

Yu Cheng

CIT 413, 115 Waterman St, Providence, RI 02906

✉ yu_cheng@brown.edu

🌐 <https://cs.brown.edu/people/ycheng79/>

Research Interests

Theoretical Machine Learning: provable and scalable learning algorithms, theoretical deep learning, high-dimensional robust statistics, learning in the presence of strategic behavior.

Optimization: non-convex optimization, convex optimization, spectral graph theory.

Game Theory: information structure design, computational social choice, equilibrium computation.

Employment

2022–present **Assistant Professor**, *Department of Computer Science, Brown University.*

2019–2022 **Assistant Professor**, *Department of Mathematics, University of Illinois at Chicago.*

Adjunct Assistant Professor, *Department of Computer Science, University of Illinois at Chicago.*

Fall 2019 **Visiting Scholar**, *School of Mathematics, Institute for Advanced Study (Princeton, NJ).*

Member of the Special Year on Optimization, Statistics, and Theoretical Machine Learning.

2017–2019 **Postdoctoral Researcher**, *Department of Computer Science, Duke University.*

Hosted by Vincent Conitzer, Rong Ge, Kamesh Munagala, and Debmalya Panigrahi.

Education

2011–2017 **Ph.D. in Computer Science**, *University of Southern California.*

Thesis title: Computational Aspects of Optimal Information Revelation. Advisor: Shang-Hua Teng.

Thesis committee: Shaddin Dughmi, David Kempe, Yan Liu, and Ben Reichardt.

2007–2011 **B.S. in Computer Science** (ACM Honors Class), *Shanghai Jiao Tong University, China.*

Honors and Awards

2018 **Best Paper Award**, WINE 2018 (the 14th Conference on Web and Internet Economics).

2017 **Finalist** (5 finalists selected), Excellence in Graduate Research Award in Machine Learning, University of Southern California.

2013 **14th place**, ACM-ICPC (International Collegiate Programming Contest) World Finals.

2011, 2012 **Champion**, ACM-ICPC Southern California Regional Contest.

2010 **Student coach** of ACM-ICPC **World Champion** team (invited to ACM Awards Banquet).

Research Grants

Jan. 2022– NSF Award CCF-2307106. *AF: Small: Faster Algorithms for High-Dimensional Robust*

Dec. 2024 *Statistics*. PI: Yu Cheng. Award Amount: \$390,971.

Publications

Conference Papers

[27] Yu Cheng, Ilias Diakonikolas, Rong Ge, Shivam Gupta, Daniel M. Kane, Mahdi Soltanolkotabi. Outlier-Robust Sparse Estimation via Non-Convex Optimization. *In Proceedings of the 36th Conference on*

- Neural Information Processing Systems (NeurIPS)*, 2022. To appear.
- [26] Hanrui Zhang, Yu Cheng, Vincent Conitzer. Efficient Algorithms for Planning with Participation Constraints. *In Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, pp. 1121–1140, 2022.
 - [25] Hanrui Zhang, Yu Cheng, Vincent Conitzer. Planning with Participation Constraints. *In Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 5260–5267, 2022.
 - [24] Ruoxu Cen, Yu Cheng, Debmalya Panigrahi, Kevin Sun. Sparsification of Directed Graphs via Cut Balance. *In Proceedings of the 48th International Colloquium on Automata, Languages and Programming (ICALP)*, pp. 45:1–45:21, 2021.
 - [23] Yu Cheng, Honghao Lin. Robust Learning of Fixed-Structure Bayesian Networks in Nearly-Linear Time. *In Proceedings of the 9th International Conference on Learning Representations (ICLR)*, OpenReview.net, 2021.
 - [22] Anilesh Krishnaswamy, Zhihao Jiang, Kangning Wang, Yu Cheng, Kamesh Munagala. Fair for All: Best-Effort Fairness Guarantees for Classification. *In Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 3259–3267, 2021.
 - [21] Hanrui Zhang, Yu Cheng, Vincent Conitzer. Classification with Few Tests through Self-Selection. *In Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 5805–5812, 2021.
 - [20] Automated Mechanism Design for Classification with Partial Verification. *In Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 5789–5796, 2021.
 - [19] Yu Cheng, Ilias Diakonikolas, Rong Ge, Mahdi Soltanolkotabi. High-Dimensional Robust Mean Estimation via Gradient Descent. *In Proceedings of the 37th International Conference on Machine Learning (ICML)*, pp. 1768–1778, 2020.
 - [18] Hanrui Zhang, Yu Cheng, Vincent Conitzer. Distinguishing Distributions When Samples Are Strategically Transformed. *In Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS)*, pp. 3187–3195, 2019.
 - [17] Yu Cheng, Zhihao Jiang, Kamesh Munagala, Kangning Wang. Group Fairness in Committee Selection. *In Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, pp. 263–279, 2019.
 - [16] Yu Cheng, Ilias Diakonikolas, Rong Ge, David P. Woodruff. Faster Algorithms for High-Dimensional Robust Covariance Estimation. *In Proceedings of the 32nd Conference on Learning Theory (COLT)*, pp. 727–757, 2019.
 - [15] Hanrui Zhang, Yu Cheng, Vincent Conitzer. When Samples Are Strategically Selected. *In Proceedings of the 36th International Conference on Machine Learning (ICML)*, pp. 7345–7353, 2019.
 - [14] Hanrui Zhang, Yu Cheng, Vincent Conitzer. A Better Algorithm for Societal Tradeoffs. *In Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pp. 2229–2236, 2019.
 - [13] Yu Cheng, Ilias Diakonikolas, Rong Ge. High-Dimensional Robust Mean Estimation in Nearly-Linear Time. *In Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 2755–2771, 2019.
 - [12] Yu Cheng, Nick Gravin, Kamesh Munagala, Kangning Wang. A Simple Mechanism for a Budget-Constrained Buyer. *In Proceedings of the 14th Conference on Web and Internet Economics (WINE)*, pp. 96–110, 2018. **Best Paper Award.**
 - [11] Yu Cheng, Ilias Diakonikolas, Daniel M. Kane, Alistair Stewart. Robust Learning of Fixed-Structure Bayesian Networks. *In Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS)*, pp. 10304–10316, 2018.

