FF FRANKFOX

San Jose, CA · frank@frankfox.dev · linkedin.com/in/ffox77

Frank Fox · Applied AI Engineer

Building how coding agents and people work together: agents make more of the build decisions as they earn it, while people set the goals and decide when they change. Co-founded two AI startups since 2023, one sold.

Claude Code | Agent autonomy | Agent evals | Context engineering | Spec-driven development | TypeScript

Experience

Constellation Network · Principal Software Engineer, AI Agents

An AIAI Holdings company · Remote · Contract · 2025 – now

- Built a **spec-driven harness for AI coding agents** for the network's community: each stage starts with curated context, and a **learning loop** files what each change teaches and retires stale rules.
- Designed how coding agents **stop only when a person is needed**: each stop is tagged by type, and the harness sends implementation questions to an AI advisor and goal questions to the person.
- Built an **A/B rig** for agent skills: it runs old and new versions of a skill on the same task in headless Claude sessions, scores both with program checks and a blinded LLM judge, and records each run's cost. Cut skill text **30%** over about 50 runs; reverted the one cut the judge flagged.
- Built **Node Pilot**, a CLI to install, monitor and upgrade validator nodes. **Sole engineer**; **200+ operators**.
- When run with no arguments, it checks what's missing on the machine and the node, walks the operator through only those steps, and **asks before any update that takes the node down**. Ran community support for operators and turned their reports into fixes.

Cabana Labs · Technical Co-Founder

Digital identity studio, three founders · 2021 – 2026

- Led Cabana's **decentralized-identity** work: identity profiles whose owners control who can read them (**Lit Protocol** grant), a confidential wallet (**Oasis** grant) and a DID registry (Alephium hackathon).
- In 2026, extended the identity work to AI agents: a **trust score for AI agents**, tested in Python simulations of 16 attack scenarios. Sybil swarms (one actor, many fake identities) scored **5%** of an honest score; a parameter sweep made taking over other users' emails and phones unprofitable.

UserHunt AI · Technical Co-Founder

SF Bay Area · 2024

- Built a **paid AI user-research product**: a founder describes an idea, interviews a panel of AI "virtual users", and gets a written report. Stripe billing. The company was accepted into **Antler** (S24) and ran a pilot with **5 market-insight agencies**.
- Built the **interview screen** for talking to several AI personas at once: a tab per persona, a **typing indicator** on each one still answering, image questions and suggested questions.
- Ran each persona as its own **OpenAI Assistant**, asked in parallel; a job runner chained the multi-step AI work, passing typed JSON between steps, and showed **live progress** in the UI.

Experience · continued

Social Twin · Technical Co-Founder
SF Bay Area · 2023

- Built a **GPT-4 service** that drafts social posts in the user's own writing style and tone: paste a writing sample, pick a topic, get drafts. Prompt variants included two-step chains that analyze the style.
- **Sold the product** about 4 months after the first build.

Constellation Network · Principal Product Engineer
San Francisco, CA · 2020 – 2021

- **Sole engineer** on **dag4.js**, the network's TypeScript SDK: 7 npm packages (keys, wallets, network calls, Ledger signing), with **50k+ downloads a year**. The Stargazer wallet runs on it.
- Led the team that built the **Stargazer wallet**, including Ledger hardware-wallet support, and wrote the developer and exchange-integration docs.

TOTVS Labs · Senior Front End Architect, Team Lead
Mountain View, CA · 2015 – 2020

- **Front-end team lead** in a stealth enterprise R&D lab. Set the standards and workflows for a product suite covering data mapping, transforms, analytics dashboards and real-time data processing.

GlobalEnglish / Pearson English · Technical Lead, Team Lead
SF Bay Area · 2010 – 2015

- **Technical lead** on an e-learning platform; brought AngularJS, Cordova, Rx and React Native into the team's stack early and mentored engineers.

Earlier experience

- **Qualcomm** · Software Engineer: a real-time status and alarm monitor for the CDMA network, and a C++ upgrade tool used at **60+ sites**.
- **Ericsson** · Sr. Software Engineer: technical lead for UI frameworks.
- **Kratos Defense** · Lead Software Engineer: a framework others used for **70 e-learning simulators**.
- **SuddenlySmart** · Senior Software Engineer: lead client-side engineer on an e-learning content system that won the **2008 Brandon Hall gold award**.

Education

B.S. Computer Engineering
UC San Diego

Skills

Agent infrastructure: Claude Code plugins, skills and hooks; MCP servers; headless agent runs.

Agent autonomy: escalation and handoff rules, AI-advisor review, approval gates, decision logs.

Evals: built A/B rigs for prompts and skills, with blinded LLM judges; ran experiments and log analyses.

LLM SDKs: Claude, OpenAI, Gemini; prompt chains, typed output, background AI jobs.

Engineering: TypeScript, React, Node.js, NestJS, Firebase, Google Cloud, Docker, CLIs.