

Jacob D. Abernethy

Associate Professor, School of Computer Science, Georgia Institute of Technology
Research Scientist, Google Research

Atlanta, Georgia · prof@gatech.edu · [Website](#) · [Google Scholar](#) · [Google Research](#) · [ORCID](#)

Research Interests

Machine learning theory; online learning and sequential decision-making; optimization; learning in games and game theory; market design; and the dynamics of learning and decision-making systems, with recent work on the theoretical and algorithmic foundations of modern AI systems.

Academic and Research Appointments

Google Research — Research Scientist December 2022–present

Major contributor to the Gemini Large-Scale Pretraining team, January 2025–present.

Georgia Institute of Technology, School of Computer Science — Associate Professor March 2020–present

20% appointment since December 2022.

Georgia Institute of Technology, School of Computer Science — Assistant Professor August 2017–February 2020

University of Michigan, Computer Science and Engineering — Assistant Professor September 2013–August 2017

University of Pennsylvania — Simons Postdoctoral Fellow October 2011–August 2013

Host: Michael Kearns.

Visiting Appointments

Simons Institute for the Theory of Computing, UC Berkeley — Visiting Scholar Spring 2022

Northwestern University, Special Quarter on Data Science and Online Markets — Visiting Professor April–June 2018

Simons Institute for the Theory of Computing, UC Berkeley — Visiting Scholar Spring 2017

Education

University of California, Berkeley — PhD in Computer Science 2011

Advisor: Peter L. Bartlett.

Toyota Technological Institute at Chicago — MS in Computer Science 2006

Massachusetts Institute of Technology — BS in Mathematics 2002

Commonwealth of Massachusetts — General Educational Development (GED) 1998

Honors and Awards

Best Student Paper Award, ACM SIGKDD 2018

ActiveRemediation: The Search for Lead Pipes in Flint, Michigan.

IBM Research Pat Goldberg Memorial Best Paper Award 2008

Competing in the Dark: An Efficient Algorithm for Bandit Linear Optimization.

Best Student Paper Award, Conference on Learning Theory 2008

Competing in the Dark: An Efficient Algorithm for Bandit Linear Optimization.

Yahoo! PhD Fellowship 2008–2010

One of four recipients of two years of doctoral support.

\$5,000 research award.

Entrepreneurship and Public-Interest Leadership

BlueConduit — Co-founder and Board Member

August 2019–present

Co-founded with Eric Schwartz to develop data and software tools for lead-service-line inventory and replacement, and to support water utilities with software solutions to infrastructure challenges.

In8, Inc. — Co-founder

March–December 2017

Co-founded with Elad Hazan to develop machine-learning systems for autonomous-vehicle control; acquired by Google in December 2017.

Flint Water Crisis and Lead-Service-Line Infrastructure

Developed statistical and machine-learning tools with Eric Schwartz and the Michigan Data Science Team to identify homes likely to have hazardous lead service lines during the Flint water crisis. The work informed field operations and grew into BlueConduit, which supports water utilities with lead-service-line inventory and replacement.

Michigan Data Science Team

Founded and advised the Michigan Data Science Team, a student-led organization using data science for high-impact community and public-interest projects. The model was highlighted by the National Academies and adapted by student groups at other universities.

Selected Press

The Irish Times — AI at forefront of huge Flint pipe reconstruction effort

May 2, 2019

NRDC — Flint Safe Water Case Update: 2019 Terms Will Improve Efficiency in Pipe-Replacement Program

February 13, 2019

The Atlantic — How a Feel-Good AI Story Went Wrong in Flint

January 3, 2019

New Scientist — Flint water crisis: How AI is finding thousands of hazardous pipes

August 22, 2018

Georgia Tech — Using Data Science to Fix the Flint Water Crisis

August 20, 2018

Michigan Today — Get the lead out

August 20, 2018

The New York Times — Michigan Allots \$87 Million to Replace Flint’s Tainted Water Pipes

March 27, 2017

Gizmodo — Google Is Helping Flint Prioritize Which Lead Pipes Need to Be Replaced

May 3, 2016

Sponsored Research and Gifts

Sponsored Research

National Science Foundation — Collaborative Research: CIF: Medium: An Algorithmic Theory of Hamiltonian Dynamics for Sampling, Optimization, and Games

