

Nika Haghtalab

Assistant Professor, EECS
University of California, Berkeley

Email : nika@berkeley.edu
Homepage: people.eecs.berkeley.edu/~nika
Address: 2021 Berkeley Way West, BWW 8028
Berkeley, California

EDUCATION

Ph.D. in Computer Science , Carnegie Mellon University Advisors: Avrim Blum, Ariel Procaccia	2013-2018
M.Math in Computer Science , University of Waterloo Advisor: Shai Ben-David	2011-2013
B.Math in Comp. Sci. and Combinatorics & Optimization , University of Waterloo, Dean's Honors List with Distinction	2008-2011

EMPLOYMENT

Assistant Professor , EECS, UC Berkeley	2021–
Assistant Professor , Computer Science, Cornell University	2019–2020
Postdoctoral Researcher , Microsoft Research – New England	2018–2019
Research Intern , Microsoft Research – New York City	06-09/2016
Research Intern , Microsoft Research – Redmond	05-08/2015

AWARDS AND HONORS

Alfred P. Sloan Fellowship	2024
Schmidt Sciences AI2050 Fellowship	2024
Google Research Scholar Award	2023
Exemplary Paper Award, AI Track, EC 2023	2023
Outstanding Paper Award, NeurIPS 2022	2022
NSF CAREER Award	2022
SIGecom honorable mention Dissertation Award	2018–2019
Best Paper Award, ICAPS 2018	2018
Carnegie Mellon University SCS Dissertation Award	2018
Carnegie Mellon University ACM Dissertation Nomination	2018
Rising Stars in EECS	2017
Siebel Scholarship	2017
Microsoft Research Ph.D. Fellowship	2016-2018
Facebook Ph.D. Fellowship (declined)	2016-2018
IBM Ph.D. Fellowship	2015-2016
Best presentation, <i>ML Symposium, New York Academy of Sciences.</i>	2017
Best poster, <i>ML Symposium, New York Academy of Sciences.</i>	2017

LEADERSHIP AND HIGHEST-LEVEL SERVICE

Steering Committee member , Association for Computational Learning Theory	2021–
Director , Association for Algorithmic Learning Theory	2021–
Program Co-Chair , Conference on Algorithmic Learning Theory	2022
Founding Co-Director , Center for the Foundations of Learning, Inference, Information, Intelligence, Math and Microeconomics at Berkeley (CLIMB)	2021–
Co-founder , Learning Theory Alliance (LET-ALL). The first mentorship program in learning theory	2020–
Co-founder , “Learning and Strategic Behavior” workshop series	2017–
Virtual Experience Co-Chair , Conference on Learning Theory (COLT)	2021
Co-Chair , NeurIPS Workshop Machine Learning and Economic Policy	2020
Co-Chair , NeurIPS Workshop Bridging Game Theory and Deep Learning	2019
Co-Chair , IJCAI Workshop on “AI for Social Good”	2019

REGULAR TEACHING

<i>CS 170: Efficient Algorithms and Intractable Problems</i> , UC Berkeley	2023–
<i>Data 102: Data, Inference, and Decisions</i> , UC Berkeley	2022–
<i>Foundations of Learning, Decisions, and Games</i> , UC Berkeley	2021–
<i>CS4780/CS5780: Introduction to Machine Learning</i> , Cornell University	2019–2020
<i>CS6781: Foundations of Modern Machine Learning</i> , Cornell University	2020
<i>Machine Learning</i> , Baha’i Institute for Higher Education	2012

RESEARCH ADVISING

PH.D. STUDENTS

Abhishek Shetty, EECS.	2019–present
Nivasini Ananthakrishnan, EECS co-advised with Jordan.	2021–present
Kunhe Yang, EECS.	2021–present
Eric Zhao, EECS co-advised with Jordan.	2021–present
Jessica Dai, EECS co-advised with Recht.	2022–present
Alex Wei, EECS co-advised with Jordan and Steinhardt.	2021–2023

