

Alan E. Mishler

New York, NY • 720-352-9671 • contact@alanmishler.com

Website: alanmishler.com • Google Scholar: bit.ly/gs_alanmishler • www.linkedin.com/in/alanmishler

Profile

ML/AI researcher and engineer with a record of shipping production systems across LLM optimization, adaptive experimentation, causal inference, sequential inference, and financial ML. Combines rigorous statistical foundations with hands-on experience building open-source LLMOps tools and deployed models and agents.

Education

08/2021	PhD in Statistics, Carnegie Mellon University
05/2017	MS in Statistics, Carnegie Mellon University
08/2016	BS in Mathematics, University of Maryland

Employment

08/25–06/26	TensorZero (leading open-source LLMOps platform), <i>Member of Technical Staff</i> - Designed and shipped algorithms for adaptive A/B/n testing, top-k arm identification, and adaptive evals with anytime-valid inference guarantees; developed supporting statistical theory - Built automated evaluation system: automated example labeling, active learning, LLM-judge feedback loops, adaptive stopping - Shipped supporting infrastructure: tooling for Autopilot (agentic LLM optimization system), durable execution, evaluation tooling
06/21–08/25	JP Morgan AI Research <i>Vice President / AI Research Lead</i> (02/24–08/25) · <i>Senior Research Associate</i> (06/21–02/24) - Shipped production ML models for pricing of consumer loans, mortgage-backed securities, other asset classes - Conducted basic and applied research across topics including adaptive experimentation, causal inference, algorithmic fairness, uncertainty quantification, and conformal inference for LLMs - Supervised multiple research interns; co-authored papers in venues including AAAI, NeurIPS, TMLR, CLeaR
05/19–08/19	Google , <i>Data Scientist Intern</i> Developed improved methodology to estimate ads lift/incrementality using combined experimental and observational data, with both empirical and theoretical results.
06/18–08/18	Google , <i>Data Scientist Intern</i> Built a machine learning pipeline to improve classification of optical character recognition (OCR) outputs.

Recent publications

2026	Alan Mishler . (2026). Even (very) noisy LLM evaluators are useful for improving AI agents. TensorZero blog post. https://www.tensorzero.com/blog/even-very-noisy-llm-evaluators-are-useful-for-improving-ai-agents/ .
2025	Alan Mishler . (2025). Bandits in your LLM Gateway: Improve LLM Applications Faster with Adaptive Experimentation (A/B Testing). TensorZero blog post. https://www.tensorzero.com/blog/bandits-in-your-llm-gateway/ .
2025	Alan Mishler , Mohsen Ghassemi, Niccolò Dalmaso, Luhao Zhang, Vamsi K. Potluru, Tucker Balch, Manuela Veloso. (2025). Auditing and enforcing conditional fairness via optimal transport. <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> . https://doi.org/10.1609/aaai.v39i16.33847 . [*Authors 1-4 contributed equally.]
2025	Harit Vishwakarma, Alan Mishler , Thomas Cook, Niccolò Dalmaso, Natraj Raman, Sumitra Ganesh. (2024). Prune 'n Predict: Optimizing LLM Decision-making with Conformal Prediction. Poster presented at the ICML Workshop on Reliable and Responsible Foundation Models. https://arxiv.org/abs/2501.00555 .
2024	Thomas Cook, Alan Mishler , Aaditya Ramdas. (2024). Semiparametric efficient inference in adaptive experiments. <i>Proceedings of the Third Conference on Causal Learning and Reasoning</i> 236. https://proceedings.mlr.press/v236/cook24a.html .