August 2024–July 2028

Principal Investigator · \$399,998 intended total · CCF-2403390

National Science Foundation — RI: Small: Training Modularized Learning Systems

September 2019–September 2022

Principal Investigator · \$465,000

Google — Development of Sequential Optimization Algorithms for Machine Learning

2018

Principal Investigator · \$100,000 unrestricted gift

Captricity — Development of Background Detection in Registration

January–August 2017

Principal Investigator · \$45,000

Google.org — Mapping Critical Data for the Flint Water Crisis

May 2016–March 2017

Principal Investigator · \$200,000

National Science Foundation — RI: Small: When Algorithms Trade: Dynamics, Limits, and Economic Implications

June 2014–June 2018

Co-Principal Investigator; PI Michael Wellman · \$516,018

National Science Foundation — CAREER: Machine Learning through the Lens of Economics (And Vice Versa) February 2015–February 2020

Principal Investigator · \$500,000

Research Gifts

Apple Inc. — Research Gift — Jacob Abernethy/Guanghui Wang — Apple Scholars

August 2025–August 2032

Principal Investigator · \$111,631.30 · AWD-007752 / GR00035361

Students and Postdoctoral Researchers

Current PhD Students

Guanghui Wang

2021–present

PhD student in Machine Learning; co-advised with Vidya Muthukumar. Apple Scholar in AI/ML, 2025.

Current Postdoctoral Researchers

John Lazarsfeld

August 2026–present

ARC Postdoctoral Fellow; hosted with Vidya Muthukumar and Juba Ziani.

PhD and MS Alumni

Tyler LaBonte

2026

Co-advised with Vidya Muthukumar; NDSEG and NSF Graduate Research Fellow. Dissertation: Generalization in Dynamic Environments: Foundations and Algorithms. Senior Researcher, Microsoft Research AI Frontiers.

Yeo Joon Youn

2024

Co-advised with Vidya Muthukumar. Dissertation: Advanced Quantization Techniques for Communication Efficiency and Privacy in Federated Learning, and Memory-efficient Fine-tuning of LLMs. AI Engineer, LinkedIn.

Zihao Hu

2023

Research Assistant Professor, Department of Mathematics, HKUST.

Bhuvesh Kumar

2023

Co-advised with Jamie Morgenstern. Research Scientist, Snap Research.

Jun-Kun Wang

2020

Assistant Professor, UC San Diego ECE and Halicioğlu Data Science Institute.

Kevin Lai

2020

Voleon Capital.

Benjamin Bray

MS 2019

Left the PhD program with an MS.

Chansoo Lee

2018

Co-advised with Ambuj Tewari. Google DeepMind.

Sindhu Kutty

2015

Previously advised by Rahul Sami. Lecturer IV, EECS, University of Michigan.

Postdoctoral Alumni

Andre Wibisono

2018–2021

Assistant Professor, Yale Computer Science.

Matus Telgarsky

2015–2017

Researcher, OpenAI.

Selected Thesis Committees

Qingsi Wang, Erik Brinkman, Elaine Wah, Kareem Amin, Audra McMillan, Nan Jiang.

Teaching

CS 8803-SP: Sequence Prediction — Georgia Tech	Spring 2026
CS 7545: Machine Learning Theory — Georgia Tech	Spring 2024
CS 7545: Machine Learning Theory — Georgia Tech	Spring 2023
CS 3510: Design and Analysis of Algorithms — Georgia Tech	Fall 2022
CS 3510: Design and Analysis of Algorithms — Georgia Tech	Fall 2021
CS 7545: Machine Learning Theory — Georgia Tech	Fall 2018
CS 4540: Advanced Algorithms for Machine Learning — Georgia Tech	Fall 2018
CS 7545: Machine Learning Theory — Georgia Tech	Fall 2017
EECS 445: Undergraduate Machine Learning — University of Michigan	Fall 2016
EECS 545: Graduate Machine Learning — University of Michigan	Spring 2016
EECS 598: Theory of Machine Learning — University of Michigan	Fall 2015
EECS 203: Introduction to Discrete Mathematics — University of Michigan	Fall 2014
EECS 203: Introduction to Discrete Mathematics — University of Michigan	Spring 2014
EECS 598: Prediction and Learning: It's Only a Game — University of Michigan	Fall 2013

Publications

Book Chapters

Abernethy, J. D., Lee, C., & Tewari, A. (2016). Perturbation techniques in online learning and optimization. In *Perturbations, Optimization, and Statistics*, 233–256.