POSTDOCTORAL RESEARCHERS

Yuval Dagan, FODSI.	2023–present
Mingda Qiao, FODSI.	2023–present
Paul Goelz, FODSI.	2024–present
Chara Podimata, FODSI. <i>Now faculty at MIT Sloan.</i>	2022–2023
Ellen Vitercik, Miller Postdoc (unofficial mentor). <i>Now faculty at Stanford</i>	2021–2022
Mahsa Derakhshan, FODSI. <i>Now faculty at Northeastern Univ.</i>	2021–2022

UNDERGRADUATE RESEARCH ADVISING

Naveen Durvasula, BEng/BBA 2023. <i>Now Ph.D student at Columbia U.</i>	2021–2023
Korinna Frangias, BEng 2023.	2022–2023

PUBLICATIONS

The authors are ordered alphabetically in almost all of my publications. Exceptions are marked by (*), in which typically the first and last authors led the project.

BOOK CHAPTERS

- BC2 M.F Balcan, and **N. Haghtalab**. Noise in Classification. In T. Roughgarden, editors, *Beyond the Worst-Case Analysis of Algorithms*, chapter 16, Cambridge University Press, United Kingdom, 2020.
- BC1 A. Blum, **N. Haghtalab**, and A.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

- J7 **N. Haghtalab**, T. Roughgarden, A. Shetty. Smoothed Analysis with Adaptive Adversaries. *Journal of the ACM*, forthcoming.
- J6 **N. Haghtalab**, M.O. Jackson, A.D. Procaccia. Belief Polarization in a Complex World: A Learning Theory Perspective. *Proc. of the National Academy of Science*, 118 (19) e2010144118, 2021.
- J5 A. Torrico, M. Singh, S. Pokutta, S. Naor, **N. Haghtalab**, N. Anari. Structured Robust Submodular Maximization: Offline and Online. *INFORMS Journal on Computing*, 33(4):1590–1607, 2021.
- J4 M. Dudík, **N. Haghtalab**, H. Luo, R.E. Schapire, V. Syrgkanis, and J. Wortman Vaughan. Oracle-efficient learning and auction design. *Journal of the ACM* 67(5):1–57, 2020.
- J3 M.F. Balcan, **N. Haghtalab**, and C. White. k -center Clustering under Perturbation Resilience. *ACM Transactions on Algorithms*, 16(2):1–30, 2020.
- J2 A. Blum, J.P. Dickerson, **N. Haghtalab**, A.D. Procaccia, T. Sandholm, and A. Sharma. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. *Operations Research*, 68(1):16–34, 2020.
- J1 **N. Haghtalab**, A. Laszka, A.D. Procaccia, Y. Vorobeychik, and Xenofon Koutsoukos. Monitoring stealthy diffusion. *Knowledge and Information Systems*, 52(3):1–29, 2017.

PEER-REVIEWED CONFERENCE PROCEEDINGS

- C44 N. Ananthakrishnan, S. Bates, M.I. Jordan, and **N. Haghtalab**. Delegating Data Collection in Decentralized Machine Learning. In *Proc. International Conference on Artificial Intelligence and Statistics*, (**AISTATS**), 2024.
- C43 J. Dai, B. Flanigan, **N. Haghtalab**, M. Jagadeesan, C. Podimata. Can Probabilistic Feedback Drive User Impacts in Online Platforms?. In *Proc. International Conference on Artificial Intelligence and Statistics*, (**AISTATS**), 2024.
- C42 C. Daskalakis, N. Golowich, **N. Haghtalab**, and A. Shetty.. Smooth Nash Equilibria: Algorithms and Complexity. In *Proc. Innovations in Theoretical Computer Science*, (**ITCS**) 2024.

