

Hongseok Namkoong

Columbia Business School
Kravis 933, 625 W 130th Street
New York, NY 10027

Email: namkoong@gsb.columbia.edu
Homepage: <http://hsnamkoong.github.io>
Google Scholar: <https://bit.ly/3hmDmjY>

Employment

Associate Professor, Decision, Risk, and Operations Division, Columbia Business School, 2026–Present
Associate Professor (Affiliate), Department of Computer Science, Columbia University, 2026–Present
Assistant Professor, Decision, Risk, and Operations Division, Columbia Business School, 2020–2026
LinkedIn Scholar, Trust & Responsible AI, LinkedIn, 2022–2025
Research Scientist, Core Data Science, Facebook, 2019–2020

Education

Ph.D. Management Science and Engineering, Stanford University, 2019
M.S. Statistics, Stanford University, 2017
B.S. Summa Cum Laude. Industrial Engineering and Mathematics, KAIST, 2013

Teaching

B8103: Business Analytics II (MBA, MS), Columbia University, 2021-2026
B9145: Topics in Trustworthy AI (PhD), Columbia University, 2020, 2023, 2025
B8784: Business and Society (MBA), Guest Instructor, Columbia University, 2024-26

Papers under review

Second+ round

- P1 E. Che and H. Namkoong. Adaptive experimentation at scale: A computational framework for flexible batches. *arXiv:2303.11582 [cs.LG]*, 2023. Under third round review in Operations Research.
- P2 S. Jeong and H. Namkoong. Assessing external validity to subpopulation shift via worst-case treatment effects. *arXiv:2007.02411 [stat.ML]*, 2022. Under third round review in Management Science. Conference version appeared in Conference on Learning Theory 2020.
- P3 N. Ye and H. Namkoong. Exchangeable sequence models quantify uncertainty over latent concepts. *arXiv:2408.03307 [stat.ML]*, 2024. Under second round review in Operations Research.
- P4 E. Che, D. Jiang, H. Namkoong, and J. Wang. A unified planning framework for short-horizon adaptive experimentation. *arXiv:2408.04570 [cs.LG]*, 2024. Under second round review in Management Science, Selected for oral presentations at the Econometric Society Interdisciplinary Frontiers: Economics and AI+ML conference and Conference on Digital Experimentation.
- P5 E. Che, J. Dong, and H. Namkoong. Differentiable discrete event simulation for queuing network control. *arXiv:2409.03740 [cs.LG]*, 2024. Major revision in Operations Research.

- P6 D. Mittal*, Y. Ma*, S. Joshi, and H. Namkoong. Adaptive measurement for posterior variance minimization. *arXiv:2502.06076 [cs.LG]*, 2025. Under second round review in Operations Research, Conference version appeared in NeurIPS 2024.
- P7 J. Liu, T. Wang, H. Lam, H. Namkoong, and J. Blanchet. DRO: A Python library for distributionally robust optimization in machine learning. *arXiv:2505.23565 [cs.LG]*, 2025. R&R at Journal of Machine Learning Research.
- P8 H. Namkoong. Empirical likelihood for nonsmooth functionals. *arXiv:2603.27743 [stat.ME]*, 2026. Major revision in Operations Research.

First round

- P9 D. Mittal, S. Zheng, J. Dong, and H. Namkoong. Data-driven stochastic modeling using autoregressive sequence models. *arxiv:2509.05839 [cs.LG]*, 2025.
- P10 T. Cai, H. Namkoong, D. Russo, and K. Zhang. Active exploration via autoregressive generation of missing data. *arXiv:2405.19466 [cs.LG]*, 2024. Selected for presentation at the Econometric Society Interdisciplinary Frontiers: Economics and AI+ML conference.
- P11 Y. Qu, H. Namkoong, and A. Zeevi. A broader view of thompson sampling. *arXiv:2510.07208 [cs.LG]*, 2026. Conference version appeared in ICML 2026.
- P12 J. Lee, H. Namkoong, and Y. Zeng. Design and scheduling of an AI-based queueing system. *arXiv:2406.06855 [math.OC]*, 2024. Selected for presentation at SIG Day 2025.
- P13 T. Cai*, Y. Fonseca*, K. Hou, and H. Namkoong. Constrained learning for causal inference and semiparametric statistics. *arXiv:2405.09493 [stat.ML]*, 2024.
- P14 J. Dong, D. Mittal, and H. Namkoong. Empirical rigor through benchmarking in operations research. *SSRN preprint*, 2026.
- P15 H. Namkoong. Local sensitivity under transport restrictions. *arXiv:2606.04276 [math.ST]*, 2026.

