KEVIN FAIRLEY

AI Researcher — Agentic Systems, Evaluation & Reliability

Los Angeles, CA | 213.855.1190 | kfairley9999@gmail.com

https://github.com/kfairley9999 | https://www.linkedin.com/in/kevin-fairley-94330b42a/

PROFESSIONAL SUMMARY

AI researcher focused on making LLM agents measurable and dependable. Builds open-source tool-calling agents, multi-agent systems, eval harnesses, regression gates, and code-enforced guardrails, directing AI coding agents, and publishes reproducible results and failure analyses. Brings hands-on server and data-workflow experience. Seeking fully remote AI Research Engineer / Applied AI roles.

RESEARCH & PROJECTS (open source: github.com/kfairley9999/<project>)

catchpr — Execution-gated AI pull-request reviewer | Python, GitHub Actions
- Comments only with a generated pytest that passes on base and fails on the PR (3/3 runs); 35 tests, CI green

leadflow-mcp — MCP server for a small-business lead pipeline | TypeScript, Cloudflare Workers, D1
- 11 tools to capture, score with a published rubric, and follow up on leads; 113 tests; live on Cloudflare

mini-agent — Tool-calling agent with self-check and eval harness | Python
- Built a provider-agnostic agent loop (OpenAI-compatible; Groq, Gemini, OpenAI) with calculator, sandboxed Python, and web-search tools, a step budget, a required self-check, and rate-limit-aware retries
- Designed a 15-task eval with automated checkers: 14/15 (93.3%) on a hosted Groq model; published failure analysis showing the self-check accepted an off-by-one answer, motivating a re-execution verifier

agent-crew — Multi-agent research desk | TypeScript, Cloudflare Workers, React
- Planner, researcher, writer, and critic agents share a blackboard; controls live outside the model: a deterministic check rejects citations to unretrieved sources, a capped critic loop, a token/dollar budget

llm-watch — LLM call ledger and eval regression gate | Python, FastAPI, React
- Logs each LLM call's tokens, latency, and cost; gates model swaps in CI with an eval suite (regex, JSON-schema, LLM-judge scorers), a paired bootstrap 95% CI, and pass-to-fail flip detection

hermes-phone-agent — Guardrails for an open-source agent's phone/SMS skills | Python
- Wrapped Nous Research's Hermes Agent telephony skill with guardrails enforced in code, not prompts: allowlist, confirmation, call/text/spend caps, quiet hours, and an append-only audit log; mock-by-default providers

fast-tokenizer — Byte-level BPE tokenizer | Rust, Python (PyO3)
- Exact GPT-2 token parity with tiktoken on 905,180 tokens; about 2.5× tiktoken's single-thread speed from Python and 155.7 MB/s with 4 threads

TECHNICAL SKILLS

AI/ML: LLM agents, tool calling, multi-agent orchestration, evals and regression gates, LLM-as-judge, guardrails, audit logs, token/cost budgets, BPE tokenization
Languages: Python, TypeScript, JavaScript, Rust, HTML/CSS; English and Spanish (bilingual)
Tools: FastAPI, React, Cloudflare Workers, SQLite, PyO3, pytest, Playwright, Docker, GitHub Actions, Git
Security: Bug bounty / vulnerability research

EXPERIENCE

Independent Security Researcher (Bug Bounty) | Bugcrowd | Remote | Nov 2026–Present
- ID-verified researcher on Bugcrowd; actively hunting on public programs
- Focus on AI/LLM applications, agent and MCP tooling, APIs, and web apps, within each program's scope and rules
- Use Semgrep, Bandit, Gitleaks, dependency audits, and passive recon; write reproducible vulnerability reports

Data & Workflow Specialist | Solvis Corp. | Commerce, CA | 2024–Present
- Validate and integrate multi-source data, automate spreadsheet workflows, and build visualizations for leadership

Equipment Manager / Logistics Specialist | U.S. Artificial Grass | Los Angeles, CA | Jan 2018–Dec 2023
- Managed inventory, freight, and equipment with computer-aided inventory, GPS, and diagnostic tools

IT Specialist / Server Management | Kesem Technology | Washington, DC | 2014–2017
- Maintained server infrastructure for high availability and security compliance; patching, network health monitoring, and OSI-model root-cause troubleshooting

CERTIFICATIONS & EDUCATION

- CompTIA A+ — earned 2014; The Hope Project, Washington, DC.
- Job Readiness & Professional Development | Chrysalis, Los Angeles, CA | 2024
- High School Diploma | James Irwin Charter, Colorado Springs, CO

Volunteer: Hospital Volunteer, Good Samaritan Hospital, Los Angeles, CA (2025)