

Daniel Chen

San Diego, CA | (858) 396-4613 | dac030@ucsd.edu | github.com/dan9574

EDUCATION

University of California San Diego

B.S. in Mathematics–Computer Science | GPA: 3.915/4.000

Expected Spring 2028

San Diego, CA

TECHNICAL SKILLS

Languages: Python, TypeScript/JavaScript, Rust, Swift, Java, SQL, C/C++

Backend: FastAPI, SQLAlchemy/Alembic, PostgreSQL, pgvector, Spring Boot, JPA, SQLite, Supabase

Frontend & Apps: React, Vite, Tauri 2, SwiftUI/SwiftData, WidgetKit, CodeMirror 6

Tools & ML: Git, GitHub Actions, pytest, Linux, PyTorch, Transformers, PEFT/QLoRA, vLLM

EXPERIENCE

Altheros Capital — Cara AI

Sep 2026 – Present

Backend Lead, Part-time (Software Engineer, Sep–Oct 2026)

Remote

- Promoted to Backend Lead of an AI companion app (FastAPI, PostgreSQL, pgvector) after three weeks of top individual delivery; assign backend tasks and review the team's pull requests.
- Designed and shipped the RAG ingestion pipeline syncing journals, profiles, and evicted chat history into the vector store: deterministic serialization with SHA-256 dedup for idempotent retries, transactional doc-id write-back, and a pending-delete queue so user saves never fail on outages.
- Added backward-compatible doc_id responses to the retrieval API and per-user reverse lookup so chat citations resolve to the originating journal entry or conversation.
- Tiered retrieval timeouts by query mode from benchmarks (rerank latency linear in top_k), verified by mutation testing; shipped 6 merged PRs across 2 repos (~4,100 lines, ~70 tests) with green CI.

Xiamen Limayao Network Technology Co., Ltd.

Jun 2026 – Sep 2026

AI/ML R&D Intern

Beijing, China

- Reworked an untested Qwen3-4B SFT codebase into a versioned Python pipeline for data maintenance, training, stratified evaluation, model merging, and production-compatible export.
- Turned 6.17M raw product records into 66,429 curated samples (validation, deduplication, candidate generation, hard negatives); a two-stage Doubao/DeepSeek label audit corrected 4,173 mislabeled records.
- Fine-tuned Qwen3-4B with QLoRA SFT, raising core validation accuracy from 58.78% to 88.49% and candidate-set compliance from 78.59% to 99.98%.
- Deployed Qwen3.5-2B to production with 2.33x vLLM throughput (45.89 vs. 19.73 samples/s) within 0.17 pp of Qwen3-4B accuracy; shipped atomic dataset releases and inference fallback fixes.

Beijing Mapvision Technology Co., Ltd.

Jun 2025 – Aug 2025

Software Engineering Intern

Beijing, China

- Built an internal full-stack platform (Spring Boot, JPA, PostgreSQL, React) with REST APIs, cookie-session auth, and user-level authorization; normalized ~440K brand records into ~400K canonical brands.

PROJECTS

Hope — Cross-Platform Time Tracker

2025 – Present

Open-Source Personal Project | Tauri 2, Rust, React, TypeScript, Swift, SQLite, Supabase | GitHub

- Rebuilt an Electron app as a Tauri 2 desktop app (macOS/Windows) plus native iPhone and Apple Watch apps in Swift (~23K lines, 150 unit tests); Rust owns all data access behind typed commands.
- Designed an offline-first sync protocol on Supabase: UUID keys, last-writer-wins soft deletes, server-sequence cursors, Postgres triggers for conflicts and a single running timer, and row-level security.
- Built click-through desktop overlay cards with native blur (NSVisualEffectView on macOS, Acrylic on Windows), a WidgetKit watch-face complication, and cross-device break reminders over WatchConnectivity.

Office Flow — Obsidian Plugin (TypeScript, CodeMirror 6) | GitHub

2026

- Built a CodeMirror 6 editor extension (decorations, widgets, keymaps) for inline task actions, a code-block DSL rendering cross-note progress cards, and hold-to-archive/restore with write-before-trim safety.