ML conferences

- P16 T. Yen, J. Poeltl, H. S. Gear, Y. Meng, J. Fan, A. Shen, Y. Liu, A. Bauyrzhan, S. Du, H. Liu, D. Guetta, and H. Namkoong. MBABench: Evaluating LLM agents on end-to-end spreadsheet tasks in finance. *arXiv:2605.22664 [cs.AI]*, 2026.
- P17 D. Mittal, T. Castellani, T. Yen, N. Ye, F. Wu, M. Chen, T. Cai, E. Koukoumidis, W. Zeng, and H. Namkoong. LatentGym: A testbed for cross-task experiential learning with controllable latent structure. *arXiv:2606.15306 [cs.LG]*, 2026.
- P18 G. Yang, Y. Chen, T. Yen, and H. Namkoong. Benchmarking in-context experiential learning through repeated product recommendations. *arXiv:2511.22130 [cs.LG]*, 2025.
- P19 T. Castellani, N. Ye, D. Mittal, T. Yen, and H. Namkoong. SynthTools: A framework for scaling synthetic tools for agent development. *arXiv:2511.09572 [cs.AI]*, 2025.

Working papers not yet submitted

- P20 H. Li, H. Namkoong, and I. Scheinfeld. When representative samples produce worse outcomes: Scale-up decisions and testing in small-budget RCTs. *arXiv:2606.13531 [stat.ME]*, 2026.
- P21 J. Wang, E. Che, D. Jiang, and H. Namkoong. AExGym: Benchmarks and environments for adaptive experimentation. *arXiv:2408.04531 [cs.LG]*, 2024.

P22 H. Namkoong, Y. Xu, and S. Zheng. Triply robust causal estimation for continuous treatments. *Working Paper*, 2025.

Publications

Journal publications

- J1 J. Liu*, T. Wang*, P. Cui, and H. Namkoong. Rethinking distribution shifts: Empirical analysis and inductive modeling for tabular data. *Management Science*, 2026.
- J2 H. Namkoong, S. Daulton, and E. Bakshy. Distilled Thompson sampling: Practical and efficient Thompson sampling via imitation learning. *Transactions on Machine Learning Research*, 2026. Selected for an oral presentation at the NeurIPS 2020 OfflineRL Workshop.
- J3 M. Li, D. Mittal, H. Namkoong, and S. Xia. Evaluating model performance under worst-case subpopulations. *Mathematics of Operations Research*, 2026. Short version appeared in NeurIPS 2021.
- J4 T. Cai, H. Namkoong, and S. Yadlowsky. Diagnosing model performance under distribution shift. *Operations Research*, 2025. Conference version appeared Symposium on Foundations of Responsible Computing 2023.
- J5 S. Joshi, I. Urteaga, W. A. C. van Amsterdam, G. Hripcsak, P. Elias, B. Recht, N. Elhadad, J. Fackler, M. P. Sendak, J. Wiens, K. Deshpande, Y. Wald, M. Fiterau, Z. Lipton, D. Malinsky, M. Nayan, H. Namkoong, S. Park, J. E. Vogt, and R. Ranganath. AI as an intervention: improving clinical outcomes relies on a causal approach to AI development and validation. *Journal of the American Medical Informatics Association*, 32(3):589–594, 01 2025.
- J6 H. Namkoong*, Y. Ma*, and P. W. Glynn. Minimax optimal estimation of stability under distribution shift. *Operations Research*, 74(1):464–483, 2025.
- J7 S. Yadlowsky, H. Namkoong, S. Basu, J. Duchi, and L. Tian. Bounds on the conditional and average treatment effect with unobserved confounding factors. *Annals of Statistics*, 50(5):2587–2615, 2022.
- J8 J. C. Duchi, T. Hashimoto, and H. Namkoong. Distributionally robust losses against mixture covariate shifts. *Operations Research*, 2022.
- J9 J. C. Duchi and H. Namkoong. Learning models with uniform performance via distributionally robust optimization. *Annals of Statistics*, 49(3):1378–1406, 2021.
- J10 J. C. Duchi, P. W. Glynn, and H. Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 46(3):946–969, 2021.
- J11 J. C. Duchi and H. Namkoong. Variance-based regularization with convex objectives. *Journal of Machine Learning Research*, 2019.

