

Jonathan Gu

Senior Machine Learning Engineer · Causal Inference, Decision Systems, Agent Infrastructure

jonathangu@gmail.com (510) 333-5539 jonathangu.com LinkedIn GitHub

Senior ML engineer and PhD economist who builds production decision systems at scale. Owns Instacart's bidding and targeting systems controlling >60% of advertiser spend. Deep experience in causal inference, experimentation, reinforcement learning, and low-ops production ML. Outside of work, builds open-source agent memory infrastructure and a neighborhood sharing marketplace.

EXPERIENCE

Instacart

2020 – Present · San Francisco

Senior Machine Learning Engineer & Economist II

- ▶ Builds and owns the optimized bidding algorithm controlling >60% of advertiser spend and indirectly influencing ~20% of total ads revenue. Advertisers specify only objective + budget; the system allocates spend to maximize ROI using reinforcement learning with transparent, interpretable bid adjustments.
- ▶ Develops causal measurement models for coupon email campaigns, quantifying incremental lift per offer and automating budget reallocation toward higher-impact coupons across millions of customers.
- ▶ Designs and deploys an in-app marketing decision system enabling advertisers and internal teams to target customers within the Instacart app, solving overlapping-intervention and sparse-data challenges through rigorous experimentation and causal inference at scale.
- ▶ Leads LLM-assisted critical-item prioritization for fulfillment routing, predicting the business cost of missing items to optimize which warehouse fulfills each order — serving real-time decisions across millions of transactions with minimal on-call overhead.
- ▶ Builds and maintains the batch and real-time inventory-availability platform: production modeling, threshold calibration, and predictions across millions of SKUs with automated feature parity and traceability.

Microsoft Research New England

2013 – 2014 · Boston

Research Assistant (Prof. Susan Athey)

- ▶ Modeled counterfactual outcomes of Bing search auction mechanisms, identifying potential revenue manipulation strategies under alternative ranking formulas. Contributed to research on auction design and causal inference in marketplace settings.

SELECTED PROJECTS

Bountiful Garden · [bountiful.garden](#) · Live

Neighborhood produce-sharing marketplace. Post your harvest, browse what's ripe nearby, pick it up off the porch. Full-stack React app with location-based matching and real-time availability.

OpenClawBrain · [openclawbrain.ai](#) · Alpha

Open-source memory layer for AI agents. Replaces keyword retrieval with a learned routing function that maps knowledge as a graph, picks context based on relevance, and improves from corrections. Beats best RAG baseline by +9.7–26.9 pp across three benchmark families using 5× less context. TypeScript, published on npm.

Project Pelican · Private R&D

Autonomous options research and execution system. End-to-end ML pipeline: market data ingestion, feature engineering, model training, position sizing, and execution — no human in the loop.

PATENTS FILED (7)

Auction design for variable-size ads · Causal/incremental targeting via synthetic treatment effects · RL-based optimized bidding · Dynamic offer targeting by purchase behavior · Distribution-free confidence intervals via bootstrapping · LLM-based retailer classification from sales data · Fulfillment optimization via predicted item criticality

EDUCATION

PhD Economics — UCLA

2014 – 2020

BA Statistics & Economics — UC Berkeley

2007 – 2011

ML & Methods: Causal inference, reinforcement learning, experimentation/A/B testing, propensity modeling, LLM integration, optimization

Engineering: Python, TypeScript, Java, Scala, Go · TensorFlow, PyTorch · Airflow, SQL, large-scale data pipelines

Research: Corrected policy gradients for graph traversal (Gu 2016) · Auction mechanism design · Causal measurement