Peer-Reviewed Journal Articles

Sánchez Vargas, F., Potluri, S. R., **Abernethy, J. D.**, & Benítez, M. E. (2026). CapuchinAI 1.0: Development of a machine learning-based touchscreen paradigm to test cognition in wild capuchins. *American Journal of Primatology*. Published online July 28, 2026. DOI: 10.1002/ajp.70194.

Wang, G., Hu, Z., Gentile, C., Muthukumar, V., & **Abernethy, J. D.** (2025). Faster margin maximization rates for generic and adversarially robust optimization methods . *Mathematical Programming* . Published online October 9, 2025. DOI: 10.1007/s10107-025-02283-4.

Wang, J.-K., **Abernethy, J. D.**, & Levy, K. Y. (2024). No-regret dynamics in the Fenchel game: a unified framework for algorithmic convex optimization . *Mathematical Programming*, 205 (1–2), 203–268.

Youn, Y., Kumar, B., & **Abernethy, J. D.** (2023). Accelerated federated optimization with quantization . *IEEE Data Engineering Bulletin*, 46 (1), 79–123.

Misra, K., Schwartz, E. M., & **Abernethy, J.** (2019). Dynamic Online Pricing with Incomplete Information Using Multiarmed Bandit Experiments. *Marketing Science* .

Abernethy, J., Chen, Y., & Vaughan, J. (2013). Efficient Market Making via Convex Optimization, and a Connection to Online Learning. *ACM Transactions on Economics and Computation* , 1(2), 1-39.

Abernethy, J., Hazan, E., & Rakhlin, A. (2012). Interior-Point Methods for Full-Information and Bandit Online Learning. *IEEE Transactions on Information Theory* , 58(7), 4164-4175.

Abernethy, J., Chapelle, O., & Castillo, C. (2010). Graph regularization methods for Web spam detection. *Machine Learning* , 81(2), 207-225.

Abernethy, J., Bach, F., Evgeniou, T., & Vert, J. (2009). A New Approach to Collaborative Filtering: Operator Estimation with Spectral Regularization. *Journal of Machine Learning Research* , 10, 803-826.

Abernethy, J., Evgeniou, T., Toubia, O., & Vert, J. (2008). Eliciting Consumer Preferences Using Robust Adaptive Choice Questionnaires. *IEEE Transactions on Knowledge and Data Engineering* , 20(2), 145-155.

Abernethy, J. D., Frongillo, R. M., & Kuty, S. (2014). On risk measures, market making, and exponential families. *ACM SIGecom Exchanges*, 13(2), 21–25.

Refereed Conference Papers

Wang, G., Syed, U., Schapire, R. E., & **Abernethy, J. D.** (2026). Multi-distribution learning: from worst-case optimality to lexicographic min-max optimality . 37th International Conference on Algorithmic Learning Theory (ALT 2026) . Accepted.

Pan, L., Ganesh, V., **Abernethy, J. D.**, Esposito, C., & Lee, W. (2025). Can transformers reason logically? A study in SAT solving . Proceedings of the 42nd International Conference on Machine Learning, PMLR 267 , 47632–47671.

Pan, L., Esposito, C., **Abernethy, J. D.**, Ganesh, V., & Lee, W. (2025). Symbolic constraint-solving capabilities of transformer large language models . Proceedings of the 1st GENZERO Workshop: Revolutionizing Autonomous Systems with Generative AI , 147–154. DOI: 10.1007/978-981-95-1050-4_18.

Hu, Z., Wang, G., Wang, X., Wibisono, A., **Abernethy, J. D.**, & Tao, M. (2024). Extragradient type methods for Riemannian variational inequality problems . Proceedings of AISTATS 2024, PMLR 238 , 2080–2088.

Abernethy, J. D., Schapire, R. E., & Syed, U. (2024). Lexicographic optimization: algorithms and stability . Proceedings of AISTATS 2024, PMLR 238 , 2503–2511.