Refereed conference proceedings¹

- C1 Y. Kobayashi, Z. Jing, J. Yao, H. Namkoong, and S. Joshi. Learning-to-measure: In-context active feature acquisition. In *Proceedings of the 42nd International Conference on Machine Learning*, 2026.
- C2 Y. Qu, H. Namkoong, and A. Zeevi. A broader view of thompson sampling. In *Proceedings of the 42nd International Conference on Machine Learning*, 2026.

¹Papers displayed in gray are superseded by long versions.

- C3 L. Hu, J. Jiao, J. Liu, Y. Ren, Z. Wen, K. Zhang, X. Zhang, X. Gao, T. He, F. Hu, Y. Liao, Z. Wang, C. Yang, Q. Yang, M. Yin, Z. Zeng, G. Zhang, X. Zhang, X. Zhao, Z. Zhu, H. Namkoong, W. Huang, and Y. Tang. FinSearchComp: Towards a realistic, expert-level evaluation of financial search and reasoning. In *Proceedings of the Fourteenth International Conference on Learning Representations*, 2026.
- C4 A. Li, H. Chen, H. Namkoong, and T. Peng. LLM generated persona is a promise with a catch. In *Advances in Neural Information Processing Systems 38*, 2025.
- C5 T. Yen*, A. W. T. Siah*, H. Chen, T. Peng, D. Guetta, and H. Namkoong. Data mixture optimization: A multi-fidelity multi-scale bayesian framework. In *Advances in Neural Information Processing Systems 38*, 2025.
- C6 D. Mittal*, A. Li*, T.-C. Yen*, D. Guetta, and H. Namkoong. Architectural and inferential inductive biases for exchangeable sequence modeling. In *Advances in Neural Information Processing Systems 38*, 2025.
- C7 K. Zhang, T. Cai, H. Namkoong, and D. Russo. Contextual Thompson sampling via generation of missing data. In *Advances in Neural Information Processing Systems 38*, 2025.
- C8 J. Wang, T. Zollo, R. Zemel, and H. Namkoong. Adaptive elicitation of latent information using natural language. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025. Best Paper Award at ICLR 2025 Workshop on Quantify Uncertainty and Hallucination in Foundation Models.
- C9 T. Zollo*, A. Siah*, N. Ye, A. Li, and H. Namkoong. PersonalLLM: Tailoring LLMs to individual preferences. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2024.
- C10 B. Hsu, C. DiCiccio, N. Pillai, and H. Namkoong. From models to systems: A comprehensive framework for AI system fairness in compositional recommender systems. In *Proceedings of the Algorithmic Fairness Through the Lens of Metrics and Evaluation*, volume 279 of *Proceedings of Machine Learning Research*, pages 8–37. PMLR, 2025.
- C11 D. Mittal*, Y. Ma*, S. Joshi, and H. Namkoong. Adaptive labeling for efficient out-of-distribution model evaluation. In *Advances in Neural Information Processing Systems 37*, 2024.
- C12 H. Chen, A. Li, E. Che, T. Peng, J. Dong, and H. Namkoong. QGym: Scalable simulation and benchmarking of queuing network controllers. In *Advances in Neural Information Processing Systems 37, Datasets and Benchmark Track*, 2024.
- C13 J. Liu*, T. Wang*, P. Cui, and H. Namkoong. On the need for a language describing distribution shifts: Illustrations on tabular datasets. In *Advances in Neural Information Processing Systems 36*, 2023.
- C14 A. Boyarsky, H. Namkoong, and J. Pouget-Abadie. Modeling interference via experiment rollout. In *Proceedings of the 24th ACM conference on Economics and Computation*, 2023.
- C15 T. Cai, H. Namkoong, and S. Yadlowsky. Diagnosing model performance under distribution shift. In *Symposium on Foundations of Responsible Computing*, 2023.
- C16 M. Wortsman, G. Ilharco, S. Y. Gadre, R. Roelofs, R. Gontijo-Lopes, A. S. Morcos, H. Namkoong, A. Farhadi, Y. Carmon, S. Kornblith, and L. Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *Proceedings of the 39th International Conference on Machine Learning*, 2022.

