Ellen Vitercik

Email: vitercik@stanford.edu
Website: vitercik.github.io

Employment

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

Education

2021	PʜD in Computer Science *Advisors: Maria-Florina Balcan and Tuomas Sandholm* *Thesis committee: Eric Horvitz, Kevin Leyton-Brown, and Ameet Talwalkar*	Carnegie Mellon University
2018	MS in Computer Science	Carnegie Mellon University
2015	BA in Mathematics, *summa cum laude*	Columbia University

Honors and awards

2024	National Science Foundation CAREER Award
2023	Exemplary Artificial Intelligence Track Paper Award *Awarded to one paper at the ACM Conference on Economics and Computation (EC)*
2022-2025	Gabilan Fellowship *Stanford University*
2022	Robert N. Noyce Faculty Fellow *Stanford University*
2022	Simons-Berkeley Research Fellowship (declined)
2021-2022	Miller Fellowship *University of California, Berkeley*
2021	ACM SIGecom Dissertation Award *ACM Special Interest Group on Economics and Computation*
2021	Distinguished Dissertation Award *Carnegie Mellon University, School of Computer Science*
2021	Victor Lesser Distinguished Dissertation—Honorable Mention *International Foundation for Autonomous Agents and Multiagent Systems*

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

Journal papers

2024   Maria-Florina Balcan, Travis Dick, Tuomas Sandholm, and Ellen Vitercik.
Learning to branch: Generalization guarantees and limits of data-independent discretization.
Journal of the ACM (JACM).
Supersedes the ICML'20 and ICML'18 papers below.

2024   Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik.
Generalization guarantees for multi-item profit maximization: Pricing, auctions, and randomized mechanisms.
Operations Research (OR).
Supersedes the EC'18 paper below.

Conference papers

2023   Wenshuo Guo, Nika Haghtalab, Kirthevasan Kandasamy, and Ellen Vitercik.
Leveraging reviews: Learning to price with buyer and seller uncertainty.
ACM Conference on Economics and Computation (EC).
Exemplary Artificial Intelligence Track Paper Award (awarded to one paper at EC 2023).

2023   Christian Borgs, Jennifer Chayes, Christian Ikeokwu, and Ellen Vitercik.
        Disincentivizing polarization in social networks.
        ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO).

2022   Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm, and Ellen Vitercik.
        Structural analysis of branch-and-cut and the learnability of Gomory mixed integer cuts.
        Conference on Neural Information Processing Systems (NeurIPS).

2022   Wenshuo Guo, Michael I. Jordan, and Ellen Vitercik.
        No-regret learning in partially-informed auctions.
        International Conference on Machine Learning (ICML).

2022   Maria-Florina Balcan, Siddharth Prasad, Tuomas Sandholm, and Ellen Vitercik.
        Improved sample complexity bounds for branch-and-cut.
        International Conference on Principles and Practice of Constraint Programming (CP).

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 Shahrasbi, 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

2023   Christian Borgs, Jennifer Chayes, Christian Ikeokwu, and Ellen Vitercik.
Disincentivizing polarization in social networks.
3^{rd} *Workshop on Adverse Impacts and Collateral Effects of AI Technologies (AiOfAi) at the International Joint Conference on Artificial Intelligence (IJCAI).*

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

Machine Learning for Discrete Optimization: Theoretical Guarantees and Applied Frontiers
2024   AAAI Conference on Artificial Intelligence

Machine Learning for Algorithm Design: Theoretical Guarantees and Applied Frontiers
2023   Cargese-Porquerolles Workshop on Combinatorial Optimization

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

Workshops

Sampling and Optimization in Discrete Space
2023 International Conference on Machine Learning (ICML)
with Hanjun Dai, Priyank Jaini, Haoran Sun, and Ruqi Zhang

Selected talks

From Large to Small Datasets: Size Generalization for Clustering Algorithm Selection
2024 Massachusetts Institute of Technology, Theory of Computing Colloquium
2024 Oregon State University, AI Seminar
2024 Stanford University, Information Systems Laboratory Colloquium
2024 AAAI Workshop on Artificial Intelligence for Operations Research

Leveraging Reviews: Learning to Price with Buyer and Seller Uncertainty
2024 Simons Institute, Data-Driven Decision Processes Reunion
2023 Stanford University, Statistics Seminar
2023 University of Wisconsin-Madison, Systems, Information, Learning and Optimization Seminar
2023 INFORMS Annual Meeting
2023 University of Chicago, Booth School of Business, Operations Workshop Series
2023 Stanford Graduate School of Business, Operations, Information & Technology Seminar
2023 University of Massachusetts, Amherst, CS Theory Seminar
2023 Cornell Tech, Urban Tech Workshop
2023 IPAM Workshop on Artificial Intelligence and Discrete Optimization

Machine Learning for Algorithm Design
2023 Stanford University, CS & EE New Research Directions Workshop
2022 Simons Institute, Data-Driven Decision Processes Boot Camp

How Much Data is Sufficient to Learn High-Performing Algorithms?
2022 University of Massachusetts, Amherst, Algorithms with Predictions Seminar
2021 Worcester Polytechnic Institute, Computer Science Colloquium
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
2022 Google, Mountain View, Search-Ads Auction Spotlight Series
2021 INFORMS Annual Meeting

