

Justin Sublette

AI Systems Engineer | Autonomous Agents, Context and Research

j.sublette@vmosaic.com · 812-374-4565

Naples, Florida, USA · Willing to relocate

Open source: [Mathematical Research](#) · [Protocol Runner](#)

Profile

AI systems builder developing autonomous research systems, persistent agent workflows and context infrastructure through substantial software projects. Combines architecture, integration and experimental judgment to preserve useful work across agents, investigate failures and turn model capabilities into working applications.

Experience

Founder and AI Systems Engineer | Vector Mosaic Inc. | 2025-present

Founded a software company spanning autonomous research, agent infrastructure and practical applications. Own problem definition, architecture, integration, evaluation and release acceptance, collaborating extensively with Codex on implementation, investigation and testing. Began practical software development in June 2025 through self-directed, AI-assisted learning. Selected systems below are projects within this work.

Mathematical Research / RH Mission | Autonomous research across replaceable agents | 2026-present

- Created a system for autonomous AI research toward a checkable proof or disproof of the Riemann Hypothesis. Designed a durable Mission whose objective survives replacement of its coordinating agent; the AI executive chooses and revises mathematical strategies, delegates research and criticism, and synthesizes results.
- Separated current strategy from cumulative scientific understanding, exploratory branches, evidence and proposed results. Designed successor orientation around qualified scientific context and exact retrieval, preserving useful partial results and failed approaches without rebuilding an ever-growing transcript.
- Built a single-writer SQLite store with immutable revisions and content-addressed sources. Integrated atomic publication of an evidence result with its authored scientific continuation and exact-request replay, preventing interruption from retaining the result while losing that paired continuation.
- Open-sourced the real research core and Mission Host under MIT with [four reproducible examples](#). A reference-assisted live successor retrieved exact prior records, retained the corrected assumptions and tightened a localization bound from 10 to less than 9; the standalone system also passed real Linux stop/restart and research-tool interaction.

RH remains open. Public results cover the standalone release with fresh state and the selected examples.

Protocol Runner | Persistent execution for AI engineering workflows | 2026-present

- Designed a system that retains plans, step status, attempts and evidence outside conversations; resumes dependent work in the same conversation and launches independent assignments as separate workers. The coordinating agent retains quality judgment and final acceptance under my direction.
- Used Runner hundreds of times for development and investigation; a retained census identified more than 4,400 worker tasks matching its prompt signature. After a worker used the wrong checkout, directed a fix binding execution to the launching Git worktree and verified the assignment in a fresh run.
- Coordinated 25 investigation assignments covering 595 commits for an engineering-history update. Separate AI review identified three of 40 accounts needing correction; preserved completed work, corrected those accounts and obtained a fresh full-draft pass before internal acceptance.
- Open-sourced [Protocol Runner](#) under MIT, with a real API/store recovery demo using simulated workers and separate live Codex CLI completion and cancellation checks. [Evidence](#).

CTX Engine | Repository context and agent evaluation | 2025-2026

- Combined keyword and semantic search with dependency analysis in a CLI that returns task-specific reading guides and recommendations, with agents inspecting the underlying sources. Added version-aware indexing and recovery that preserves previous searchable data while marking it stale.
- In 30 internal code-navigation tasks from the project's own repository, retrieved 69% of reference-relevant files versus 34% for keyword search; 85% of selected files were relevant.
- Built answer-comparison harnesses with protected reference answers and frozen holdout conditions. Repaired invalid run conditions and completed a clean rerun. An experimental replacement did not improve answers on 24 same-repository held-out tasks; retained the standard configuration rather than adopting from development results alone.

Life Links | Private physical-world context for AI clients | 2026-present

- Created a web application for owner-curated places, objects, documents, routines and calendars. Shared ownership, revision and privacy rules across the UI, page-bound WebMCP and OAuth-protected remote MCP access.
- Qualified deployed Codex and ChatGPT workflows that created, read and edited private context across conversations while the site was closed. Tested exact-target deletion confirmation, replay and revocation; repaired defects exposed by bounded live Google and Outlook calendar round trips.
- Redesigned Collection loading: startup API requests fell from 55 to 7 for the same 48-member Collection. Median first-content rendering improved 43-57% across four matched desktop/mobile-width conditions, with three observations each. Released a sanitized MIT-licensed application and submitted it to the OpenAI WebMCP Challenge.

Pre-Note Delta | Agent access to authenticated vendor documentation | 2026

- Built browser extraction, authenticated capture and a searchable local library so coding agents could work from Epic developer contracts. A retained collection contained 191 of 192 selected API documents, supporting an Epic-integrated chart-change application and automated testing. Deployed an AWS demonstration environment and prepared hospital pilot materials; no hospital rollout was established.

Technical foundation

TypeScript, Python, Node.js, React, PowerShell; PostgreSQL, SQLite; MCP, REST/OpenAPI, OAuth/PKCE; semantic and keyword retrieval, dependency analysis, controlled agent evaluations; Playwright, browser debugging protocols, Git worktrees, Windows/Linux, containers, AWS and Railway. Experience centers on existing-model integration and agent systems; model construction and training remain a development area.

Earlier experience and education

Manufacturing and pipeline construction, 2012-2018: shift-team leadership and field coordination. Independent market research and trading, 2018-2025.

CCNA: certified April 7, 2025; valid through April 7, 2028. **Indiana University, SPEA:** policy analysis studies, 2013-Spring 2015; no degree completed. **Columbus North High School:** Academic Honors diploma.