- C17 M. Wortsman, G. Ilharco, J. W. Kim, M. Li, S. Kornblith, R. Roelofs, R. Gontijo-Lopes, H. Hajishirzi, A. Farhadi, H. Namkoong, and L. Schmidt. Robust fine-tuning of zero-shot models. In *Proceedings of the 32nd IEEE Conference on Computer Vision and Pattern Recognition*, 2022. Selected for a full oral presentation.
- C18 M. Li, H. Namkoong, and S. Xia. Evaluating model performance under worst-case subpopulations. In *Advances in Neural Information Processing Systems 34*, 2021.
- C19 H. Namkoong*, R. Keramati*, S. Yadlowsky*, and E. Brunskill. Off-policy policy evaluation for sequential decisions under unobserved confounding. In *Advances in Neural Information Processing Systems 33*, 2020.
- C20 S. Jeong and H. Namkoong. Robust causal inference under covariate shift via worst-case subpopulation treatment effect. In *Conference on Learning Theory*, 2020.
- C21 M. O’Kelly*, A. Sinha*, H. Namkoong*, J. Duchi, and R. Tedrake. Scalable end-to-end autonomous vehicle testing via rare-event simulation. In *Advances in Neural Information Processing Systems 31*, 2018.
- C22 R. Volpi*, H. Namkoong*, J. Duchi, V. Murino, and S. Savarese. Generalizing to unseen domains via adversarial data augmentation. In *Advances in Neural Information Processing Systems 31*, 2018.
- C23 T. Hashimoto, M. Srivastava, H. Namkoong, and P. Liang. Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning*, 2018. Best Paper Runner-up Award.
- C24 A. Sinha*, H. Namkoong*, and J. Duchi. Certifiable distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018. Selected for a full oral presentation; 2% of submissions.
- C25 H. Namkoong and J. C. Duchi. Variance regularization with convex objectives. In *Advances in Neural Information Processing Systems 30*, 2017. Best Paper Award.
- C26 H. Namkoong, A. Sinha, S. Yadlowsky, and J. C. Duchi. Adaptive sampling probabilities for non-smooth optimization. In *International Conference on Machine Learning*, pages 2574–2583, 2017.
- C27 H. Namkoong and J. C. Duchi. Stochastic gradient methods for distributionally robust optimization with f -divergences. In *Advances in Neural Information Processing Systems 29*, 2016.

Technical reports & software

- T1 Y. Ma and H. Namkoong. A sensitivity approach to causal inference under limited overlap. *arxiv:2511.22003 [stat.ML]*, 2025.
- T2 Y. Zeng*, J. Liu*, H. Lam, and H. Namkoong. LLM embeddings improve test-time adaptation to tabular $Y|X$ -shifts. *arXiv:2410.07395 [cs.LG]*, 2024.
- T3 G. Ilharco, M. Wortsman, N. Carlini, R. Taori, A. Dave, V. Shankar, H. Namkoong, J. Miller, H. Hajishirzi, A. Farhadi, and L. Schmidt. Openclip: Open-sourced implementation of clip, Jul 2021. URL <https://doi.org/10.5281/zenodo.5143773>.
- T4 A. Sinha*, H. Namkoong*, R. Volpi, and J. Duchi. Certifying some distributional robustness with principled adversarial training. *Technical Report*, 2020.

Honors & Awards

“Statistics of Robust Optimization: A Generalized Empirical Likelihood Approach” featured in Mathematics of Operations Research’s 50th anniversary as its selection of 50 publications in its history

Best Paper Award at ICLR 2025 Workshop on Quantify Uncertainty and Hallucination in Foundation Models

Interdisciplinary Seed Grant Program Competition Award, Columbia University, 2024

Amazon Research Award, 2022

Best Paper Finalist for “Robust fine-tuning of zero-shot models”, *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022

Best Student Paper Award for “Statistics of Robust Optimization: A Generalized Empirical Likelihood Approach”, *INFORMS Applied Probability Society*, 2018

Best Paper Runner Up Award for “Fairness Without Demographics in Repeated Loss Minimization”, *International Conference on Machine Learning (ICML)*, 2018

Best Paper Award for “Variance Based Regularization with Convex Objectives”, *Neural Information Processing Systems (NeurIPS)*, 2017

Samsung Fellowship, 2013–2018

Department Fellowship, Management Science and Engineering, Stanford, 2013–2018

KAIST President’s Award (graduated top of class in the School of Engineering), 2013

