

# CHRISTOPHER WOOD

Software Engineer | San Francisco, CA

[christopherhwood.com](http://christopherhwood.com) | [github.com/christopherhwood](https://github.com/christopherhwood) | [linkedin.com/in/christopherhwood](https://linkedin.com/in/christopherhwood)

---

Software engineer with 10+ years of experience spanning systems engineering at Meta, applied AI engineering, and full-stack product development. After six years at Meta, spent two years on independent and contract work: shipped a TUI for an autonomous coding agent from concept to launch, built and open-sourced a coding agent SDK, and conducted independent research applying diffusion models to the ARC-AGI reasoning benchmark. Passionate about bridging AI research and production software.

## EXPERIENCE

### Founder / Software Engineer | qckfx

Jan 2025 - Present

- Open-sourced an AI coding agent SDK (TypeScript, predating Claude Code SDK) supporting local and remote sandboxed execution. Modular, OpenAI-compatible framework supporting Anthropic, OpenAI, Google, and local model providers.
- Built black-box iOS regression testing tool (Swift macOS app, Go proxy, XCUITest) that treats application code as opaque, recording touch inputs and network traffic from iOS simulators for deterministic e2e test replay without SDK access.

### Contract Software Engineer | Harvey (Legal AI)

2025 - 2026

- Joined pre-beta, built the design system and polished the UI for Harvey's iOS application. Fixed performance issues ahead of App Store launch.

### Contract Software Engineer | Random Labs

2025

- Built the terminal user interface (TUI) for Slate, the company's autonomous coding agent designed for long-running, highly autonomous development tasks.

### Independent Software Engineer | Various Projects

2024 - 2026

- AI PDF Extraction Platform: Full-stack app (Node.js, PostgreSQL, React) using multiple AI models to extract structured content from PDFs into editable tables with CSV export. Recursive self-improving prompt flow: user corrections fed back into the prompt generator to improve extraction accuracy over time.
- Ecommerce Image Generator: Deployed AI image generation using ComfyUI pipelines and LoRA fine-tunes of Flux diffusion models on Kubernetes. Used by real customers.
- Open-sourced: LLM-powered website builder, frontend component builder agent, Node.js coding agent (predating Devin).

### Software Engineer | Meta

2018 - 2024

- XROS (Team Lead): Led 5 engineers on a novel microkernel AR/VR operating system. Built the UI component framework, input management, scene graph, and application lifecycle model supporting multi-application shared-screen rendering. Architected process isolation so third-party apps drove UI through an IPC layer (Cap'n Proto) without direct compositor access. Developed cross-service debugging tools across the distributed microkernel.
- Newsfeed Ads: Took a data science finding (showing ads at top-of-feed could generate ~\$x.xB annually) from study to production. Designed the implementation plan, evaluating cached-ad and lightweight-ranking approaches to fit within strict latency and relevance guardrails. Debugged video ad rendering bottlenecks and layout performance issues that were blowing latency budgets.
- Ray-Ban Glasses: Debugged the multi-hop logging pipeline (glasses to mobile app to backend) resolving dropped logs, memory pressure on-device, and bandwidth contention with user media transfers. Built LLM-powered unit test generation tooling for the Ray-Ban development team.

### Founder / Software Engineer | Earlyworm

2022 - 2023

- Built and deployed a language learning app on Kubernetes with microservices architecture. Fine-tuned LLMs, NER, clustering, summarization, recommendation systems, and vector search over embeddings.

### iOS Team Lead | DayDayCook

2015 - 2018

- Led iOS team for a recipe platform (millions of downloads, 30+ countries, 5 Apple platforms). Featured repeatedly on App Store front pages; HK Top Ten Apps 2016. Company raised \$30M; IPO'd on NYSE.

## INDEPENDENT RESEARCH

### Diffusion Models for ARC-AGI

July 2024

- Designed a novel approach to ARC-AGI using per-task latent diffusion with custom autoencoder and UNet architectures. Iterated through six architectural variations; breakthrough was a spatial latent representation preserving grid structure. Full retrospective: [christopherhwood.com/Diffusion-Models-for-ARC-AGI-A-Retrospective](http://christopherhwood.com/Diffusion-Models-for-ARC-AGI-A-Retrospective)

- Solved problems from the original dataset over a year before the broader diffusion-for-ARC research thread emerged. The ARCHitects later placed 2nd in ARC Prize 2025 using a conceptually related masked diffusion approach.

## TECHNICAL SKILLS

---

**Languages:** Python, TypeScript/JavaScript, Swift, C++, Go

**AI/ML:** Diffusion models (Flux, LoRA, ComfyUI), UNet architectures, attention mechanisms, LLM agent systems, embeddings, NER, clustering, RAG, PyTorch

**Infrastructure:** Node.js, PostgreSQL, React, Docker, Kubernetes, microservices, distributed systems, API design

**Platforms:** iOS/macOS (Swift, XCUITest), web (full-stack), AR/VR operating systems

## EDUCATION

---

**University of Virginia** - B.A. in Economics and Chinese (Double Major), 2013

**Hopkins-Nanjing Center** - Graduate Certificate in International Relations, 2014