- [10] Yu Cheng, Rong Ge. Non-Convex Matrix Completion Against a Semi-Random Adversary. *In Proceedings of the 31st Conference on Learning Theory (COLT)*, pp. 1362–1394, 2018.
- [9] Yu Cheng, Wade Hann-Caruthers, Omer Tamuz. A Deterministic Protocol for Sequential Asymptotic Learning. *In Proceedings of 2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 1735–1738, 2018.
- [8] Yu Cheng, Shaddin Dughmi, David Kempe. On the Distortion of Voting with Multiple Representative Candidates. *In Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pp. 973–980, 2018.
- [7] Yu Cheng, Shaddin Dughmi, David Kempe. Of the People: Voting Is More Effective with Representative Candidates. *In Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pp. 305–322, 2017.
- [6] Xi Chen, Yu Cheng, Bo Tang. Well-Supported versus Approximate Nash Equilibria: Query Complexity of Large Games. *In Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*, pp. 57:1–57:9, 2017.
- [5] Yu Cheng, Ilias Diakonikolas, Alistair Stewart. Playing Anonymous Games using Simple Strategies. *In Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 616–631, 2017.
- [4] Yu Cheng, Xi Chen, Bo Tang. On the Recursive Teaching Dimension of VC Classes. *In Proceedings of the 30th Conference on Neural Information Processing Systems (NeurIPS)*, pp. 2164–2171, 2016.
- [3] Umang Bhaskar, Yu Cheng, Young Kun Ko, Chaitanya Swamy. Hardness Results for Signaling in Bayesian Zero-Sum and Network Routing Games. *In Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pp. 479–496, 2016.
- [2] Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, Shang-Hua Teng. Mixture Selection, Mechanism Design, and Signaling. *In Proceedings of the 56th Symposium on Foundations of Computer Science (FOCS)*, pp. 1426–1445, 2015.
- [1] Dehua Cheng, Yu Cheng, Yan Liu, Richard Peng, Shang-Hua Teng. Efficient Sampling for Gaussian Graphical Models via Spectral Sparsification. *In Proceedings of the 28th Conference on Learning Theory (COLT)*, pp. 364–390, 2015.

Journal Papers

- [2] Yu Cheng, Nick Gravin, Kamesh Munagala, Kangning Wang. A Simple Mechanism for a Budget-Constrained Buyer. *In ACM Transactions on Economics and Computation (TEAC)*, 9(2), pp. 10:1–10:25, 2021.
- [1] Yu Cheng, Zhihao Jiang, Kamesh Munagala, Kangning Wang. Group Fairness in Committee Selection. *In ACM Transactions on Economics and Computation (TEAC)*, 8(4), pp. 23:1–23:18, 2020.

Talks

Scalable and Provably Robust Algorithms for Machine Learning.

- July 2022 *IFDS (Institute for Foundations of Data Science) Seminar*, University of Wisconsin-Madison, Madison, WI.
- May 2022 Brown University, virtual.
- Apr. 2022 University of Virginia, virtual.
- Apr. 2022 University of Michigan, Ann Arbor, MI.
- Apr. 2022 University of California, Santa Barbara, Santa Barbara, CA.
- Apr. 2022 Rutgers University, virtual.