Undergraduate Research Award, First Place, Department of Industrial and Systems Engineering, 2012

Professional Service

Editorial Service

2026–Present	Associate Editor, Operations Research
2021–Present	Area chair, NeurIPS
2023–Present	Area chair, ICML

Reviewing

Journals *Operations Research, Management Science, Journal of the American Statistical Association, Journal of the Royal Statistical Society: Series B, Mathematical Programming, SIAM Journal on Mathematics of Data Science, Journal of Machine Learning Research, Transactions on Pattern Analysis and Machine Intelligence, Automatica.*

Conferences *Neural Information Processing Systems, International Conference on Machine Learning, Conference on Learning Theory, Conference on Algorithmic Learning Theory, MSOM Special Interest Group*

Workshop Organization

- 2026 Co-organizer with Rich Zemel, Workshop on “Continual Learning,” AI in Business, Columbia Business School
- 2025 Co-organizer, three INFORMS invited sessions on AI-driven decision-making
- 2025 Organizer, Digital Future Initiative workshop on “Agentic AI in Corporate Finance”
- 2025 Co-organizer, Columbia Catalyst Workshop, Data Science Institute
- 2024 Co-organizer, ACIC Invited Session on “External Validity”
- 2023 Co-organizer, NeurIPS workshop on “Distribution shifts (DistShift)”
- 2023 Co-Organizer, three INFORMS invited sessions on experimentation
- 2023 Organizer, Digital Future Initiative workshop on “Challenges in Operationalizing Responsible AI”
- 2023 Organizer, INFORMS APS invited session “Frontiers in Sequential Learning”
- 2022 Co-organizer, NeurIPS workshop on “Distribution shifts: connecting methods and applications (DistShift)”
- 2022 Organizer, INFORMS invited session “Causality & Robustness”
- 2022 Co-organizer, INFORMS invited session “Experimentation Design”
- 2022 INFORMS HAS Pierskalla Award committee member
- 2022 Co-organizer, ICDM workshop on “Algorithms Towards Ethical and Privacy Challenges in Social Media Recommendation Systems (AESM)”
- 2021 Organizer, mentoring program for PhD students at the NeurIPS workshop on “Distribution shifts: connecting methods and applications (DistShift)”
- 2021 Co-organizer, NeurIPS workshop on “Distribution shifts: connecting methods and applications (DistShift)”
- 2021 Co-organizer, JSM invited session on “distributional robustness, validity, causality, and generalizability”
- 2019 Co-organizer, INFORMS invited session on “AI and machine learning”

Academic Committees

- AI Vision Group, Columbia Business School, 2026–Present
- Ph.D. Coordinator, Columbia University, 2025–2026
- Member, Faculty Hiring Committee, Columbia University, 2021-22, 2023-24
- Member, Ph.D. Program Review Committee, Columbia University, 2021
- Member, Ph.D. Admissions Committee, Columbia University, 2020–2025

Public Speaking For Columbia

- “AI progress: opportunities, limitations, and what comes next,” Deming Center Board Meeting, 2026
- AI session, EY Chief Risk Officer Program, Columbia Business School
- AI Lab, Value Investing Program, 2026
- Guest Instructor with Jeff McMillan (Head of Firmwide AI, Morgan Stanley), Business and Society, 2026
- “OIN Workshop on Generative AI and Productivity”, Deming Center and Business in AI, 2025
- “Agentic AI and the future of work”, CBS Women’s Circle Alumnae Event, 2025
- “Teaching and research in the age of agentic AI”, CBS Board Meeting, 2025

AI Agent Workshop, CBS Alumni Club of New York, 2025

Panelist, CBS Alumni Club of New York, 2025

Guest Instructor with Lan Guan (CAIO at Accenture), Business and Society, 2025

“AI: Opportunities and Risks”, CBS Alumni Event With Dean Maglaras, 2025

“AI: Opportunities and Risks”, Chazen Institute for Global Business Board Meeting, 2024

Guest Instructor with Scott Silver (VP at YouTube), Business and Society, 2024

Panelist, Showcase on Large Language Models, Columbia Center of AI Technology, 2023

Moderator, Responsible AI and Content Moderation Bernstein Center for Leadership and Ethics 2023

