

Ellen Vitercik

vitercik@berkeley.edu
vitercik.github.io

Employment

- 2022- Assistant Professor Stanford University
Management Science and Engineering
Computer Science
- 2021-2022 Miller Fellow University of California, Berkeley
Hosts: Jennifer Chayes and Michael Jordan

Education

- 2021 PhD in Computer Science Carnegie Mellon University
Advisors: Maria-Florina Balcan and Tuomas Sandholm
Thesis committee: Eric Horvitz, Kevin Leyton-Brown, and Ameet Talwalkar
- 2018 MS in Computer Science Carnegie Mellon University
- 2015 BA in Mathematics, *summa cum laude* Columbia University
GPA: 4.01/4.33
- 2014 Budapest Semesters in Mathematics
GPA: 4.25/4.33

Honors and awards

- 2022-2025 Gabilan Fellowship
Stanford University
- 2022 Simons-Berkeley Research Fellowship (declined)
- 2021-2022 Miller Fellowship
University of California, Berkeley
- 2019 Best Presentation by a Student or Postdoctoral Researcher
ACM Conference on Economics and Computation (EC)
- 2019 Early Career Invited Lecture Award
UBC Science
- 2019-2021 IBM PhD Fellowship
- 2019-2020 Fellowship in Digital Health
Carnegie Mellon University's Center for Machine Learning and Health
- 2019 Exemplary Artificial Intelligence Track Paper Award
Awarded to one paper at the ACM Conference on Economics and Computation (EC)

- 2017 Teaching Assistant of the Year Award
Carnegie Mellon University's Machine Learning Department
- 2016-2019 National Science Foundation Graduate Research Fellowship
- 2016-2017 Microsoft Research Women's Fellowship
- 2015-2021 National Physical Science Consortium Fellowship (declined)
- 2015-2017 Kellett Fellowship (declined)
Full scholarship for postgraduate study at Oxford
- 2014 Phi Beta Kappa Junior Inductee
Awarded to the top 2% of the graduating Columbia College class
- 2012 Columbia University Class of 1956 Scholarship

Publications

CONFERENCE PAPERS

- 2021 Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm, and Ellen Vitercik.
Sample Complexity of Tree Search Configuration: Cutting Planes and Beyond.
Conference on Neural Information Processing Systems (NeurIPS).
- 2021 Ellen Vitercik and Tom Yan.
Revenue maximization via machine learning with noisy data.
Conference on Neural Information Processing Systems (NeurIPS).
- 2021 Maria-Florina Balcan, Dan DeBlasio, Travis Dick, Carl Kingsford, Tuomas Sandholm, and Ellen Vitercik.
How much data is sufficient to learn high-performing algorithms? Generalization guarantees for data-driven algorithm design.
ACM Symposium on Theory of Computing (STOC).
- 2021 Andrés Muñoz Medina, Umar Syed, Sergei Vassilvitskii, and Ellen Vitercik.
Private optimization without constraint violations.
International Conference on Artificial Intelligence and Statistics (AISTATS).
- 2021 Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik.
Generalization in portfolio-based algorithm selection.
AAAI Conference on Artificial Intelligence.
- 2020 Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik.
Refined bounds for algorithm configuration: The knife-edge of dual class approximability.
International Conference on Machine Learning (ICML).
- 2020 Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik.
Learning to optimize computational resources: Frugal training with generalization guarantees.
AAAI Conference on Artificial Intelligence.