- C41 N. Haghtalab, N. Immerlica, B. Lucier, M. Mobius, and D. Mohan. Communicating with Anecdotes. In *Proc. Innovations in Theoretical Computer Science, (ITCS)*, 2024.
- C40 A. Wei, **N. Haghtalab**, J. Steinhardt. Jailbroken: How Does LLM Safety Training Fail?. In *Proc. of the Conference on Neural Information Processing Systems, (NeurIPS)*, 2023. (Oral).
- C39 **N. Haghtalab**, C. Podimata, K. Yang. Calibrated Stackelberg Games: Learning Optimal Commitments Against Calibrated Agents. In *Proc. of the Conference on Neural Information Processing Systems, (NeurIPS)*, 2023.
- C38 A. Bhatt, **N. Haghtalab**, A. Shetty. Smoothed Analysis of Sequential Probability Assignment. In *Proc. of the Conference on Neural Information Processing Systems, (NeurIPS)*, 2023.
- C37 **N. Haghtalab**, M.I. Jordan, E. Zhao. A Unifying Perspective on Multi-Calibration: Game Dynamics for Multi-Objective Learning. In *Proc. of the Conference on Neural Information Processing Systems, (NeurIPS)*, 2023.
- C36 M. Jagadeesan, M.I. Jordan, J. Steinhardt, **N. Haghtalab**. Improved Bayes Risk Can Yield Reduced Social Welfare Under Competition. In *Proc. of the Conference on Neural Information Processing Systems, (NeurIPS)*, 2023.
- C35 P. Awashti, **N. Haghtalab**, and E. Zhao. Open Problem: The Sample Complexity of Multi-Distribution Learning for VC Classes. In *Proc. 36th Annual Conference on Learning Theory (COLT)*, 2023.
- C34 N. Durvasula, **N. Haghtalab**, and M. Zampetakis. Smoothed Analysis of Online Non-parametric Auctions. In *Proc. 24th ACM Conference on Economics and Computation (EC)*, 2023.
- C33 W. Guo, **N. Haghtalab**, K. Kandasamy, and E. Vitercik. Leveraging Reviews: Learning to Price with Buyer and Seller Uncertainty. In *Proc. 24th ACM Conference on Economics and Computation (EC)*, 2023.
- EXEMPLARY PAPER AWARD IN THE AI TRACK.**
- C32 M. Derakhshan, N. Durvasula, and **N. Haghtalab**. Stochastic Minimum Vertex Cover in General Graphs: a $3/2$ -Approximation. In *Proc. of the ACM Symposium on Theory of Computing (STOC)*, 2023.
- C31 M. Jagadeesan, M.I. Jordan, **N. Haghtalab**. Competition, Alignment, and Equilibria in Digital Marketplaces. In *Proc. of the AAAI Conference on Artificial Intelligence (AAAI)*, 2023 (*).
- C30 **N. Haghtalab**, Y. Han, A. Shetty, and K. Yang. Oracle-Efficient Online Learning for Beyond Worst-Case Adversaries. In *Proc. of the Conference on Neural Information Processing Systems, (NeurIPS)*, 2022.
- C29 **N. Haghtalab**, M.I. Jordan, and E. Zhao. On-Demand Sampling: Learning Optimally from Multiple Distributions. In *Proc. of the Conference on Neural Information Processing Systems, (NeurIPS)*, 2022.
- WINNER OF OUTSTANDING PAPER AWARD.**
- C28 **N. Haghtalab**, T. Lykouris, S. Nietert, A. Wei. Learning in Stackelberg Games with Non-myopic Agents. In *Proc. of the ACM Conference on Economics and Computation (EC)*, 2022.