External Service

NSF panel, 2024

Ad hoc reviewer, Methodology, Measurement, and Statistics Program, NSF, 2024

Research Supervision

Ph.D. Advising

Daksh Mittal (DRO, co-advised with Jing Dong), Naimeng Ye (DRO, co-advised with Tianyi Peng), Thomson Yen (DRO), Isaac Scheinfeld (DRO, co-advised with Hannah Li), Tommaso Castellani (CS), Max Bennett (CS), Hanming Yang (DRO)

Alumni Ari Boyarsky (Ph.D., 2026, at Cubist Systematic Strategies), Tiffany Cai (Ph.D., 2026, at Hudson River Trading), Yuanzhe Ma (Ph.D., 2026, co-advised with Garud Iyengar and Jay Sethuraman, at Meta), Ethan Che (Ph.D., 2025, co-advised with Jing Dong, at Meta Adaptive Experimentation), Yibo Zeng (Ph.D., 2024, co-advised with Henry Lam, at Meta Virtual Reality Labs)

Research Scientists

Minghui Chen

Postdocs

Fangyu Wu (DFI AI Agents Initiative), Qian Xie (DFI AI Agents Initiative)

Alumni Yanlin Qu (postdoc, 2024-26, co-advised with Assaf Zeevi, Assistant Professor at CUHK Shenzhen), Kelly Zhang (postdoc, 2023-24, co-advised with Dan Russo, Assistant Professor at Imperial College London)

Undergraduates and MS students

Gilbert Yang (Columbia CS), Joshua Fan (Columbia CS)

Alumni Tony Chen (Ph.D. candidate at Princeton CS), Elise Han (Ph.D. candidate at Stanford CS), Leon Li (Ph.D. candidate at NYU CS), Jimmy Wang (Ph.D. candidate at Berkeley CS), Shunri Zheng (Ph.D. candidate at UIUC IE), Yaqin Chen (Ph.D. candidate at Northwestern IEMS), Tommaso Castellani (Ph.D. candidate at Columbia CS)

Ph.D. Thesis Committee

David Cheik (DRO), Samuel Deng (CS), Nicholas Galbraith (Statistics), Yuanling (Judy) Gan (DRO), Yian Huang (Statistics), Luofeng Liao (IEOR), Yookoon Park (CS), Sudeep Raja Putta (IEOR), Claudia Shi (CS), Madhu Sridharan (IEOR), Tianyu Wang (IEOR), Shangzhou (Shawn) Xia (DRO), Yunbei Xu (DRO), Junzhe Zhang (CS)

Invited Talks

2026 OIT Division, Stanford Graduate School of Business
2026 Operations Management, Ross School of Business, University of Michigan
2026 Decision Sciences, Fuqua School of Business, Duke University
2026 Graduate School of Management and Technology, Korea University
2026 New Horizons on Model Transportability and Data Integration, Institute for Mathematical and Statistical Innovation, University of Chicago
2026 Workshop on External Validity in the Social Sciences, Yale University, RI
2026 Department of Operations Research and Information Engineering, Cornell University
2025 Panelist at NeurIPS 2025 Workshop on “ML x OR: Mathematical Foundations and Operational Integration of Machine Learning for Uncertainty-Aware Decision-Making”
2025 Keynote at Conference on Digital Experimentation, Cambridge, MA
2025 Operations department, Kellogg School of Business, Northwestern University
2025 INFORMS Annual Meeting 2025, Atlanta, GA
2025 Data Science Institute, Columbia University
2025 Keynote at Symposium on Mathematical Foundations of Trustworthy Learning, Switzerland
2025 BIRS-CMO Workshop on Adversarial ML, Oaxaca, Mexico
2025 Rotman School of Management, University of Toronto
2025 INFORMS Applied Probability Society Conference, Atlanta, GA
2025 SIG Day (Service Management), London Business School
2025 Workshop in Learning, London Business School
2025 Adaptive Experimentation Workshop, Meta, NYC
2025 Department of Industrial Engineering and Operations Research, UC Berkeley
2025 Operations and Logistics Division, University of British Columbia
2024 Simons Institute Workshop on Domain Adaptation and Related Areas, UC Berkeley
2024 Methods Seminar, Department of Political Science, UC Berkeley
2024 Central Applied Sciences, Meta
2024 TOPS Department Seminar, Stern School of Business, New York University
2024 INFORMS Annual Meeting, Seattle, WA
2024 Econometric Society Interdisciplinary Frontiers: Economics and AI+ML, Cornell University
2024 Keynote at Deployable RL Workshop, RL Conference, Amherst, MA
2024 International Symposium on Mathematical Programming, Montreal
2024 INFORMS Optimization Society Conference, Houston, TX
2024 Workshop on Optimal Transport: Theory to Applications, Humboldt University, Germany
2023 IMS International Conference on Statistics and Data Science, Lisbon, Portugal
2023 NeurIPS Tutorial “Modeling and Exploiting Data Heterogeneity under Distribution Shifts”, New Orleans, LA
2023 Stochastic Networks and Applied Probability, and Performance Seminar
2023 Management Sciences and Information Systems, Rutgers Business School
2023 Invited Session, INFORMS Annual Meeting, Phoenix, AZ
2023 Department of Industrial Engineering and Management Science, Northwestern University
2023 Big Data and Machine Learning in Econometrics, Finance, and Statistics, University of Chicago
2023 IBM Watson, NY