Abernethy, J. D., Agarwal, A., Marinov, T. V., & Warmuth, M. K. (2024). A mechanism for sample-efficient in-context learning for sparse retrieval tasks . Proceedings of ALT 2024, PMLR 237 , 3–46.

Luo, F., Liu, L., Wang, G. G., Kumar, V., Ashton, M. S., **Abernethy, J. D.**, et al. (2023). Artificial intelligence for climate smart forestry: a forward looking vision . Proceedings of IEEE CogMI 2023 , 1–10. DOI: 10.1109/CogMI58952.2023.00011.

Hu, Z., Wang, G., & **Abernethy, J. D.** (2023). Minimizing dynamic regret on geodesic metric spaces . Proceedings of COLT 2023, PMLR 195 , 4336–4383.

Wang, G., Hanashiro, R., Guha, E. K., & **Abernethy, J. D.** (2023). On accelerated perceptrons and beyond . International Conference on Learning Representations (ICLR 2023) .

Hu, Z., Wang, G., & **Abernethy, J. D.** (2023). Riemannian projection-free online learning . Advances in Neural Information Processing Systems 36 , 41980–42014.

Wang, G., Hu, Z., Muthukumar, V., & **Abernethy, J. D.** (2023). Faster margin maximization rates for generic optimization methods . Advances in Neural Information Processing Systems 36 , 62488–62518.

Abernethy, J. D., Awasthi, P., Kleindessner, M., Morgenstern, J., Russell, C., & Zhang, J. (2022). Active sampling for min-max fairness . Proceedings of ICML 2022, PMLR 162 , 53–65.

Kumar, B., **Abernethy, J. D.**, & Saligrama, V. (2022). ActiveHedge: Hedge meets active learning . Proceedings of ICML 2022, PMLR 162 , 11694–11709.

Wang, G., Hu, Z., Muthukumar, V., & **Abernethy, J. D.** (2022). Adaptive oracle-efficient online learning . Advances in Neural Information Processing Systems 35 , 23398–23411.

Xu, Y., Kumar, B., **Abernethy, J.**, (2021, December). Observation-Free Attacks on Stochastic Bandits. In Advances in Neural Information Processing Systems 34 (NeurIPS 2021).

Wang, J. K., **Abernethy, J.**, (2021, November). Understanding How Over-Parametrization Leads to Acceleration: A case of learning a single teacher neuron. In Proceedings of the 13 th Asian Conference on Machine Learning. PMLR 157:17-32, 2021.

Abernethy, J. D., Lai, K., Wibisono, A. (2021, January). Fast convergence of fictitious play for diagonal payoff matrices. In Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA).

Wang, J-K., Lin, C-H, Abernethy (2021, July). A Modular Analysis of Provable Acceleration via Polyak’s Momentum: Training a Wide ReLU Network and a Deep Linear Network. In Proceedings of the 38th International Conference on Machine Learning, PMLR 139:10816-10827, 2021.

Abernethy, J. D., Lai, K. A., & Wibisono, A. (2021). Last-iterate convergence rates for min-max optimization: convergence of Hamiltonian gradient descent and consensus optimization . Proceedings of ALT 2021, PMLR 132 , 3–47.

Wang, J.-K., Lin, C.-H., & **Abernethy, J. D.** (2020). Escaping saddle points faster with stochastic momentum . International Conference on Learning Representations (ICLR 2020) .

Abernethy, J. D., Jung, Y. H., Lee, C., McMillan, A., & Tewari, A. (2019). Online learning via the differential privacy lens . Advances in Neural Information Processing Systems 32 , 8892–8902.

