

Kaushal Paneri

Senior Data & Applied Scientist · Microsoft

kaushalpaneri@gmail.com · LinkedIn · GitHub · Google Scholar

Senior Data & Applied Scientist blending research and production ML — 9+ years of software engineering and 7+ years specializing in large-scale applied ML.

- **6-year production ML stewardship** of a click-calibration stack at Microsoft, spanning multiple generations (Probit → Random Forest → ONNX → full MLOps on Azure ML)
- **End-to-end product ownership** of an auction-simulation platform, scaled from research prototype to **\$2M+ annual advertiser spend influence** over 3 years as primary owner
- **16 peer-reviewed publications** at NeurIPS, AAAI, IEEE Transactions on Big Data, and ACM RecSys (first-author oral at CONSEQUENCES 2024)

SKILLS

Research: Causal inference, reinforcement learning, foundation models, off-policy estimation, counterfactual inference, large-scale recommender systems, time-series forecasting, Bayesian methods.

Tools: Claude Code, Python, C#, PyTorch, scikit-learn, ONNX, Azure ML, MLOps, BOTorch, Git, GitHub Copilot, SQL.

EXPERIENCE

Senior Data & Applied Scientist · Microsoft Advertising

Sep 2024 – Present

- Own the forecasting models that drive automated-bidding decisions across the **entire marketplace** — a foundational product contribution powering advertiser spend decisions at scale.
- Drive bid-recommendation and performance-estimation pipelines at the core of the automated-bidding stack — influencing **~\$15M in annual advertiser spend**.
- Lead novel research combining **foundation-model forecasts** with **model predictive control** to make automated bidding more adaptive and efficient — targeting publication at a top ML venue.
- Established **CI/CD and security engineering** for the team's ML codebase — automated testing and validation on every change; identified and resolved critical dependency vulnerabilities that could have impacted production.

Data & Applied Scientist 2 · Microsoft Advertising

Sep 2020 – Aug 2024

- Published **SigIS (Simulator-Guided Importance Sampling)** as **first-author oral** at ACM RecSys CONSEQUENCES 2024 — a new method for tuning large-scale recommenders by combining simulation with importance sampling. Co-authored companion **MixtureIS** paper at the same venue.
- Owned and modernized the production **click-calibration ML stack for 4+ consecutive years**, evolving it across generations (Probit → Random Forest → ONNX → full MLOps) and designing a generalized feature-extraction architecture that became the platform-wide default.
- Launched and grew an **auction-simulation product line** from research prototype to production as primary owner — reached **\$1M+ annual advertiser spend influence** in the first 7 months post-launch.
- Brought **Causal Transfer Random Forest (CTRF)** from research paper to production click-prediction pipelines; co-presented the production integration at a Microsoft Research seminar with the original authors.
- Mentored **3 summer research interns** over 4 years — one project productionized into **\$400K annual savings**, another converted into the peer-reviewed **SigIS** paper.

Data & Applied Scientist · Microsoft Advertising

Sep 2019 – Sep 2020

- Delivered a stalled cross-team counterfactual-evaluation project within the first 3 months — ramping up a distributed big-data platform and a large C#/F# codebase from scratch.
- Built the **transfer-learning infrastructure** (local comparison environments, distributed training, hyperparameter optimization) that directly enabled the team's subsequent migration from the legacy click-prediction model to Random Forest in production.

Visiting Lecturer · Northeastern University, Khoury College of Computer Sciences 2022 – 2024

▸ Taught **Supervised Machine Learning** (graduate-level, 2 offerings) and **Data Analytics and Machine Learning II** (undergraduate).

Researcher · TCS Research Lab, Delhi 2014 – 2017

▸ Designed and built a **Bayesian Network Inference Engine** adopted across **20+ client organizations**.
▸ Published **7 peer-reviewed papers** (IEEE VIS, VAST, FUSION, ICDMW), won **4 international data challenges**, and produced research output that contributed to **2 granted US patents**.

OPEN SOURCE

crux — Agentic Tool Manager for Claude Code · github.com/crux-cli/crux 2025 – Present

▸ Authored and solely maintain an open-source Python CLI that lets developers orchestrate agentic workflows on top of **Claude Code** — with sandboxed tool execution, OS-keychain secrets management, and **MCP** (Model Context Protocol) server + skills integration.

PUBLICATIONS & PATENTS

1. System and Method for Visual Bayesian Data Fusion · **Patent US 10,430,417** 2019
2. Multi-Sensor Visual Analytics · **Patent US 10,013,634** 2018
3. Combining Open-box Simulation and Importance Sampling for Tuning Large-Scale Recommenders · 2024
CONSEQUENCES @ ACM RecSys (oral)
4. Adaptive Mixture Importance Sampling for Automated Ads Auction Tuning · **CONSEQUENCES @ ACM RecSys** 2024
5. Integrating Markov Processes with Structural Causal Modeling Enables Counterfactual Inference in Complex Systems · **NeurIPS** 2019
6. Regularizing Fully Convolutional Networks for Time Series Classification by Decorrelating Filters · **AAAI 2019** 2018
7. Speeding Up DNNs Using HPL Based Fine-Grained Tiling for Distributed Multi-GPU Training · **Boston Area Computer Architecture Research Workshop 2018** 2017
8. Visual Statistical Analysis of Environmental Sensor Data · **IEEE VAST 2017** 2016
9. Crop Planning Using Stochastic Visual Optimization · **IEEE VDS 2017** 2016
10. Visual Bayesian Fusion to Navigate a Data Lake · **IEEE FUSION 2016** 2015
11. Bayesian Visual Analysis of the Indian Labour Market · **CoDS Data Challenge** 2015
12. Multi-Sensor Visual Analytics Supported by Machine-Learning Models · **IEEE ICDMW 2015** 2014
13. Spatio-Temporal Analysis of Ethnic Conflicts and Human Rights Violations in Africa and the Middle East
14. Leveraging Structured Biological Knowledge for Counterfactual Inference: A Case Study of Viral Pathogenesis · **IEEE Transactions on Big Data (co-first author)** 2020
15. Integrating Markov Process and Structural Causal Models Enables Counterfactual Inference in Complex Systems · **Northeastern University MS thesis** 2019

PROFESSIONAL SERVICE & PEER REVIEW

Program Committee Member · CONSEQUENCES 2025 workshop @ ACM RecSys 2024

Applied Scientist Interviewer · Microsoft 2024 – Present

Technical Peer Reviewer · Microsoft, cross-team research reviews 2019 – Present

EDUCATION

MS, Data Science (Thesis Track) · Northeastern University 2017 – 2019

B.Tech., Computer Engineering · Charusat University, Gujarat, India 2010 – 2014