2023 Joint Statistical Meetings, Toronto
 2023 New England Statistics Symposium
 2023 INFORMS Revenue Management & Pricing Conference, London, UK
 2023 INFORMS Applied Probability Society Conference, Nancy, France
 2023 Department of Management Science and Engineering, Stanford University
 2023 AI Matrix Seminar, UT San Antonio
 2023 Operations Research Center, MIT
 2022 Conference on Digital Experimentation, Cambridge
 2022 INFORMS Annual Meeting, Indianapolis
 2022 IFDS Workshop on Distributional Robustness in Data Science, University of Washington
 2022 International Conference on Continuous Optimization (Bethlehem, PA)
 2022 Department of Industrial & Systems Engineering and Computer Science, KAIST
 2022 College of Business and AI, KAIST
 2022 International Conference on Econometrics and Statistics, Kyoto
 2022 Data Science Lab, MIT
 2022 Department of Economics, Columbia University
 2022 Workshop on Foundations of Stable, Generalizable and Transferable Statistical Learning, Mathematical Sciences Research Institute
 2021 Department of Statistics, Columbia University
 2021 LinkedIn
 2021 Workshop on “distributional robustness, validity, causality, and generalizability”, Joint Statistical Meetings
 2021 Empirical Inference Department, Max Planck Institute for Intelligent Systems
 2021 Department of Mathematics, KAIST
 2021 School of Data Science, Seoul National University
 2021 Data Science Institute, Columbia University
 2021 Decision Science Group, McCombs School of Business, UT Austin
 2020 Samsung Advanced Institute of Technology, Seoul
 2020 Google Brain, Cambridge
 2020 Conference on Information Sciences and Systems
 2020 American Causal Inference Conference
 2020 SIAM Conference on Mathematics of Data Science
 2019 Uber Marketplace and Uber Eats, San Francisco
 2019 OIT Division, Graduate School of Business, Stanford University
 2019 INFORMS Annual Meeting, Seattle
 2019 Stitchfix, San Francisco
 2019 Department of Computer Science, University of Wisconsin-Madison
 2019 Department of Industrial and Systems Engineering, University of Wisconsin-Madison
 2019 School of Operations Research and Industrial Engineering, Cornell Tech
 2019 Machine Learning and Statistics Group, Microsoft Research New England
 2019 Operations and Statistics Group, MIT Sloan School of Management
 2019 Department of Operations Research and Industrial Engineering, UT Austin
 2019 Machine Learning Department, Carnegie Mellon University
 2019 Heinz College, Carnegie Mellon University
 2019 Department of Industrial Engineering and Operations Research, Columbia University
 2019 Decisions, Risk and Operations Division, Columbia Business School
 2019 Department of Electrical and Computer Engineering, Purdue University
 2019 Operations Management Division, Booth School of Business, University of Chicago
 2019 Data Sciences and Operations, Marshall School of Business, University of Southern California
 2018 Department of Industrial and Operations Engineering, University of Michigan
 2018 Three invited talks, INFORMS Annual Meeting (Phoenix, AZ)

2018 Oral Presentation, International Conference on Learning Representations (Vancouver, Canada)
2017 Oral Presentation, Neural Information Processing Systems (Long Beach, CA)
2016 Department of Industrial and Systems Engineering, KAIST
2016 Young Researchers Workshop, School of ORIE, Cornell University

Outside Activities – None