- Abernethy, J. D.**, Cummings, R., Kumar, B., Taggart, S., & Morgenstern, J. (2019). Learning auctions with robust incentive guarantees . *Advances in Neural Information Processing Systems* 32 , 11587–11597.
- Rivera Cardoso, A., **Abernethy, J. D.**, Wang, H., & Xu, H. (2019). Competing against Nash equilibria in adversarially changing zero-sum games . *Proceedings of ICML 2019*, PMLR 97 , 921–930.
- Wang, J. K., & **Abernethy, J. D.** (2018). Acceleration through optimistic no-regret dynamics. In *Advances in Neural Information Processing Systems* (pp. 3824-3834).
- Abernethy, J.**, Chojnacki, A., Farahi , A., Schwartz, E., & Webb, J. (2018, July). ActiveRemediation: The Search for Lead Pipes in Flint, Michigan. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 5-14). ACM.
- Abernethy, J.**, Lai, K. , Levy, K. Y., & Wang, J. K. (2018). Faster Rates for Convex-Concave Games. In *Proceedings of the 31st Conference On Learning Theory* (Vol. 75, pp. 1595-1625). PMLR.
- Abernethy, J. D.**, & Wang, J. K. (2017). On frank-wolfe and equilibrium computation. In *Advances in Neural Information Processing Systems* (pp. 6584-6593).
- Chojnacki, A., Dai, C., Farahi, A., Shi, G., Webb, J., Zhang, D.T., **Abernethy, J.** and Schwartz, E. (2017, August). A data science approach to understanding residential water contamination in flint. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1407-1416). ACM.
- Azab, M., Mihalcea, R., & **Abernethy, J.** (2016, November). Analysing Ratemyprofessors evaluations across institutions, disciplines, and cultures: The tell-tale signs of a good professor. In *International Conference on Social Informatics* (pp. 438-453). Springer, Cham.
- Li, J., Ozog, P., **Abernethy, J.**, Eustice, R. M., & Johnson-Roberson, M. (2016, October). Utilizing high-dimensional features for real-time robotic applications: Reducing the curse of dimensionality for recursive bayesian estimation. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 1230-1237). IEEE.
- Abernethy, J. D.**, Amin, K., & Zhu, R. (2016). Threshold bandits, with and without censored feedback . *Advances in Neural Information Processing Systems* 29 , 4889–4897.
- Abernethy, J.**, Lahaie, S., & Telgarsky, M. (2016, July). Rate of Price Discovery in Iterative Combinatorial Auctions. In *Proceedings of the 2016 ACM Conference on Economics and Computation* (pp. 809-809). ACM.
- Abernethy, J.**, & Hazan, E. (2016, June). Faster Convex Optimization: Simulated Annealing with an Efficient Universal Barrier. In *International Conference on Machine Learning* (pp. 2520-2528).
- Abernethy, J.**, Chen, Y., Ho, C ., & Waggoner, B . (2015). Low-cost learning via active data procurement. *The Sixteenth ACM Conference on Economics and Computation* (pp. 619-636). New York: ACM.
- Abernethy, J.**, & Johnson-Roberson, M. (2015). Financialized methods for market-based multi-sensor fusion. *IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 900-907). Hamburg, Germany: IEEE.
- Abernethy, J.**, Lahaie, S., & Telgarsky, M. (2015). Price Discovery in Subgradient Combinatorial Auctions. *The Third Conference on Auctions, Market Mechanisms and Their Applications* . Chicago, Illinois.
- Waggoner, B ., Frongillo, R., & **Abernethy, J.** (2015). A Market Framework for Eliciting Private Data. *Advances in Neural Information Processing Systems* 28 (NIPS 2015) . Montreal, CA.
- Abernethy, J.**, Lee, C. , Tewari, A. (2015). Fighting Bandits with a New Kind of Smoothness . *Advances in Neural Information Processing Systems* 28 (NIPS 2015) . Montreal, CA.
- Wang, Q., Sheng, S., **Abernethy, J.**, & Liu, M. (2014). Jamming Defense Against a Resource-Replenishing Adversary in Multi-channel Wireless Systems. *The 12th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks* . Hammamet, Tunisia: IEEE.
- Abernethy, J.**, Kutty, S., Lahaie, S., & Sami, R. (2014). Information aggregation in exponential family markets. *The Fifteenth ACM Conference on Economics and Computation* (pp. 395-412). New York: ACM.
- Abernethy, J.**, Frongillo, R., Li, X., Wotman Vaughan, J. (2014). A general volume-parameterized market making framework. *The Fifteenth ACM Conference on Economics and Computation* (pp. 413-430). New York: ACM.
- Abernethy, J.**, Lee, C., Sinha, A., & Tewari, A. (2014). Online Linear Optimization via Smoothing. *The 27 th annual Conference on Learning Theory* June 13-15. Barcelona, Spain.
- Abernethy, J.**, Bartlett, P., Frongillo, R., & Wibisono, A. (2013). How to Hedge an Option Against an Adversary: Black-Scholes Pricing is Minimax Optimal. *Neural Information Processing Systems* 2013 (pp. 2346-2354). Lake Tahoe, Nevada.