- C27 **N. Haghtalab**, T. Roughgarden, A. Shetty. Smoothed Analysis with Adaptive Adversaries. In *Proc. of the Symposium on Foundations of Computer Science (FOCS)*, 2021.
- C26 A. Blum, **N. Haghtalab**, R. Phillips, H. Shao. One for One, or All for All: Equilibria and Optimality of Collaboration in Federated Learning. In *Proc. of the International Conference on Machine Learning, (ICML)*, 2021.
- C25 **N. Haghtalab**, T. Roughgarden, A. Shetty. Smoothed Analysis of Online and Differentially Private Learning. In *Proc. of the Conference on Neural Information Processing Systems, (NeurIPS)*, 2020.
- C24 **N. Haghtalab**, N. Immorlica, B. Lucier, J. Wang. Maximizing Welfare with Incentive-Aware Evaluation Mechanisms. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- C23 L.T. Liu, A. Wilson, **N. Haghtalab**, A. Kalai, C. Borgs, and J. Chayes. The Disparate Equilibria of Algorithmic Decision Making when Individuals Invest Rationally. In *Proc. of the ACM Conference on Fairness, Accountability, and Transparency, (FAT*)*, 2020(*).
- C22 **N. Haghtalab**, C. Musco, and B. Waggoner. Toward a Characterization of Loss Functions for Distribution Learning. In *Proc. of the Conference on Neural Information Processing Systems, (NeurIPS)*, 2019.
- C21 A. Blum, **N. Haghtalab**, M. Hajiaghayi, and S. Seddighin. Computing Stackelberg Equilibria of Large General-Sum Games. In *Proc. of the International Symposium on Algorithmic Game Theory (SAGT)*, 2019.
- C20 C. Borgs, J. Chayes, **N. Haghtalab**, A. Kalai, E. Vitercik. Algorithmic Greenlining: An Approach to Increase Diversity. In *Proc. of the ACM Conference on AI, Ethics, and Society (AIES)*, 2019.
- C19 N. Anari, **N. Haghtalab**, S. Naor, S. Pokutta, M. Singh, and A. Torricco. Robust Submodular Maximization: Offline and Online Algorithms. In *Proc. of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- C18 **N. Haghtalab**, S. Mackenzie, A.D. Procaccia, O. Salzman, S. Srinivasa. The Provable Virtue of Laziness in Motion Planning. In *Proc. of the International Conference on Automated Planning and Scheduling (ICAPS)*, 2018.
WINNER OF THE BEST PAPER AWARD .
- C17 **N. Haghtalab**, R. Noothigattu, and A.D. Procaccia. Weighted Voting Via No-Regret Learning. In *Proc. of the AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- C16 A. Blum and **N. Haghtalab**. Generalized topic modeling. In *Proc. of the AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- C15 O. Dekel, A. Flajolet, **N. Haghtalab**, P. Jaillet. Online Learning with a Hint. In *Proc. of the Conference on Neural Information Processing Systems, (NeurIPS)*, 2017.
- C14 A. Blum, **N. Haghtalab**, A.D. Procaccia, and M. Qiao. Collaborative PAC Learning. In *Proc. of the Conference on Neural Information Processing Systems, (NeurIPS)*, 2017.

- C13 M. Dudík, **N. Haghtalab**, H. Luo, R.E. Schapire, V. Syrgkanis, and J. Wortman Vaughan. Oracle-efficient learning and auction design. In *Proc. of the Symposium on Foundations of Computer Science (FOCS)*, 2017.
- C12 P. Awasthi, A. Blum, **N. Haghtalab**, Y. Mansour. Efficient PAC learning from the Crowd. In *Proc. of the Conference on Learning Theory (COLT)*, 2017.
- C11 A. Blum, I. Caragiannis, **N. Haghtalab**, A.D. Procaccia, E.B. Procaccia, and R. Vaish. Opting into optimal matchings. In *Proc. of the Symposium on Discrete Algorithms (SODA)*, 2017.
- C10 P. Awasthi, M.F. Balcan, **N. Haghtalab**, and H. Zhang. Learning and 1-bit compressed sensing under asymmetric noise. In *Proc. of the Conference on Learning Theory (COLT)*, 2016.
- C9 M.F. Balcan, **N. Haghtalab**, and C. White. k -center clustering under perturbation resilience. In *Proc. of the International Colloquium on Automata, Languages, and Programming, (ICALP)*, 2016.
- C8 **N. Haghtalab**, F. Fang, T. Nguyen, A. Sinha, A.D. Procaccia, and M. Tambe. Three strategies to success: Learning adversary models in security games. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2016 (*).
- C7 P. Awasthi, M.F. Balcan, **N. Haghtalab**, and R. Uner. Efficient learning of linear separators under bounded noise. In *Proc. of the Conference on Learning Theory (COLT)*, 2015.
- C6 **N. Haghtalab**, A. Laszka, A.D. Procaccia, Y. Vorobeychik, and X. Koutsoukos. Monitoring stealthy diffusion. In *IEEE International Conference on Data Mining (ICDM)*, 2015.
- C5 M.F. Balcan, A. Blum, **N. Haghtalab**, and A.D. Procaccia. Commitment without regrets: Online learning in stackelberg security games. In *Proc. of the Conference on Economics and Computation (EC)*, 2015.
- C4 A. Blum, J.P. Dickerson, **N. Haghtalab**, A.D. Procaccia, T. Sandholm, and A. Sharma. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. In *Proc. of the Conference on Economics and Computation (EC)*, 2015.
- C3 A. Blum, **N. Haghtalab**, and A.D. Procaccia. Learning optimal commitment to overcome insecurity. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- C2 S. Ben-David and **N. Haghtalab**. Clustering in the presence of background noise. In *Proc. of the International Conference on Machine Learning (ICML)*, 2014.
- C1 A. Blum, **N. Haghtalab**, and A.D. Procaccia. Lazy defenders are almost optimal against diligent attackers. In *Proc. of the AAAI Conference on Artificial Intelligence (AAAI)*, 2014.