- Mar. 2022 Northeastern University, Boston, MA.
High-Dimensional Robust Statistics: Faster Algorithms and Optimization Landscape.
- Mar. 2021 *CS Theory Lunch*, University of Southern California, virtual.
Of the People: Voting with Representative Candidates.
- June 2022 *Summer School on Game Theory and Social Choice*, City University of Hong Kong, virtual.
 July 2020 *Workshop on the Distortion and Information-Efficiency Tradeoffs (with EC 2020)*, virtual.
High-Dimensional Robust Mean Estimation via Gradient Descent.
- July 2020 *International Conference on Machine Learning (ICML)*, virtual.
Robustness and Strategic Aspects of Machine Learning.
- Dec. 2019 *IIS-Haihua Frontier Seminar*, Tsinghua University, Beijing, China.
- Oct. 2019 *Seminar on Theoretical Machine Learning*, Institute for Advanced Study, Princeton, NJ.
- Aug. 2019 *Algorithms Seminar*, Google, Mountain View, CA.
- May 2019 *Nanjing Theory Day*, Nanjing, China.
- Mar. 2019 *ITCS Seminar*, Shanghai University of Finance and Economics, Shanghai, China.
- Mar. 2019 University of Virginia, Charlottesville, VA.
- Mar. 2019 Pennsylvania State University, State College, PA.
- Mar. 2019 *Machine Learning Seminar*, Google, New York, NY.
- Mar. 2019 Washington University in St. Louis, St. Louis, MO.
- Feb. 2019 University of Utah, Salt Lake City, UT.
- Feb. 2019 University of Illinois at Chicago, Chicago, IL.
- Feb. 2019 Rensselaer Polytechnic Institute, Troy, NY.
- Feb. 2019 *John Hopcroft Center Seminar*, Shanghai Jiao Tong University, Shanghai, China.
Faster Algorithms for High-Dimensional Robust Covariance Estimation.
- June 2019 *Conference on Learning Theory (COLT)*, Phoenix, AZ.
High-Dimensional Robust Mean Estimation in Nearly-Linear Time.
- Feb. 2019 *Information Theory and Applications (ITA) Workshop*, San Diego, CA.
- Jan. 2019 *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, San Diego, CA.
- Oct. 2018 *ACO Student Seminar*, Georgia Institute of Technology, Atlanta, GA.
- Aug. 2018 *Workshop on Computational Efficiency and High-Dimensional Robust Statistics*, Toyota Technological Institute at Chicago, Chicago, IL.
Non-Convex Matrix Completion Against a Semi-Random Adversary.
- Aug. 2018 *Machine Learning Seminar*, Google, New York, NY.
- July 2018 *Conference on Learning Theory (COLT)*, Stockholm, Sweden.
- Apr. 2018 Microsoft Research, Redmond, WA.
On the Distortion of Voting with Multiple Representative Candidates.
- Feb. 2018 *AAAI Conference on Artificial Intelligence (AAAI)*, New Orleans, LA.
Playing Anonymous Games using Simple Strategies.
- July 2017 *China Theory Week*, Shanghai, China.
- Jan. 2017 *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Barcelona, Spain.
- Nov. 2016 *Southern California Theory Day*, Pasadena, CA.
Of the People: Voting Is More Effective with Representative Candidates.

- June 2017 *ACM Conference on Economics and Computation (EC)*, Cambridge, MA.
Computational Aspects of Optimal Information Revelation.
- Apr. 2017 *Algorithms Seminar*, Google, New York, NY.
- Mar. 2017 *CS-ECON Seminar*, Duke University, Durham, NC.
- Mar. 2017 *SISL (Social and Information Sciences Laboratory) Seminar*, California Institute of Technology, Pasadena, CA.
Random-Walk Sparsification, Newton's Method, and Gaussian Sampling.
- Mar. 2017 *Algorithms Seminar*, Duke University, Durham, NC.
Well-Supported vs. Approximate Nash Equilibria: Query Complexity of Large Games.
- Jan. 2017 *Innovations in Theoretical Computer Science Conference (ITCS)*, Berkeley, CA.
Hardness Results for Signaling in Bayesian Zero-Sum and Network Routing Games.
- July 2016 *ACM Conference on Economics and Computation (EC)*, Maastricht, The Netherlands.
Mixture Selection, Mechanism Design, and Signaling.
- Oct. 2015 *Symposium on Foundations of Computer Science (FOCS)*, Berkeley, CA.

Advising and Mentoring

Current Ph.D. Students

Xing Gao, Mathematics, University of Illinois at Chicago, 2020–present (co-advised with Lev Reyzin).

Former Undergraduate Students

Haichen Dong, Shanghai Jiao Tong University, 2021–2022. Now a Ph.D. student at Princeton University.

Honghao Lin, Shanghai Jiao Tong University, 2019–2020. Now a Ph.D. student at Carnegie Mellon University.