- 2019 Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik.
 Estimating approximate incentive compatibility.
ACM Conference on Economics and Computation (EC).
Exemplary Artificial Intelligence Track Paper Award (awarded to one paper at EC 2019).
Best Presentation by a Student or Postdoctoral Researcher (EC 2019).
Invited to the ACM Transactions on Economics and Computation (TEAC) Special Issue for EC 2019.
- 2019 Daniel Alabi, Adam Kalai, Katrina Ligett, Cameron Musco, Christos Tzamos, and Ellen Vitercik.
 Learning to prune: Speeding up repeated computations.
Conference on Learning Theory (COLT).
- 2019 Christian Borgs, Jennifer Chayes, Nika Haghtalab, Adam Kalai, and Ellen Vitercik.
 Algorithmic greenlining: An approach to increase diversity.
AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES).
- 2018 Maria-Florina Balcan, Travis Dick, and Ellen Vitercik.
 Dispersion for data-driven algorithm design, online learning, and private optimization.
IEEE Symposium on Foundations of Computer Science (FOCS).
- 2018 Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik.
 A general theory of sample complexity for multi-item profit maximization.
ACM Conference on Economics and Computation (EC).
- 2018 Maria-Florina Balcan, Travis Dick, Tuomas Sandholm, and Ellen Vitercik.
 Learning to branch.
International Conference on Machine Learning (ICML).
- 2018 Bernhard Haeupler, Amirbehshad Shahrashbi, and Ellen Vitercik.
 Synchronization strings: Channel simulations and interactive coding for insertions and deletions.
International Colloquium on Automata, Languages and Programming (ICALP).
- 2017 Maria-Florina Balcan, Vaishnavh Nagarajan, Ellen Vitercik, and Colin White.
 Learning-theoretic foundations of algorithm configuration for combinatorial partitioning problems.
Conference on Learning Theory (COLT).
- 2016 Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik.
 Sample complexity of automated mechanism design.
Conference on Neural Information Processing Systems (NeurIPS).
- 2016 Maria-Florina Balcan, Ellen Vitercik, and Colin White.
 Learning combinatorial functions from pairwise comparisons.
Conference on Learning Theory (COLT).
- WORKSHOP PAPERS**
- 2020 Andrés Muñoz Medina, Umar Syed, Sergei Vassilvitskii, and Ellen Vitercik.
 Private optimization without constraint violations.
Theory and Practice of Differential Privacy Workshop (TPDP).
- 2019 Andrés Muñoz Medina, Umar Syed, Sergei Vassilvitskii, and Ellen Vitercik.
 Private linear programming without constraint violations.
Privacy in Machine Learning Workshop (PriML) at the Conference on Neural Information Processing Systems (NeurIPS).

- 2019 Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik.
A general theory of sample complexity for multi-item profit maximization.
ACM/INFORMS Workshop on Market Design at the Conference on Economics and Computation (EC).
- 2019 Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik.
Estimating approximate incentive compatibility.
Workshop on Machine Learning in the Presence of Strategic Behavior at the Conference on Economics and Computation (EC).
- 2018 Maria-Florina Balcan, Travis Dick, and Ellen Vitercik.
Dispersion for private optimization of piecewise Lipschitz functions.
Workshop on Privacy in Machine Learning and Artificial Intelligence at the International Conference on Machine Learning (ICML).
- 2018 Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik.
A general theory of sample complexity for multi-item profit maximization.
AAMAS-IJCAI Workshop on Agents and Incentives in Artificial Intelligence.
- 2017 Maria-Florina Balcan, Travis Dick, and Ellen Vitercik.
Differentially private algorithm configuration.
Workshop on Private Secure Machine Learning at the International Conference on Machine Learning (ICML).
- 2017 Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik.
Sample complexity of multi-item profit maximization.
Workshop on Algorithmic Game Theory and Data Science at the Conference on Economics and Computation (EC).

Tutorials and workshops

New Frontiers of Automated Mechanism Design for Pricing and Auctions

- 2021 AAAI Conference on Artificial Intelligence
with Maria-Florina Balcan and Tuomas Sandholm
- 2020 AAAI Conference on Artificial Intelligence
with Tuomas Sandholm
- 2019 ACM Symposium on Theory of Computing (STOC)
with Maria-Florina Balcan and Tuomas Sandholm
- 2019 Conference on Economics and Computation (EC)
with Maria-Florina Balcan and Tuomas Sandholm
- 2019 AAAI Conference on Artificial Intelligence
with Maria-Florina Balcan and Tuomas Sandholm
- 2018 International Conference on Machine Learning (ICML)
with Maria-Florina Balcan and Tuomas Sandholm under the title Machine Learning in Automated Mechanism Design for Pricing and Auctions

Selected talks

How Much Data is Sufficient to Learn High-Performing Algorithms?

