

James Wiesebron

jwiesebron@pm.me | South Lake Tahoe, CA | github.com/jwbron

Engineer focused on AI agents and productivity tooling for knowledge work: adoption and training, agent guardrails, evaluations and cost-efficient AI usage. Backed by ten years of platform and data engineering, including work with sensitive student, health and financial data.

Experience

Khan Academy: Software Engineer

2022–2026

- Drove AI adoption and set best practices across Khan Academy. Taught engineers and non-engineers, including designers, to use AI tools, and built tools they used in their daily work.
- Owned AI spend and cut costs by making usage more efficient rather than reducing it. Evaluated vendors, then designed and ran an AI code-review harness on github actions and the claude agent SDK, with specialized reviewers, validated findings and scoped follow-ups. It replaced a commercial bot at about one fifth of the cost, with deeper reviews that saved engineers time.
- Built coding-agent harnesses with OS-level sandboxing, credential brokering and scoped permissions. Enforced required workflow checks in code and routed agent-reported tooling failures to a shared backlog and jira, with human triage for ambiguous reports.
- Built evaluation infrastructure with repeatable A/B runs, regression monitoring, transcript capture and cost tracking by model, separating model failures from problems in test setup and scoring.
- Built and operated autoscaling, self-hosted github actions runners on cloud run that halved compute costs for CI. Drove terraform adoption with a reusable framework and migrated services from app engine, compute engine and kubernetes engine to cloud run.
- Wrote RFCs and ADRs, gave internal talks, interviewed and onboarded engineers, and led incidents on call. Led the migration of khanmigo support feedback to reforge analytics with a TPM.

Asana: Data Engineer (L5)

2021–2022

- Designed and built a metrics analysis framework and dashboards that helped data scientists and analysts investigate changes, forecast consistently across teams, and maintain company-wide visibility into metric definitions.
- Developed feature-store infrastructure and tooling for data scientists to iterate on ML models, including experiment tracking, data quality assessment, and scalable model evaluation.

Rally Health: Senior Software Engineer II

2019–2021

- Led a team through the rewrite of a healthcare search engine on Scala and elasticsearch during its migration from elasticsearch 2 to 6, and rebuilt a high-volume batch ETL pipeline as distributed streaming with kafka.

Intuit: Senior Software Engineer

2018–2019

- Designed and built a backend web service using event-driven, distributed message processing and scoring for the quickbooks rewards program.

Capital One: Software Engineer, Technology Development Program

2016–2018

- Built a Scala and akka orchestration framework for containerized ML models behind credit card application and fraud decisions, partnering with data scientists (full-time, 2017–2018).

Skills

AI adoption and training, vendor evaluation, agent harnesses and guardrails, evals, Python, Go, TypeScript, GCP, cloud run, github actions, terraform, kubernetes, kafka, claude agent SDK, CI/CD.

Education

University of North Carolina at Chapel Hill: Bachelor's degree in Computer Science, May 2017. Teaching assistant for Data Structures.