

Nika Haghtalab

BWW 8052
2121 Berkeley Way, UC Berkeley,
Berkeley, CA 94704

☎ +1 (412) 641 9495
✉ nika@berkeley.edu
🌐 www.cs.cornell.edu/~nika

Research Interests

Theory of Machine Learning, Algorithmic Economics, Social and Economic Aspects of Machine Learning.

Education

- 2013–2018 **Ph.D. in Computer Science**, *Carnegie Mellon University*.
Advisors: Avrim Blum and Ariel Procaccia.
Thesis Committee: Maria-Florina Balcan, Tim Roughgarden, and Robert Schapire.
- 2011–2013 **M.Math in Computer Science**, *University of Waterloo*.
Advisor: Shai Ben-David, thesis topic: Clustering with Background Noise.
Thesis Readers: Daniel Brown and Daniel Lizotte.
- 2008–2011 **B.Math in Computer Science and Combinatorics & Optimization**, *University of Waterloo*.
Dean's Honors List with Distinction.

Employment

- 2021– **Assistant Professor**, *Electrical Engineering and Computer Sciences Department*, UC Berkeley.
- 2019–2020 **Assistant Professor**, *Computer Science Department*, Cornell University.
- 2018–2019 **Postdoctoral Researcher**, Microsoft Research – New England.
- May–June **Visiting Student Researcher**, *Computer Science Department*, Stanford.
2017 Hosted by Tim Roughgarden.
- June–Sept. **Research Intern**, Microsoft Research — New York City.
2016 Worked with Miro Dudik, Rob Schapire, Jenn Wortman Vaughan, and Vasilis Syrgkanis.
Worked on problems related to oracle-efficient online learning and online auction design.
- May–Aug. **Research Intern**, Microsoft Research — Redmond.
2015 Worked with Ofer Dekel on problems related to beyond-worst case models in online learning.
- 2011–2013 **Graduate Research Assistant**, *Computer Science Department*, University of Waterloo.
Worked under the supervision of Shai Ben-David on problems related to clustering.

Selected Awards and Honors

- 2018-2019 SIGecom honorable mention Dissertation Award.
- 2018-2019 Carnegie Mellon University SCS Dissertation Award.
- 2017–2018 Siebel Scholarship.
- 2016–2018 Microsoft Research Ph.D. Fellowship.

2016–2018 Facebook Ph.D. Fellowship (Declined).

2015–2016 IBM Ph.D. Fellowship.

Publications (Authors are Ordered Alphabetically)

Book Chapters

- BC2 Maria-Florina Balcan, and **Nika Haghtalab**. Noise in Classification. In *Tim Roughgarden, editors, Beyond the Worst-Case Analysis of Algorithms, chapter 16, Cambridge University Press, Cambridge, United Kingdom, 2020*.
- BC1 Avrim Blum, **Nika Haghtalab**, and Ariel D. Procaccia. Learning to Play Stackelberg Security Games. In *Ali E. Abbas, Milind Tambe, and Detlof von Winterfeldt, editors, Improving Homeland Security Decisions, chapter 25. Cambridge University Press, Cambridge, United Kingdom, 2017*.

Journal Articles

- J5 Nima Anari, **Nika Haghtalab**, Seffi Naor, Sebastian Pokutta, Mohit Singh, and Alfredo Torrico. Structured Robust Submodular Maximization: Offline and Online Algorithms. *INFORMS Journal on Computing*, 2020 (forthcoming).
- J4 Miroslav Dudík, **Nika Haghtalab**, Haipeng Luo, Robert E Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. Oracle-efficient learning and auction design. *Journal of the ACM*, 67, 5, Article 26, October 2020.
- J3 Maria-Florina Balcan, **Nika Haghtalab**, and Colin White. k -center Clustering under Perturbation Resilience *ACM Trans. on Algorithms*, 16, 2, Article 22, April 2020.
- J2 Avrim Blum, John P. Dickerson, **Nika Haghtalab**, Ariel D. Procaccia, Tuomas Sandholm, and Ankit Sharma. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. *Operations Research*, Volume 68, Issue 1, 2020.
- J1 **Nika Haghtalab**, Aron Laszka, Ariel D. Procaccia, Yevgeniy Vorobeychik, and Xenofon Koutsoukos. Monitoring stealthy diffusion. *Knowledge and Information Systems*, pages 1–29, 2017.

