

Teran Long

ML Engineer - Fraud Detection

Core Technologies: Python, Spark, Scala, C#/Java, PySpark, ONNX, ML Applications, MLOps, LLM (RAG, agents), LangGraph, VectorDbs, Chroma, Postgres, Azure Cloud, Feature stores, SQL, Data Pipelines, Microservices, A/B Testing, FastAPI, Prefect

📞 +1 4257613859

✉ teran.long@gmail.com

LinkedIn linkedin.com/in/teranlong

KEY SKILLS

Production ML at global scale:

Experience with **pb-scale** multi-region data pipelines & **99.99% SLA** Azure services powering real-time mission-critical ML decisions

ML Platform Design:

Partnered with data scientists and product stakeholders to design and build low-friction experimentation-to-production ML platform

MLOps & Infrastructure:

Hands-on ownership of the full ML lifecycle including **ONNX** deployment, feature stores, observability & rollback for high-SLA workloads

Applied ML & model iteration:

Extended production models with new signals, enabled **XGBoost** transition, and led eval, calibration, evaluation via **shadow & A-B testing**

Leadership & cross team impact:

History of outsized impact for my level - Co-led ML modernization, mentored interns, introduced key technologies, and influenced long-term platform architecture

EDUCATION

University of Washington

BS in Computer Science – 2020

- Direct Admit CS Dept. (top 3%)
- 3.64 GPA. Early graduation.
- Focus on ML, DL and NLP.
- Founding officer of ML Society UW

Highline College

AS in Computer science - 2017

- (College in high school program)
- CS Dept. award winner at age 17 cis.highline.edu/2017-pas-award-recipients

GAP YEAR a personal note:

Hi- I graduated **two years early**, then jumped straight into **Microsoft**. After the pandemic, I took that time back to travel and invest in personal growth. **Now I'm back to the US and eager for my next challenge.**

Along the way I rode a **motorcycle 2,000 miles** across **Vietnam, trekked 120 miles** in the **Himalayas**, advanced my **Spanish A1→B1** in **Bolivia**, learned to **lead climb** in **Laos**, **scuba dive & hold my breath 4 mins** in **Indonesia**, and **surf** in the **Philippines**.

EXPERIENCE

Software Engineer II

Sept 2022 – Jan 2024

Microsoft – Identity Security and Protection

- **Co-led ML modernization initiative**, redesigning model management, delivery, and feature computation across **PB-scale, multi-region data**; designed and shipped v1 of a multi-region feature store backed by AzureML.
- **Introduced and implemented ONNX** to decouple training from ML.NET, enabling flexible model training and achieving **+5% recall, +2% precision, and 50% faster** time-to-production.
- **Drove cross-functional architecture decisions** for ML pipelines, improving build velocity and shortening model iteration cycles across the product group.
- **Led intern project II** to evaluate open-source feature stores; selected and prototyped **Feathr**, demonstrating business value and accelerating adoption.

Software Engineer I

April 2020 – Sept 2022

Microsoft – Identity Security and Protection

- **Subject Matter Expert** for training and productionization of online/offline identity risk models powering **50B daily sign-ins**.
- **Owned the end-to-end ML lifecycle**—training, inference, evaluation, deployment—improving precision/recall through signal integration and experimentation.
- Built standardized **validation, rollback, and live-site monitoring** pipelines, improving reliability, observability, and telemetry quality.
- **Led intern project I** to apply **active learning** within **Spark ML** pipelines, improving efficiency of expensive labeling workflows across distributed systems.

Internships:

Microsoft - 2019 | Starbucks - 2018 | Centurylink - 2017

PROJECTS

Agentic AI – Ongoing personal projects

- *(In progress)* Building an autonomous agent for card game strategy and trading; integrated multi-source data ingestion (RAG) and persistent memory and agentic decision making via LangGraph. Using FastApi, Pydantic, ChromaDb, Prefect, Postgres, Docker.

Adversarial Deep Learning – Deep Learning @UW

- applied transfer learning on a pretrained ImageNet classifier to **train adversarial images**, then **hardened the original classifier** against these attacks

Spatial Clustering – Microsoft Hackathon

- Microsoft Hackathon Project using **DBSCAN** algorithm to find optimal pick-up and drop-off points on campus.