Ph.D. Committees Memberships (not as advisor)

Abdelrahman Hosny, Ph.D. in Computer Science, Brown University, 2022 (advisor: Sherief Reda).

Shelby Heinecke, Ph.D. in Mathematics, University of Illinois at Chicago, 2020 (advisor: Lev Reyzin).

Teaching

Courses Taught

Instructor, Algorithmic Aspects of Machine Learning (undergraduate), Brown University. Spring '23.

Instructor, Robust Algorithms for Machine Learning (graduate), Brown University. Fall '22.

Instructor, Computer Algorithms I (undergraduate), University of Illinois at Chicago (UIC). Fall '21, Fall '20.

Instructor, Codes and Cryptography (undergraduate), UIC. Fall '21, Fall '20, Spring '20.

Instructor, Spectral Graph Theory (graduate), UIC. Spring '20.

Teaching Assistant, Computational Microeconomics (graduate), Duke University. Fall '18.

Teaching Assistant, Randomized Algorithms (graduate), University of Southern California (USC). Fall '15.

Teaching Assistant, Convex and Combinatorial Optimization (graduate), USC. Fall '13.

Teaching Assistant, Discrete Methods in Computer Science (undergraduate), USC. Fall '12, Fall '11.

Competitive Programming

Student Coach, University of Southern California (USC), 2011–2015.

USC won the ACM-ICPC Southern California Regional Contest four years in a row (2011 to 2014).

Student Coach, Shanghai Jiao Tong University (SJTU), 2009–2011.

SJTU won the ACM-ICPC World Final Champion in 2010.

Professional Activities

Workshop Organization

Co-Organizer, Fall 2021 Special Quarter on *Robustness in High-dimensional Statistics and Machine Learning* at IDEAL (Institute for Data, Econometrics, Algorithms, and Learning, a collaborative institute in Chicago).

Program Committee Memberships

International Conference on Machine Learning (ICML), 2022. Selected as a top reviewer (top 10%).

Conference on Economics and Computation (EC), 2022.

International Conference on Artificial Intelligence and Statistics (AISTATS), 2022. Selected as a top reviewer (top 10%).

Conference on Web and Internet Economics (WINE), 2021.

Conference Reviewing

International Conference on Artificial Intelligence and Statistics (AISTATS). Conference on Learning Theory (COLT). Conference on Economics and Computation (EC). European Symposium on Algorithms (ESA). International Colloquium on Automata, Languages, and Programming (ICALP). International Conference on Machine Learning (ICML). Innovations in Theoretical Computer Science (ITCS). Symposium on Algorithmic Game Theory (SAGT). Symposium on Discrete Algorithms (SODA). Symposium on Theory of Computing (STOC). Conference on Web and Internet Economics (WINE).

Journal Reviewing

Artificial Intelligence (AIJ). Bernoulli Journal. Games and Economic Behavior (GEB). IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD). IEEE Transactions on Information Theory. Information Sciences. Journal of Machine Learning Research (JMLR). SIAM Journal on Computing (SICOMP). Social Network Analysis and Mining (SNAM).

Grants Reviewing

Panelist (virtual), NSF, *Computing and Communication Foundations (CCF): Algorithmic Foundations (AF)*, 2023.

Panelist (virtual), NSF, *Stimulating Collaborative Advances Leveraging Expertise in the Mathematical and Scientific Foundations of Deep Learning (SCALE MoDL)*, 2021.

Institutional Service

Ph.D. Admissions Committee, Department of Computer Science, Brown University, 2023.

Graduate Admissions Committee, Department of Mathematics, University of Illinois at Chicago, 2019–2021.

MCS Masters Exam Coordinator, Department of Mathematics, University of Illinois at Chicago, 2021.

Ph.D. Admissions Committee, Department of Computer Science, Duke University, 2019.

Organizer, Algorithms Seminars, Duke University, 2017–2019.

Co-Organizer, CS-ECON Seminar Series, Duke University, 2017–2019.

Organizer, Theory Reading Group, University of Southern California, 2014–2016.

Organizer, USC Programming Contest, University of Southern California, 2013–2015.

Internships

May–Aug. *Google*, New York, NY.

2013 Mentor: Konstantin Voevodski.

Projects: Local clustering algorithms (Geo team), parallel algorithms for set cover (Research team).

May–Aug. *Google*, Mountain View, CA.

2012 Mentor: Wen Xu.

Project: Load balancing algorithms (Infrastructure team).

July 2010–Feb. 2013 *Microsoft Research Asia*, Beijing, China.
Mentor: Chin-Yew Lin.
Project: Structured queries and knowledge base (Web Search and Mining group).
Sep. 2009–Jan. 2010 *APEX Data & Knowledge Management Lab*, Shanghai, China
Mentor: Yong Yu.
Project: 3D reconstruction from multiple images.