Refereed Conference Proceedings

- C25 **Nika Haghtalab**, Tim Roughgarden, Abhishek Shetty. Smoothed Analysis of Online and Differentially Private Learning In *Proc. of the 33rd Annual Conference on Neural Information Processing Systems, (NeurIPS)*, 2020.
- C24 **Nika Haghtalab**, Nicole Immorlica, Brendan Lucier, Jack Wang. Maximizing Welfare with Incentive-Aware Evaluation Mechanisms In *Proc. of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- C23 Lydia T. Liu, Ashia Wilson, **Nika Haghtalab**, Adam Kalai, Christian Borgs, and Jennifer Chayes. The Disparate Equilibria of Algorithmic Decision Making when Individuals Invest Rationally. In *Proc. of the 3rd Annual ACM Conference on Fairness, Accountability, and Transparency, (FAT*)*, 2020.
- C22 **Nika Haghtalab**, Cameron Musco, and Bo Waggoner. Toward a Characterization of Loss Functions for Distribution Learning. In *Proc. of the 32nd Annual Conference on Neural Information Processing Systems, (NeurIPS)*, 2019.

- C21 Avrim Blum, **Nika Haghtalab**, Mohammadtaghi Hajiaghayi, and Saeed Seddighin. Computing Stackelberg Equilibria of Large General-Sum Games. In *Proc. of the 12th International Symposium on Algorithmic Game Theory (SAGT)*, 2019.
- C20 Christian Borgs, Jennifer Chayes, **Nika Haghtalab**, Adam Tauman Kalai, Ellen Vitercik. Algorithmic Greenlining: An Approach to Increase Diversity. In *Proc. of the 2nd ACM Conference on AI, Ethics, and Society (AIES)*, 2019.
- C19 Nima Anari, **Nika Haghtalab**, Seffi Naor, Sebastian Pokutta, Mohit Singh, and Alfredo Torrico. Robust Submodular Maximization: Offline and Online Algorithms. In *Proc. of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- C18 **Nika Haghtalab**, Simon Mackenzie, Ariel D. Procaccia, Oren Salzman, Siddhartha S. Srinivasa. The Provable Virtue of Laziness in Motion Planning. In *Proc. of the 28th International Conference on Automated Planning and Scheduling (ICAPS)*, 2018. Winner of the best paper award.
- C17 **Nika Haghtalab**, Ritesh Noothigattu, and Ariel D. Procaccia. Weighted Voting Via No-Regret Learning. In *Proc. of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- C16 Avrim Blum and **Nika Haghtalab**. Generalized topic modeling. In *Proc. of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- C15 Ofer Dekel, Authur Flajolet, **Nika Haghtalab**, Patrick Jaillet. Online Learning with a Hint. In *Proc. of the 30th Annual Conference on Neural Information Processing Systems, (NeurIPS)*, 2017.
- C14 Avrim Blum, **Nika Haghtalab**, Ariel D. Procaccia, and Mingda Qiao. Collaborative PAC Learning. In *Proc. of the 30th Annual Conference on Neural Information Processing Systems, (NeurIPS)*, 2017.
- C13 Miroslav Dudík, **Nika Haghtalab**, Haipeng Luo, Robert E Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. Oracle-efficient learning and auction design. In *Proc. of the 58th Annual Symposium on Foundations of Computer Science (FOCS)*, 2017.
- C12 Pranjal Awasthi, Avrim Blum, **Nika Haghtalab**, Yishay Mansour. Efficient PAC learning from the Crowd. In *Proc. of the 30th Conference on Learning Theory (COLT)*, 2017.
- C11 Avrim Blum, Ioannis Caragiannis, **Nika Haghtalab**, Ariel D. Procaccia, Eviatar B. Procaccia, and Rohit Vaish. Opting into optimal matchings. In *Proc. of the 28th Symposium on Discrete Algorithms (SODA)*, pages 2351–2363, 2017.
- C10 Pranjal Awasthi, Maria-Florina Balcan, **Nika Haghtalab**, and Hongyang Zhang. Learning and 1-bit compressed sensing under asymmetric noise. In *Proc. of the 29th Conference on Learning Theory (COLT)*, pages 152–192, 2016.
- C9 Maria-Florina Balcan, **Nika Haghtalab**, and Colin White. k-center clustering under perturbation resilience. In *Proc. of the 43rd International Colloquium on Automata, Languages, and Programming, (ICALP)*, pages 68:1–68:14, 2016.