Theoretical Foundations of Machine Learning for Cutting Plane Selection
2022 Stanford University, Women's Theory Forum

Machine Learning for Tree Search Configuration: Cutting Planes and Beyond
2022 Simons Foundation Symposium on New Directions in Theoretical Machine Learning

Automated Algorithm and Mechanism Configuration
2022 Conference on Economics and Computation (EC)

Estimating Approximate Incentive Compatibility

2022 Algorithmic Game Theory: Past, Present, and Future (Workshop for Noam Nisan's 60th Birthday)
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

Sample Complexity of Tree Search Configuration: Cutting Planes and Beyond

2022 AAAI Workshop on Machine Learning for Operations Research
2022 STOC Workshop on Algorithms with Predictions

Private Optimization Without Constraint Violations

2022 Workshop on Algorithms for Learning and Economics (WALE)
2021 International Conference on Artificial Intelligence and Statistics (AISTATS)

Data-Driven Auction Design

2022 Miller Institute, UC Berkeley

Theoretical Foundations of Data-Driven Algorithm Design

2021 Google Learning Theory Workshop

Automated Parameter Optimization for Integer Programming

2021 AutoML Workshop at the International Conference on Machine Learning

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)

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

Primary instructor

2024 *Machine Learning for Discrete Optimization*, Stanford MS&E 236 / CS 225
2023 *Introduction to Probability*, Stanford MS&E 120
2023 *Machine Learning for Algorithm Design*, Stanford MS&E 331 / CS 331
2022 *Introduction to Probability*, Stanford MS&E 120

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

Research mentoring

PhD students

2023- Anders Wikum

2023- Ishani Karmarkar (co-advised with Aaron Sidford)

Masters students

2023- Alexandre Hayderi (Stanford University)

Undergraduate students

2023- Yash Dalmia (Stanford University)

2021-2023 Korinna Frangias (UC Berkeley)

2021-2023 Andrew Lin (UC Berkeley)

2021-2023 David Zhang (UC Berkeley)

2018-2019 Rong He (Carnegie Mellon University)

2017 Mengxiao Zhang (Peking University)

Outreach

2021- Co-organizer of the Learning Theory Alliance
Mentorship initiative for the machine learning theory community.

2024 Co-chair of *Mentoring Chats* at the International Conference on Learning Representations (ICLR)

2022 Teaching Assistant at the Institute for Advanced Studies' *Women and Mathematics* program

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 activities

Workshop chair
Conference on Economics and Computation (EC) 2024

Senior program committee
Conference on Learning Theory (COLT) 2024

Area chair
International Conference on Machine Learning (ICML) 2024

Program committee
Conference on Economics and Computation (EC) 2023
Conference on Fairness, Accountability, and Transparency (FAccT) 2023
Conference on Web and Internet Economics (WINE) 2021
Innovations in Theoretical Computer Science (ITCS) 2023
International Conference on Algorithmic Learning Theory (ALT) 2022, 2023
Symposium on Discrete Algorithms (SODA) 2024

Journal editorial boards
2024- Action Editor, Transactions on Machine Learning Research (TMLR)

Journal reviewing
ACM Transactions on Economics and Computation (TEAC) 2020, 2021
Artificial Intelligence (AIJ) 2019, 2021
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 2019
INFORMS Journal on Computing 2019
INFORMS Journal on Optimization 2022
Journal of the ACM (JACM) 2020, 2021, 2023
Management Science 2022
Nature Advances 2023
Operations Research (OR) 2020, 2021
SIAM Journal on Mathematics of Data Science (SIMODS) 2019

Conference reviewing
AAAI Conference on Artificial Intelligence 2021
Conference on Artificial Intelligence, Ethics, and Society (AIES) 2019
Conference on Economics and Computation (EC) 2020
Conference on Learning Theory (COLT) 2018
Conference on Neural Information Processing Systems (NeurIPS) 2017, 2018, 2019, 2020, 2021
European Symposium on Algorithms (ESA) 2020
Innovations in Theoretical Computer Science (ITCS) 2021, 2022
International Colloquium on Automata, Languages and Programming (ICALP) 2022
International Conference on Artificial Intelligence and Statistics (AISTATS) 2019
International Conference on Learning Representations (ICLR) 2022
International Conference on Machine Learning (ICML) 2017, 2018, 2019, 2020
International Conference on Randomization and Computation (RANDOM) 2018
International Joint Conference on Artificial Intelligence (IJCAI) 2016
Symposium on Discrete Algorithms (SODA) 2018, 2020, 2021, 2023
Symposium on Foundations of Computer Science (FOCS) 2019
Symposium on Principles of Distributed Computing (PODC) 2016
Symposium on Theory of Computing (STOC) 2017, 2020, 2021, 2024
Conference on Web and Internet Economics (WINE) 2018

Session Chair
INFORMS Annual Meeting, 2018

University service

Stanford University
2022 PhD Admissions Committee Member, Computer Science Department

Carnegie Mellon University

2017-2018   PhD Admissions Committee Member, Computer Science Department

2016-2017   Co-coordinator of the Artificial Intelligence Lunch and Seminar

Typeset in X͟ɟTEX