

IFECHUKWUDENI TEDDY OWEH

281-725-1576 | teddyoweh@gmail.com | linkedin.com/in/teddyoweh | github.com/teddyoweh | teddyoweh.net

EDUCATION

Morgan State University

B.S. Computer Science & Mathematics, Concentration in Artificial Intelligence · GPA 3.9

Expected May 2027

Baltimore, MD

WORK EXPERIENCE

Google

Software Engineer Intern · Python, TypeScript, Lit, Polymer, google3, Zipline, LLM Agents, MCP

May 2026 – Aug 2026

San Francisco, CA

- Built agents on Jetski, Google's internal agent harness, that prototype UI directly against the internal web stack (Lit, Polymer, google3).
- Built a component knowledge graph: AST-parsed the TypeScript component library to extract props, types, and composition edges, indexed as a snapshot and served to the agents through an MCP server they query at generation time.
- Worked with and deployed across the Workspace, Gemini, and Google Search teams: built the pipeline for agents to deploy prototypes to a Zipline server behind a live preview link, migrated the core component library onto google3's new framework, and added MCP servers to the harness.

LinkedIn

Software Engineer Intern, Identity · Python, PyTorch, Hadoop

May 2025 – Aug 2025

Sunnyvale, CA

- Fine-tuned LinkedIn's internal foundation model with a contrastive objective over positive profile pairs, extending its embeddings to job history, titles, skills, and certifications.
- Ran Hadoop batch jobs generating embeddings across 500M member profiles, producing the vectors behind semantic persona search.
- Built the benchmarking harness (MRR, NDCG, AUC, serving latency) and measured a ~10% ranking-quality lift over the production baseline at unchanged latency, across the member recommendation and job matching pipelines.

USC Information Sciences Institute (ISI)

Research, Data-Center-Scale ML Infrastructure · Python, Rust, C++, GPU Clusters

May 2024 – Sept 2024

Los Angeles, CA

- Built "Git for models", a version-controlled registry for ML model checkpoints: content-addressed storage under a commit DAG, with branching and diff/merge across model variants.
- Designed the storage engine: chunk-level deduplication and delta encoding over serialized weights, cutting checkpoint storage 73% across the registry.
- Built multi-node sync on the shared GPU cluster that ships only the chunks a node is missing, cutting 12 TB/month of transfer; adopted by 8 research teams across the ISI lab.

Apple Inc.

Machine Learning & Software Engineering Intern, Apple Home · Python, C, TypeScript, PyTorch, React, FastAPI

May 2023 – Sept 2023

Cupertino, CA

- Built a hybrid Python/C ML library (ensemble models cross-validated against VAE latents to handle multicollinearity), reaching 94.3% accuracy on Apple Home device crash and proximity prediction.
- Designed backend services and telemetry pipelines behind internal APIs, cutting p99 serving latency from 78 ms to 22 ms, with model-metric dashboards and A/B comparison surfaces on top in React + TypeScript.
- Shipped a multi-threaded REST API, an analytics web UI, and the MCQI failure-metric system; presented root-cause findings to Apple Home leadership and the org VP.

NASA Marshall Space Flight Center (USLI)

Research Payload Engineer, "Eyes in the Cloud" · Python, C++, SDR, React

Aug 2022 – May 2023

Huntsville, AL

- Built payload comms for a sub-scale rocket: RF telecommand, a bespoke SDR receiver, a custom UDP server for real-time flight analytics, and onboard image processing at 4,600 ft apogee; integrated a BMP388 + IMU stack with a React ground-station UI, validated through launch and recovery.

PROJECTS

Spawn Labs · TypeScript, Next.js, Python, Kubernetes, GKE, Docker, LLM Agents, MCP

spawnlabs.ai

- General-purpose agent platform that runs modern coding harnesses (Claude Code, Codex, and custom ones) as containerized workloads on a Kubernetes cluster on GKE, deployed across 41 businesses: agents that research, build, and deploy full-stack apps, drive a real browser, and run long-horizon jobs.
- Handles image build and per-run sandboxed execution, with persistent volumes, storage, and Kubernetes secrets for credential management.

herds.run · Python, Distributed Scheduling, macOS Sandboxing, TLS Networking

herds.run

- Turns idle Macs into a private fleet of programmable cloud runtimes, giving agents what Linux containers cannot (Xcode, iOS Simulator, codesigning, native Apple apps, real Safari): a work-stealing scheduler fans one prompt across the fleet, with up to 24 disposable sandboxes per Mac on isolated HOME, TMPDIR, and toolchains, each published on a named subdomain over TLS and zero inbound ports.

mantis-agent · Python, asyncio, MCP, OpenTelemetry

mantisagent.cc

- Custom agent harness in the vein of Claude Code and Codex, running one agent loop across frontier and open-source models (OpenAI, Claude, Gemini, Qwen, DeepSeek, Kimi), with one-click hosting of open models on inference providers; handles context, session persistence, MCP servers, streaming, usage tracking, and OpenTelemetry tracing. On PyPI under Apache 2.0.

TECHNICAL SKILLS

Languages: Python, TypeScript, C++, C, Rust, Swift, SQL, Assembly (x86, AVR)

ML: PyTorch, TensorFlow, scikit-learn, Pandas, NumPy

Agents: LLM agents, MCP, Claude Code, Codex, sandboxes, VMs

Infrastructure: Docker, Kubernetes, Modal, AWS (EC2/S3), PostgreSQL, MongoDB, Linux/Unix, Git

Web: React, Next.js, Lit, Polymer, React Native, Node.js, FastAPI, Flask, TailwindCSS