- C8 **Nika Haghtalab**, Fei Fang, Thanh Hong Nguyen, Arunesh Sinha, Ariel D. Procaccia, and Milind Tambe. Three strategies to success: Learning adversary models in security games. In *Proc. of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 308–314, 2016.
- C7 Pranjal Awasthi, Maria-Florina Balcan, **Nika Haghtalab**, and Ruth Urner. Efficient learning of linear separators under bounded noise. In *Proc. of The 28th Conference on Learning Theory (COLT)*, pages 167–190, 2015.
- C6 **Nika Haghtalab**, Aron Laszka, Ariel D. Procaccia, Yevgeniy Vorobeychik, and Xenofon D. Koutsoukos. Monitoring stealthy diffusion. In *2015 IEEE International Conference on Data Mining (ICDM)*, pages 151–160, 2015.
- C5 Maria-Florina Balcan, Avrim Blum, **Nika Haghtalab**, and Ariel D. Procaccia. Commitment without regrets: Online learning in stackelberg security games. In *Proc. of the 16th Conference on Economics and Computation (EC)*, pages 61–78, 2015.
- C4 Avrim Blum, John P. Dickerson, **Nika Haghtalab**, Ariel D. Procaccia, Tuomas Sandholm, and Ankit Sharma. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. In *Proc. of the 16th Conference on Economics and Computation (EC)*, pages 325–342, 2015.
- C3 Avrim Blum, **Nika Haghtalab**, and Ariel D. Procaccia. Learning optimal commitment to overcome insecurity. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 1826–1834, 2014.
- C2 Shai Ben-David and **Nika Haghtalab**. Clustering in the presence of background noise. In *Proc. of the 31st International Conference on Machine Learning (ICML)*, pages 280–288, 2014.
- C1 Avrim Blum, **Nika Haghtalab**, and Ariel D. Procaccia. Lazy defenders are almost optimal against diligent attackers. In *Proc. of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 573–579, 2014.

Workshops and Symposiums

- W7 Pranjal Awasthi, Avrim Blum, **Nika Haghtalab**, Yishay Mansour. Efficient PAC learning from the Crowd In *Forecasting Workshop*, Concurrent with EC-17, 2017.
- W6 Avrim Blum, Ioannis Caragiannis, **Nika Haghtalab**, Ariel D. Procaccia, Eviatar B. Procaccia, and Rohit Vaish. Opting into optimal matchings. In *Workshop on Mechanism Design for Social Good*, Concurrent with EC-17, 2017.
- W5 Miroslav Dudík, **Nika Haghtalab**, Haipeng Luo, Robert E Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. Oracle-efficient Online learning and Auction Design. *3rd Workshop on Algorithmic Game Theory and Data Science*, Concurrent with EC-17, 2017.
- W4 Miroslav Dudík, **Nika Haghtalab**, Haipeng Luo, Robert E Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. Oracle-efficient Online learning and Auction Design. *11th Machine Learning Symposium, New York Academy of Sciences*, 2017.
Winner of a best spotlight presentation award.

- W3 Avrim Blum and **Nika Haghtalab**. Generalized topic modeling. *11th Machine Learning Symposium, New York Academy of Sciences*, 2017.
Winner of a best poster award.
- W2 Avrim Blum and **Nika Haghtalab**. Towards a General Theory of Multi-view Mixture Models. *ICML Workshop on Multi-View Representation Learning*, 2016.
- W1 Pranjali Awasthi, Maria-Florina Balcan, **Nika Haghtalab**, and Hongyang Zhang. Learning and 1-bit compressed sensing under asymmetric noise. *ICML Workshop on Non-convex Optimization*, 2016.

