. Incentives for expressing opinions in online polls. In Proceedings of the 9th ACM Conference on Electronic Commerce (EC 2008), 2008.

Radu Jurca and Boi Faltings. Mechanisms for making crowds truthful. J. Artif. Int. Res., 34(1), March 2009.

R. Jurca and B. Faltings. Incentives for answering hypothetical questions. In Proceedings of the 1st Workshop on Social Computing and User Generated Content (SC 2011). ACM, 2011.

Ece Kamar and Eric Horvitz. Incentives for truthful reporting in crowdsourcing. In Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 3, pages 1329–1330. International Foundation for Autonomous Agents and Multiagent Systems, 2012.

N. Lambert and Y. Shoham. Truthful surveys. Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE 2008), 2008.

Goran Radanovic and Boi Faltings. A robust bayesian truth serum for non-binary signals. In Proceedings of the 27th AAAI Conference on Artificial Intelligence, AAAI 2013, number EPFL-CONF-197486, pages 833–839, 2013.

Goran Radanovic and Boi Faltings. Incentives for truthful information elicitation of continuous signals. In Twenty-Eighth AAAI Conference on Artificial Intelligence, 2014.

Heterogeneous Users

Agarwal, Mandal, Parkes. "Peer Prediction with Heterogeneous Users."
http://www.cs.toronto.edu/~nisarg/papers/het-pp.pdf

Focal mechanisms

Radu Jurca and Boi Faltings. Minimum payments that reward honest reputation feedback. In Proceedings of the 8th ACM Conference on Electronic Commerce (EC 2006), 2006.

Yuqing Kong, Katrina Liggett, and Grant Schoenebeck. Putting peer prediction under the micro(economic)scope and making truth-telling focal. 2016.

Yuqing Kong and Grant Schoenebeck. Equilibrium selection in information elicitation without verification via information monotonicity. 2016.

Machine Learning variation:

Yang Cai, Constantinos Daskalakis, and Christos H Papadimitriou. Optimum statistical estimation with strategic data sources. arXiv preprint arXiv:1408.2539, 2014.

Pricing and Task Assignment:

Yang Liu and Mingyan Liu. An online learning approach to improving the quality of crowdsourcing. In ACM SIGMETRICS, 2015.

Singer, Y. and Mittal, M., 2013, May. Pricing mechanisms for crowdsourcing markets. In Proceedings of the 22nd international conference on World Wide Web (pp. 1157-1166). ACM.

Gagan Goel, Afshin Nikzad, and Adish Singla. Mechanism design for crowdsourcing markets with heterogeneous tasks. In Second AAAI Conference on Human Computation and Crowdsourcing, 2014.

Eliciting Private Signals:

Arpita Ghosh and Aaron Roth. Selling privacy at auction. In Proceedings of the 12th ACM conference on Electronic commerce, pages 199–208, 2010.

Yiling Chen, Stephen Chong, Ian A Kash, Tal Moran, and Salil Vadhan. Truthful mechanisms for agents that value privacy. In Proceedings of the 14th ACM Conference on Electronic Commerce (EC 2013), pages 215–232.

Arpita Ghosh, Katrina Ligett, Aaron Roth, and Grant Schoenebeck. Buying private data without verification. In Proceedings of the 15th ACM Conference on Economics and Computation (EC 2014), pages 931–948, 2014.. ACM, 2013.

Katrina Ligett and Aaron Roth. Take it or leave it: Running a survey when privacy comes at a cost. In PaulW. Goldberg, editor, Internet and Network Economics, volume 7695 of Lecture Notes in Computer Science, pages 378–391. Springer Berlin Heidelberg, 2012.

Aaron Roth, The Sensitive Surveyor Problem.

Aggregation:

Arpita Ghosh, Satyen Kale, and Preston McAfee. Who moderates the moderators?: Crowdsourcing abuse detection in user-generated content. In Proceedings of the 12th ACM Conference on Electronic Commerce, EC '11, pages 167–176, New York, NY, USA, 2011. ACM.

Nilesh Dalvi, Anirban Dasgupta, Ravi Kumar, and Vibhor Rastogi. Aggregating crowdsourced binary ratings. In Proceedings of the 22Nd International Conference on World Wide Web, WWW '13, pages 285–294, 2013.

David R Karger, Sewoong Oh, and Devavrat Shah. Iterative learning for reliable crowdsourcing systems. In Advances in neural information processing systems, pages 1953–1961, 2011.

Lots of interesting literature here (too many to list).

Peer Grading

More Complex Functions

Cheng, Justin, and Michael S. Bernstein. "Flock: Hybrid crowd-machine learning classifiers." In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pp. 600-611. ACM, 2015.
https://pdfs.semanticscholar.org/f1a2/e20e5ed83128511020417a4284e4c475cfef.pdf

Jacob Abernethy, Rafael Frongillo. "A Collaborative Mechanism for Crowdsourcing Prediction Problems." NIPS '11. http://papers.nips.cc/paper/4382-a-collaborative-mechanism-for-crowdsourcing-prediction-problems.pdf

Incentives for User Generated Content:

David Easley and Arpita Ghosh. Incentives, gamification, and game theory: an economic approach to badge design. In Proceedings of the fourteenth ACM conference on Electronic commerce, pages 359–376. ACM, 2013.

Arpita Ghosh and Preston McAfee. Incentivizing high-quality user-generated content. In Proceedings of the 20th international conference on World wide web, pages 137–146. ACM, 2011.

Jain, S., Chen, Y. and Parkes, D.C., 2009, July. Designing incentives for online question and answer forums. In *Proceedings of the 10th ACM conference on Electronic commerce* (pp. 129-138). ACM.

Cavallo, R. and Jain, S., 2012, June. Efficient crowdsourcing contests. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2* (pp. 677-686). International Foundation for Autonomous Agents and Multiagent Systems.

Jain, S. and Parkes, D.C., 2013. A game-theoretic analysis of the ESP game. *ACM Transactions on Economics and Computation*, *1*(1), p.3.

Prediction Markets

Abernethy, J., Kutty, S., Lahaie, S. and Sami, R., 2014, June. Information aggregation in exponential family markets. In *Proceedings of the fifteenth ACM conference on Economics and computation* (pp. 395-412). ACM.

Empirical Testing of Information Elicitation Mechanisms/:

Leslie K John, George Loewenstein, and Drazen Prelec. Measuring the prevalence of questionable research practices with incentives for truth telling. Psychological science, 2012.

Servan-Schreiber, Emile, Justin Wolfers, David M. Pennock, and Brian Galebach. "Prediction markets: Does money matter?." Electronic markets 14, no. 3 (2004): 243-251. https://pdfs.semanticscholar.org/49e8/8ce5f20955cdf5b73911e467026c119fbe29.pdf

Adam Marcus, Eugene Wu, David R Karger, Samuel Madden, and Robert C Miller. Crowdsourced databases: Query processing with people. CIDR, 2011.

Many more can be found.

Behavioral Economics:

https://www.brookings.edu/wp-content/uploads/2016/06/01-voter-expectations-wolfers.pdf

Scoring Rules:

Jacob D Abernethy and Rafael M Frongillo. A characterization of scoring rules for linear properties. In COLT, pages 27–1, 2012.

A Philip Dawid. The geometry of proper scoring rules. Annals of the Institute of Statistical Mathematics, 59(1):77–93, 2007.