Abernethy, J., Kale, S. (2013). Adaptive Market Making via Online Learning. Neural Information Processing Systems 2013 (pp. 2058-2066). Lake Tahoe, Nevada.

McMahan, B., **Abernethy, J. (2013).** Minimax optimal algorithms for unconstrained linear optimization . Neural Information Processing Systems 2013 (pp. 2724-2732). Lake Tahoe, Nevada.

Abernethy, J. D., Amin, K., Draief, M., & Kearns, M. (2013). Large-scale bandit problems and KWIK learning . Proceedings of the 30th International Conference on Machine Learning, PMLR 28 , 588–596.

Abernethy, J., Frongillo, R., & Wibisono, A. (2012). Minimax option pricing meets black-scholes in the limit. The 44th Symposium on Theory of Computing Conference (pp. 1029-1040). New York, New York: ACM.

Abernethy, J., & Frongillo, R. (2012). A Characterization of Scoring Rules for Linear Properties. The 25 th Conference on Learning Theory (pp. 27.1-27.13). Edinburgh, Scotland.

Abernethy, J., & Frongillo, R. (2011). A Collaborative Mechanism for Crowdsourcing Prediction Problems. Neural Information Processing Systems 2011 (pp. 2600-2608). Granada, Spain.

Abernethy, J., Chen, Y., & Wortman Vaughan, J. (2011). An optimization-based framework for automated market-making. The 12th ACM Conference on Electronic Commerce (pp. 297-306). San Jose, California.

Abernethy, J. D., Bartlett, P. L., & Hazan, E. (2011). Blackwell approachability and no-regret learning are equivalent . Proceedings of COLT 2011, PMLR 19 , 27–46.

Abernethy, J., & Warmuth, M. (2010). Repeated Games against Budgeted Adversaries. The 24th Annual Conference on Neural Information Processing Systems (pp. 1-9). Vancouver, British Columbia, Canada.

Abernethy, J., Bartlett, P., Buchbinder, N., & Stanton, I. (2010). A regularization approach to metrical task systems. Algorithmic Learning Theory, 21st International Conference (pp. 270-284). Canberra, Australia.

Abernethy, J., Agarwal, A., Bartlett, P., & Rakhlin, A. (2009). A Stochastic View of Optimal Regret through Minimax Duality. The 22nd Conference on Learning Theory. Montreal, Quebec, Canada.

Abernethy, J., Rakhlin, A. (2009). Beating the adaptive bandit with high probability. The 22nd Conference on Learning Theory . Montreal, Quebec, Canada.

Abernethy, J., Warmuth, M., & Yellin, J. (2008). When Random Play is Optimal Against an Adversary. The 21st Annual Conference on Learning Theory (pp. 437-446). Helsinki, Finland.

Abernethy, J., Hazan, E., Rakhlin, A. (2008). Competing in the Dark: An Efficient Algorithm for Bandit Linear Optimization. The 21st Annual Conference on Learning Theory (pp. 263-274). Helsinki, Finland.

Abernethy, J., Bartlett, P., Rakhlin, A., Tewari, A. (2008). Optimal Strategies and Minimax Lower Bounds for Online Convex Games. The 21st Annual Conference on Learning Theory (pp. 415-424). Helsinki, Finland.

Abernethy, J., Chapelle, O., & Castillo, C. (2008). Web spam identification through content and hyperlinks. The Fourth International Workshop on Adversarial Information Retrieval on the Web (pp. 41-44). New York, NY: ACM.

Rakhlin, A., **Abernethy, J., & Bartlett, P. (2007).** Online discovery of similarity mappings. Machine Learning, Proceedings of the Twenty-Fourth International Conference (pp. 767-774). New York: ACM.

Abernethy, J., Bartlett, P., & Rakhlin, A. (2007). Multitask Learning with Expert Advice. The 20th Annual Conference on Learning Theory (pp. 484-498). San Diego, CA.

