

Hannah Lim Junior AI Engineer

Austin, TX · hannah.lim@gmail.com · +1 512-555-0173 · linkedin.com/in/hannahlim

PROFILE SUMMARY

- **Junior AI Engineer** with **2 years** shipping **LLM-powered features** for a **marketing-content platform**, across **prompt design, structured outputs, and evaluation**, specializing in **prompt test suites, eval dashboards, and RAG quality checks**.
- Hands-on across primary language (**Python**), model APIs (**OpenAI** and **Anthropic**), orchestration (**LangChain**), evaluation (**Promptfoo**), vector search (**Pinecone**), and serving (**FastAPI**), with working knowledge of LangGraph and Hugging Face.
- Growing expertise in **prompt regression testing, JSON-schema output validation, and golden-set evaluation**, using **weekly eval reviews** and **per-prompt cost tracking** to ship **features that behave the same way on Friday as they did on Monday**.
- Works daily with a **senior AI engineer, product managers, and content leads** in two-week sprints, joining eval reviews, release checks, and the customer feedback triage for 3 LLM-facing features.
- Started in this field by testing prompts by hand at an internship: **knows exactly how a model fails in production**, and now builds the suites that catch it before customers do.

TECHNICAL SKILLS

Languages:	Python, SQL, TypeScript (basics), Bash, Git
LLM Tooling:	OpenAI SDK, Anthropic SDK, LangChain, LangGraph (basics), structured outputs, function calling
Prompt Engineering:	few-shot design, system prompts, JSON-schema outputs, length budgets, brand-voice constraints
Evaluation:	Promptfoo, golden sets, LLM-as-judge (basics), regression suites, cost per prompt
Retrieval:	Pinecone, embeddings, chunking strategies, recall@k, citation checks
Serving & Data:	FastAPI, Streamlit, Docker (basics), pandas, PostgreSQL
Cloud:	AWS (Lambda, S3, basic IAM), GitHub Actions

EDUCATION

University of Texas at Austin	B.S. in Computer Science	Austin, TX · Aug 2019 - May 2023
-------------------------------	--------------------------	----------------------------------

WORK EXPERIENCE

Jasper AI	Junior AI Engineer	Austin, TX · Aug 2023 - Present
• Shipped 9 LLM-app features for the marketing-content surface in Python and FastAPI under senior review, each one behind a feature flag and a prompt regression suite. • Wrote 54 prompt tests in Promptfoo, including golden-set fixtures and JSON-schema validators for structured outputs, catching 12 regressions before they reached customers. • Own 3 prompt test suites covering brand voice, schema conformance, and length budgets, reviewed weekly with the senior engineer and run on every pull request in CI. • Built a Streamlit eval dashboard that 4 product managers and 2 content leads use to compare model outputs, golden-set diffs, and cost per prompt across 3 model providers. • Found a chunking bug in the Pinecone retrieval evaluation notebook that was inflating citation errors, and the fix cut unsupported citations by 22% on the internal eval set.		
Glean	AI Engineering Intern, then Junior AI Engineer	Palo Alto, CA · May 2022 - Jul 2023
• Ran 4 prompt experiments on the enterprise-search assistant in Python with the OpenAI SDK, including LangChain chains for follow-up rephrasing. • Contributed 30 golden-set tests to the citation-grounding pipeline and paired with a senior engineer on the retrieval-quality dashboards the team still uses. • Wrote a FastAPI internal endpoint exposing the assistant's structured-output schema, which 2 downstream teams used to build their own integrations.		