KEYNOTES, SELECTED INVITED TALKS, AND HIGHLY-VISIBLE TUTORIALS

Video recording and slides are available for some tutorials on <https://people.eecs.berkeley.edu/~nika/#activities>

Keynote: Multi-Objective Learning, a unifying framework for fairness, robustness, and collaboration, *Algorithmic Learning Theory Conference [video]*

2023

Tutorial: Learning and Incentives , <i>Simons Institute Bootcamp on Learning and Games semester [part 1, part 2, part 3, part 4], Oxford's workshop on Social foundations for statistics and machine learning [video]</i> .	2022-2023
Tutorial: Multi-Distribution Learning, for Fairness, Robustness, and Collaboration , <i>Simons Institute Bootcamp on Data-Driven Processes [video]</i> .	2022-present
Sequential Decision Making Beyond the Worst-Case , <i>Workshop on Algorithms and Learning in Economics</i>	2022
Beyond worst-case adversaries in machine learning , <i>Bay Area Optimization, Stanford ISL seminar, MIT ORC Seminar, Stanford Theory seminar, International Information Theory Workshop, TRIPODS meeting</i>	2021-present
Multi-objective learning: A unifying framework for robustness, fairness, and collaboration/Collaborative Machine Learning , <i>Stanford RAIN seminar, Columbia IEOR-DRO Seminar, NeurIPS Intel TO3 meeting, BEARS</i>	2021-present
Learning and Persuading with Anecdotes , <i>Quarterly CS+Econ Workshop Northwestern University, Harvard AI for Economics seminar,</i>	2021-2022
Learning for Decision Making: Dynamics and Economics , <i>EECS Colloquium University of California–Berkeley, University of Washington</i>	2020
Foundations of Machine Learning, by the people, for the people , <i>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, Microsoft Research</i>	2018–2019
Oracle-Efficient Online Learning , <i>Cornell Probability Seminar, Duke Theory Lunch, Harvard Econ-CS Seminar, INFORMS, New York Academy of Sciences, Stanford Theory Seminar, University of Massachusetts Amherst, University of Pennsylvania Theory Lunch</i>	2017-2019
Opting into Optimal Matchings , <i>CalTech, Carnegie Mellon Theory Lunch, Google, MIT Algorithm & Complexity Seminar, Simons Institute Uncertainty Seminar</i>	2016–2018
Efficient PAC Learning from the Crowd , <i>Economics and Computation 2017 Forecasting workshop, TTIC Workshop on Learning in Presence of Strategic Behavior, INFORMS</i>	2017–2019
Learning with Unstructured Noise , <i>Conference on Learning Theory, University of Pennsylvania, Carnegie Mellon University, University of California, San Diego.</i>	2015–2017