Abernethy, J. D., Langford, J., & Warmuth, M. K. (2006). Continuous experts and the binning algorithm . Proceedings of COLT 2006 , 544–558. DOI: 10.1007/11776420_40.

Abernethy, J. D., & Mannor, S. (2011). Does an efficient calibrated forecasting strategy exist? Proceedings of the 24th Annual Conference on Learning Theory (COLT), 809–812.

Broader-Impact Writing

Abernethy, J. D., Candelon, F., Evgeniou, T., Gupta, A., & Lostanlen, Y. (2024). Bring human values to AI . Harvard Business Review, March–April 2024 .

Selected Preprints and Technical Reports

Havrilla, A., Hughes, E., Samvelyan, M., & **Abernethy, J. D. (2025).** SPARQ: Synthetic problem generation for reasoning via quality-diversity algorithms . arXiv preprint arXiv:2506.06499 .

Google DeepMind Gemini Team (**including Abernethy, J.**). (2025). Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities . arXiv technical report arXiv:2507.06261 .

Kodali, N., **Abernethy, J. D.**, Hays, J., & Kira, Z. (2017). On convergence and stability of GANs . arXiv preprint arXiv:1705.07215 .

Extended Abstracts and Notes

Abernethy, J. D. (2010). Can we learn to gamble efficiently? Proceedings of the 23rd Annual Conference on Learning Theory (COLT), 318–319.

Abernethy, J. D., & Rakhlin, A. (2009). An efficient bandit algorithm for $\sqrt{T}$ regret in online multiclass prediction. Proceedings of the 22nd Annual Conference on Learning Theory (COLT).

Abernethy, J. D., & Warmuth, M. K. (2009). Minimax games with bandits. Proceedings of the 22nd Annual Conference on Learning Theory (COLT).

Selected Invited Talks

Google Paris — Building Algorithms by Playing Games	May 2019
Institut Henri Poincaré, Paris — Building Algorithms by Playing Games	May 2019
Universitat Pompeu Fabra Machine Learning Group — Building Algorithms by Playing Games	May 2019
Boston University CISE Seminar — Building Algorithms by Playing Games	April 2019
University of Southern California ML Seminar — Building Algorithms by Playing Games	March 2019
Data Learning and Inference Conference, South Africa — Building Algorithms by Playing Games	January 2019
Google AI for Social Good launch event — Data Science Remediation Efforts in the Flint Water Crisis	October 2018
NeurIPS Workshop on Learning in Games — Building Algorithms by Playing Games	December 2018
Google Zürich — Building Algorithms by Playing Games	July 2018
National Academies Workshop on Data Integration and Public Health — Data Science Remediation Efforts in the Flint Water Crisis	February 2018
Alan Turing Institute Workshop in Online Learning — Minimax Solutions, Random Playouts, Perturbations	October 2017
Georgia Tech ISyE Seminar — Simulated Annealing and Interior-Point Path Following	August 2017
NYU Machine Learning Seminar — Simulated Annealing and Interior-Point Path Following	April 2017
Conference on Big Data and Healthcare — Data Science Remediation Efforts in the Flint Water Crisis	May 2017

Professional Service

External Professional Service

- Program Co-Chair, Conference on Learning Theory (COLT), 2020.
- Senior Area Chair, Conference on Neural Information Processing Systems (NeurIPS), 2021, 2022, and 2026.
- Senior Area Chair, International Conference on Machine Learning (ICML), 2025 and 2026.
- Senior PC / Area Chair, ICML, 2015–2019.
- Senior PC / Area Chair, NeurIPS, 2013, 2014, and 2017–2019.
- Open Problems Chair, COLT, 2015 and 2019.
- Program Committee, COLT, 2015 and 2017–2019.
- Organizer, NeurIPS Workshop on Transactional Machine Learning and E-Commerce, 2014.
- Organizer, ICML Workshop on Markets, Mechanisms, and Multi-Agent Models, 2012.
- Panelist, National Science Foundation Small panel, February 2026.

Georgia Tech Service

- Faculty recruiting and graduate admissions, Georgia Tech.
- Faculty advisor, Data Science at Georgia Tech (DSGT).
- Faculty advisor, Georgia Tech Squash Club.