- 2021 Worcester Polytechnic Institute, Computer Science Colloquium (Forthcoming)
- 2021 Purdue University, Theory Seminar
- 2021 Stanford University, Statistics Seminar
- 2021 Machine Learning for Algorithms Workshop, Foundation of Data Science Institute
- 2021 ACM Symposium on Theory of Computing (STOC)

- 2021 IPAM Workshop on Deep Learning and Combinatorial Optimization
2020 NeurIPS Workshop on Learning Meets Combinatorial Algorithms
2020 Stanford University CS Theory Lunch
2020 Columbia University Theory Seminar
- Generalization Guarantees For Multi-item Profit Maximization: Pricing, Auctions, And Randomized Mechanisms**
2021 INFORMS Annual Meeting *(Forthcoming)*
- Theoretical Foundations of Data-Driven Algorithm Design**
2021 Google Learning Theory Workshop *(Forthcoming)*
2021 Simons Institute, Meet the Fellows Welcome Event
- Automated Parameter Optimization for Integer Programming**
2021 AutoML Workshop at the International Conference on Machine Learning
- Private Optimization Without Constraint Violations**
2021 International Conference on Artificial Intelligence and Statistics (AISTATS)
- Integrating Machine Learning into Algorithm Design**
2021 University of Texas at Austin, Computer Science Seminar
2021 New York University, Computer Science Colloquium
2021 Columbia University, Computer Science Colloquium
2021 University of British Columbia, Computer Science Seminar
2021 University of Waterloo, Computer Science Seminar
2021 Harvard University, Computer Science Colloquium
2021 Princeton University, Computer Science Department Colloquium
2021 University of California, Los Angeles, Computer Science Seminar
2021 California Institute of Technology, Frontiers in Computing and Mathematical Sciences Symposium
2021 MIT Sloan, Operations Research and Statistics Seminar
2021 Stanford University, Management Sciences and Engineering Seminar
2021 Georgia Institute of Technology, School of Computer Science Seminar
2021 Microsoft Research New England, Seminar
2020 Columbia University, Industrial Engineering and Operations Research Seminar
- Generalization in Portfolio-Based Algorithm Selection**
2021 AAAI Conference on Artificial Intelligence
- Refined Bounds for Algorithm Configuration: The Knife-Edge of Dual Class Approximability**
2020 INFORMS Annual Meeting
2020 International Conference on Machine Learning
- Machine Learning as a Tool for Algorithm Design**
2020 Carnegie Mellon University, Open House for Admitted PhD Students
2019 University of British Columbia, Early Career Invited Lecture
- Learning to Prune: Speeding up Repeated Computations**
2020 Carnegie Mellon University, Open House for Admitted PhD Students
2019 Conference on Learning Theory (COLT)
- Estimating Approximate Incentive Compatibility**
2020 Young Researcher Workshop on Economics and Computation, Tel-Aviv University
2019 INFORMS Annual Meeting
2019 Carnegie Mellon University, Theory Lunch

2019 Conference on Economics and Computation (EC)
2019 EC Workshop on Machine Learning in the Presence of Strategic Behavior

Learning to Branch

2019 Cornell ORIE Young Researchers Workshop
2018 Carnegie Mellon University
2018 International Conference on Machine Learning (ICML)

A General Theory of Sample Complexity for Multi-Item Profit Maximization

2019 EC ACM/INFORMS Workshop on Market Design
2018 INFORMS Annual Meeting
2018 China Theory Week
2018 AAMAS-IJCAI Workshop on Agents and Incentives in Artificial Intelligence
2018 Conference on Economics and Computation (EC)