Talks

- 2020 **Learning for Decision Making: Dynamics and Economics**, *University of California–Berkeley, University of Washington*.
- 2019 **Incentive-Aware Classification**, *Simons Institute*.
- 2018–2019 **Foundations of Machine Learning, by the people, for the people**, *CalTech, Carnegie Mellon University, Cornell University, Duke University, GeorgiaTech ISyE, GeorgiaTech CS, Harvard University, Northwestern University, Stanford University, University of California–Berkeley, University of California–San Diego, University of Michigan, University of Pennsylvania, University of Wisconsin–Madison, University Massachusetts–Amherst*.
- 2017–2019 **Oracle-Efficient Online Learning**, *Cornell Probability Seminar, Economics and Computation AGT-Data Science workshop, Duke Theory Lunch, Foundations of Computer Science 2017, Harvard Econ-CS Seminar, INFORMS, New York Academy of Sciences, Stanford Theory Seminar, University of Massachusetts Amherst, University of Pennsylvania Theory Lunch*.
- 2016–2018 **Opting into Optimal Matchings**, *CalTech, Carnegie Mellon Theory Lunch, Economics and Computation 2017 Mechanism Design for Social Good workshop, Google, MIT Algorithm & Complexity Seminar, Simons Institute Uncertainty Seminar*.
- 2017–2019 **Efficient PAC Learning from the Crowd**, *Economics and Computation 2017 Forecasting workshop, TTIC Workshop on Learning in Presence of Strategic Behavior, INFORMS*.
- Mar. 2017 **Generalized Topic Modeling**, *AAAI’18, Carnegie Mellon Theory Lunch, Carnegie Mellon AI Seminar, International Conference on Machine Learning 2016 Workshop on Co-training, New York Academy of Sciences, University of California–San Diego*.
- May 2016 **Learning Halfspaces and 1-bit Compressed Sensing under Noise**, *Carnegie Mellon Theory Lunch, Conference on Learning Theory 2016, University of Pennsylvania Theory Lunch*.
- May 2016 **Three Strategies to Success: Learning Adversary Models in Security Games**, *International Joint Conference on Artificial Intelligence 2016*.
- Nov. 2015 **Monitoring Stealthy Diffusions**, *International Conference on Data Mining 2015*.
- Aug. 2015 **Online Learning with Hints**, *Microsoft Research–Redmond*.
- Jul. 2015 **Online Learning in Stackelberg Security Games**, *Economics and Computation 2015*.
- Jun. 2015 **k -Center Clustering Under Perturbation Resilience**, *Microsoft Research–Redmond*.

Jan.–Feb. **Learning to play Stackelberg Games**, *Carnegie Mellon Theory Lunch, University of Southern California*, 2015

Teaching

- Spring '20 **Instructor**, *Cornell University*.
CS6781: Foundations of Modern Machine Learning
- Fall '19-20 **Instructor**, *Cornell University*.
CS4780/CS5780: Introduction to Machine Learning
- Fall '15 **CS Teaching Assistant**, *Carnegie Mellon University*.
CS 10-806: Foundations of Machine Learning and Data Science
- Fall '14 **CS Teaching Assistant**, *Carnegie Mellon University*.
CS 15-381: Introduction to Artificial Intelligence
- Fall '12 **Instructor**, *Baha'i Institute for Higher Education*.
MCS4117-F12: Machine Learning

Professional Activities

Chairmanship

- 2020- **Co-Chair**, *Learning Theory Alliance*, The first mentorship program in learning theory.
- Dec. 2020 **Co-Chair**, *NeurIPS Workshop Machine Learning and Economic Policy*.
- Dec. 2019 **Co-Chair**, *NeurIPS Workshop Bridging Game Theory and Deep Learning*.
- Aug. 2019 **Co-Chair**, *IJCAI Workshop on "AI for Social Good"*.
- Jun. 2019 **Co-Chair**, *EC Workshop on "Learning in Presence of Strategic Behavior"*.
- Aug. 2018 **Co-Chair**, *TTIC Workshop on "Learning in Presence of Strategic Behavior"*.
- Dec. 2017 **Co-Chair**, *NeurIPS Workshop on "Learning in Presence of Strategic Behavior"*.

Senior Program Committee Member

- 2020 **International Conference on Algorithmic Learning Theory (ALT) 2020**.
- 2020 **Conference On Learning Theory (COLT) 2020**.
- 2020 **ACM Symposium on Theory of Computing (STOC), 2020**.
- 2019 **International Conference on Artificial Intelligence and Statistics, AISTATS, 2019**.
- 2019 **International Conference on Algorithmic Learning Theory (ALT) 2019**.
- 2018 **International Conference on Algorithmic Learning Theory (ALT) 2018**.

Program Committee Member

- 2019 **ACM Conference on Economics and Computation (EC) 2020**.
- 2019 **ACM Conference on Economics and Computation (EC) 2019**.
- 2017 **International Conference on Machine Learning (ICML) 2017**.
- 2017 **AAAI Conference on Artificial Intelligence (AAAI) 2017**.
- 2016 **International Conference on Machine Learning (ICML) 2016**.

Journal Refereeing

ACM Transactions on Economics and Computations (**TEAC**), Journal of Artificial Intelligence Research (**JAIR**), Journal of Machine Learning (**ML**).

Conference Refereeing

ACM Conference on Economics and Computation (**EC**), Conference On Learning Theory (**COLT**), Annual IEEE Symposium on Foundations of Computer Science (**FOCS**), Annual Conference on Neural Information Processing Systems (**NeurIPS**), ACM-SIAM Symposium on Discrete Algorithms (**SODA**), ACM Symposium on Theory of Computing (**STOC**), Theoretical Aspects of Rationality and Knowledge (**TARK**).