

ZIYIN FENG

New York City | [linkedin.com/in/ziyin-awesome](https://www.linkedin.com/in/ziyin-awesome)

SUMMARY

Machine Learning Engineer with 9+ years of experience building large-scale systems at the intersection of deep learning, optimization, and product. At Amazon, I build prediction models serving billions of ad requests across global marketplaces, architect real-time ML serving infrastructure, and have shipped systems from auction algorithms to streaming ad pipelines. Looking for roles where I can apply deep ML engineering and systems thinking to hard problems at scale.

EXPERIENCE

Machine Learning Engineer — Ads Ranking & Prediction (ARP)

Jul 2021 – Present

Amazon · Sponsored Products

New York, NY

- **Deep Learning & Model Consolidation** — Build and optimize deep learning models for click-through rate prediction serving Sponsored Products across all Amazon marketplaces. Led co-serve consolidation to unify models across regions, reducing infrastructure cost while improving prediction quality globally.
- **Streaming Architecture Migration** — Migrated the entire ad serving workflow from RESTful request-response to a streamlined reactive architecture, enabling multi-response streaming that reduced end-to-end latency and unlocked new ad formats requiring incremental rendering.
- **Tier 1 Auction Service (Mercury)** — Owned and operated Mercury, the Tier 1 service for Sponsored Products auction and allocation processing billions of daily ad requests at 99.99% uptime. Responsible for how ads are ranked and served across Amazon search and browse pages.
- **Ranking Algorithms & Experimentation** — Designed auction algorithms optimizing ad allocation under budget and relevance constraints, contributing to \$500M in annual revenue growth. Built an experimentation platform automating A/B test lifecycle end-to-end, cutting launch time from 2 weeks to 2 days and enabling 200+ org-wide experiments.
- **ML Serving Infrastructure** — Partnered with applied scientists to integrate model updates into real-time serving pipelines, implementing feature extraction, scoring interfaces, and candidate ranking at low-latency scale.

Full Stack Engineer

Jan 2017 – Sep 2021

EdLab, Teachers College Columbia University

New York, NY

- Led development of the EdLab website and a custom CMS serving the Teachers College community, working with researchers to ship educational technology products used by thousands of students and faculty.
- Built data-driven web applications with React, Redux, and Node.js; designed RESTful APIs and data pipelines supporting content management, user analytics, and research tools.
- Introduced modern frontend tooling, automated testing, and CI/CD practices that improved team development velocity and release confidence.

TECHNICAL SKILLS

Languages	Java, Python, JavaScript/TypeScript, SQL
ML & Data	Deep Learning, CTR Prediction, Model Training & Serving, Feature Engineering, Scoring Pipelines, A/B Testing
Systems	Distributed Systems, Reactive Streaming, Microservices, Real-time Ranking & Serving, Low-latency Pipelines
AWS	CDK, Lambda, DynamoDB, Step Functions, S3, CloudFront, CloudFormation, EC2, SQS
Frontend	React, Redux, Node.js, HTML/CSS
Tools	Git, CI/CD, Docker, Agile/Scrum, Monitoring & Observability

EDUCATION

Stevens Institute of Technology

2014 – 2016

M.S. Computer Science

Dalian University of Technology

2010 – 2014

B.E. Software Engineering