Dispersion for Data-Driven Algorithm Design, Online Learning, and Private Optimization

2018 Northwestern Quarterly Theory Workshop

Learning-Theoretic Foundations of Algorithm Configuration for Combinatorial Partitioning Problems

2018 INFORMS Annual Meeting

Sample Complexity of Multi-Item Profit Maximization

2017 Harvard University, Economics and CS Research Seminar
2017 Dagstuhl Workshop on *Game Theory Meets Computational Learning Theory*
2017 Workshop on Algorithmic Game Theory and Data Science at the Conference on Economics and Computation (EC)

Differentially Private Algorithm and Auction Configuration

2017 Carnegie Mellon University, Theory Lunch

Foundations of Application-Specific Algorithm Configuration

2017 Massachusetts Institute of Technology, Machine Learning Tea
2017 Microsoft Research New England, Machine Learning Lunch
2016 Carnegie Mellon University, Artificial Intelligence Lunch

Learning Submodular Functions from Pairwise Comparisons

2017 Carnegie Mellon University, Open House for Admitted PhD Students
2016 Conference on Learning Theory (COLT)

Sample Complexity of Automated Mechanism Design

2016 University of Pennsylvania, Theory Lunch
2016 Carnegie Mellon University, Theory Lunch

Teaching

Teaching assistant

2020 *Research and Innovation in Computer Science*, Carnegie Mellon University
2017 *Introduction to Machine Learning*, Carnegie Mellon University
2015 **Won the Machine Learning Department's Teaching Assistant of the Year Award.**
Computer Science Theory, Columbia University

Guest lecturer

- 2018 Machine Learning and Differential Privacy
Carnegie Mellon University course on Advanced Introduction to Machine Learning
- 2017 Introduction to Auction Design via Machine Learning
Carnegie Mellon University course on Advanced Introduction to Machine Learning
- 2017 Introduction to Research in Machine Learning
Carnegie Mellon University course on Research and Innovation in Computer Science

Mentoring

- 2018-2019 Rong He
Undergraduate student from Carnegie Mellon University.
- 2017 Mengxiao Zhang
Undergraduate student from Peking University.

Outreach

- 2021 Co-organizer of the Learning Theory Mentorship Initiative
Workshops at ALT'21 and COLT'21
- 2015-2020 Volunteer Instructor for Carnegie Mellon University TechNights
Workshop for middle school girls.
Sessions led: "Strategic Voting", "Game Theory", "Smashing Computers", and "Logic Puzzles".
- 2019 Session leader for Carnegie Mellon University OurCS
Workshop for undergraduate women in computer science.
Session led: "Machine Learning for Automated Algorithm Configuration".
- 2014-2015 Workshop Leader for Columbia University's Computer Science Emerging Scholars Program

Professional Services

- 2018-2020 Organizer of the Carnegie Mellon University Learning Theory Reading Group
- 2018 INFORMS Annual Meeting Session Chair
Session on Machine Learning and Optimization for Automated Mechanism Design.
- 2017-2018 PhD Admissions Committee Member at Carnegie Mellon University
- 2016-2017 Co-coordinator of the Carnegie Mellon University Artificial Intelligence Lunch and Seminar

Program Committee

- 2022 International Conference on Algorithmic Learning Theory (ALT)
- 2021 Conference on Web and Internet Economics (WINE)

Journal reviewing

INFORMS JOC '19; SIMODS '19; TPAMI '19; AIJ '19, '21; TEAC '20, '21; OR '20, '21; JACM '20, '21.

Conference reviewing

IJCAI '16; PODC '16; COLT '18; RANDOM '18; WINE '18; AIES '19; AISTATS '19; FOCS'19; ICML '17-'20; ESA'20; EC'20; SODA '18, '20, '21; ITCS '21; AAAI '21; STOC '17, '20, '21; NeurIPS '17-'21.

Typeset in Xe_{La}TeX