

Obuli Krishnaraj Venkatesan

San Jose, CA | 408-623-4074 | obulikrishnaraj@gmail.com
github.com/obuliv | linkedin.com/in/obulikrishnaraj | shastraw.ai

Summary

Co-Founder of modelteam.ai and founding engineer at Whisp AI. Built AI products end-to-end across AI/ML, fullstack, and infrastructure. The last 3 years of hands-on AI work build on ~17 years of experience designing large-scale graph, data, and ML systems at LinkedIn, Drawbridge, and Yahoo. Most effective as an early engineer, shipping quickly with modern AI tooling.

Technical Skills

AI/ML: LLMs, Fine-tuning (HuggingFace, CodeT5+), AI Agents, Video Understanding, ML Pipelines

Languages: Python, Java, Scala, JavaScript/React

Systems & Data: Apache Spark, Apache Giraph, Graph Processing, Hadoop, MapReduce, HBase, Distributed Systems, Vector DBs

Infra & Tools: Google Cloud, Docker, Flask, React, Claude Code

Experience

Whisp AI - Founding Engineer

Jun 2025 - Apr 2026

- Sole fullstack and AI engineer at an angel-backed video-AI startup - owned product direction, UI (React), backend services, and cloud infrastructure.
- Built and architected the core long-video understanding pipeline enabling reliable Q&A over 1hr+ videos - owning system architecture, infrastructure, deployment, and data/DB schema design and setting overall technical direction for an agentic system.
- Automated quality evaluation by building a CI/CD benchmarking pipeline that scoped each run from Git tags - re-running agent evals or reprocessing the entire pipeline as needed - so every change was measured against benchmarks without manual steps or unnecessary recomputation.
- Built a semantic video search system from scratch (two-level vector retrieval + LLM reranker); internal benchmarks showed retrieval quality comparable to leading commercial tools at a fraction of the compute cost.

modelteam.ai - Co-Founder

Aug 2023 - Jun 2025

- Co-founded an AI platform that assessed software engineers by analyzing their real code contributions and Git history - a deliberate pivot into hands-on AI/ML after years in large-scale systems.
- Built and fine-tuned a 770M-parameter CodeT5+ model to classify skills from source code; owned the full lifecycle - dataset construction, labeling, taxonomy design, training, evaluation, and deployment.
- Reframed skill prediction as next-token generation over custom skill tokens so model confidence could serve as a code-quality signal.
- Caught train/test leakage from repository memorization and redesigned evaluation around repository-level holdouts (from user-level) for honest generalization. Built the full stack (Flask APIs, React, GCP); ran pilot with a design partner before winding down without product-market fit.

LinkedIn - Staff Software Engineer

Jun 2019 - May 2023

- Designed a privacy-preserving Group Identity solution for LinkedIn Marketing Solutions (the ads business) - converting the identity graph into group-level representations to keep ad targeting viable under Apple-consent and other privacy/legal changes.
- Partnered with legal and privacy researchers to balance utility with privacy constraints, preserving ad-targeting effectiveness that would otherwise have been lost.
- Built a graph-algorithms + ML module that doubled Identity Graph coverage at 80% precision.
- Led post-acquisition migration of Drawbridge's identity platform onto LinkedIn's infrastructure, enabling long-term operation within LinkedIn's ecosystem. Mentored the team through the transition.

Drawbridge (acquired by LinkedIn) - Architect

Jul 2012 - Jan 2018, Oct 2018 - Jun 2019

- Joined as one of the first ~10 engineers and stayed through the LinkedIn acquisition, ending as the most senior individual contributor on the engineering team - a full startup-to-exit arc.
- Drove a paradigm shift from MapReduce to graph processing for cross-device identity resolution - the company's core ad-tech product - at billions-of-nodes scale; introduced Apache Giraph and designed novel pairing/clustering algorithms.
- Cut core pipeline runtimes from days to hours (clustering ~2 days→<4h; ID-IP pairing ~3 days→<4h), enabling daily/weekly refreshes and improving match quality.

Yahoo! - Software Engineer

Jun 2006 - Jul 2012

- Tuned mobile web-search relevancy with ML and graph features; early Hadoop adopter and contributor to Apache Oozie, the workflow scheduler for Hadoop.

Patent

System to group internet devices based upon device usage - US 9,514,248

Selected Project

Graphedart · 2026

- Fully automated graph + ML trading system, built solo end-to-end using modern AI-assisted development tools. Models the Russell 3000 as a daily directed graph - price co-movement, news co-occurrence, and LLM-extracted SEC-filing relationships (supplier/customer/competitor) - ranked by gradient-boosted models plus a learning-to-rank reranker.
- Walk-forward training with weekly retrain-and-promote gating and rigorous time-boundary leakage prevention; runs unattended on automated paper trading.

Education

B.E. Computer Science & Engineering - Government College of Technology, Coimbatore, India