

HAMED ALIKHANI, PHD

+1 (979) 676-7452, Austin, TX | Authorized to work in the U.S. (No sponsorship required)

hamed.alikhani.phd@gmail.com | [LinkedIn](#) | [GitHub](#) | [Google Scholar](#)

SUMMARY

PhD-level **Senior AI Engineer** specializing in **Agentic AI Architecture** and **LLM Safety**. Proven ability to design, deploy, and optimize **enterprise-scale multi-agent systems**, including high-impact work at McGraw Hill and Kraft Heinz. I focus on the core pillars of production-grade AI: **Security & Guardrails, Cost Efficiency, and System Reliability**.

Expert in **Multi-Agent Orchestration**, tool calling, and hardening AI pipelines against adversarial attacks, with demonstrated improvements in **risk classification performance** and **system-level cost optimization**. My technical expertise spans Semantic Routing, Hybrid Concurrency (Async/Threading), FastAPI microservices, and MLOps, delivering measurable outcomes including a **\$100M annual profit impact** from forecasting pipelines.

CORE SKILLS

Agentic AI & LLM Systems: Multi-Agent Orchestration & Tool Routing (LangChain, CrewAI), LLM Safety & Guardrails (Prompt Shield, Bedrock), Tool/Function Calling, Structured Validation (Pydantic), Semantic Routing

RAG & Retrieval Systems: RAG Ingestion Pipelines, Structure-Aware Chunking, Hybrid Retrieval (BM25 + Semantic Search), Context Assembly, Token & Latency Optimization

Evaluation & Observability: LLM-as-a-Judge Evaluation, Retrieval Evaluation (Precision/Recall), Red-Teaming, LangSmith Tracing & Experimentation

Backend & Production: FastAPI, Microservices, Hybrid Concurrency (Async/Threading), Docker, CI/CD (GitHub Actions), Azure Functions

ML & Data Pipelines: Transformers, LSTM, Time-Series Modeling, Optimization, Distributed Computing (PySpark), Data Engineering, Python (Advanced), SQL

Cloud & Infra: Azure, AWS Guardrail, Databricks, Vector Databases (Weaviate, Chroma), Vast.ai

EXPERIENCE

Senior AI Engineer (Contract) | McGraw Hill

June 2025 – Present

- Led production-grade **LLM safety and document ingestion systems** for PDF/DOCX/PPTX, achieving **96% sensitivity/ 90% specificity** while maintaining **\$0.00145 per 1K characters** validation cost.
- Built and evaluated a **production RAG pipeline** with structure-aware chunking and **hybrid retrieval (BM25 + semantic search)**, achieving **>90% accuracy/recall** while optimizing latency and token usage.
- Standardized **structure-aware parsing and normalization** for RAG ingestion, creating token-efficient representations that **reduced cost** while preserving academic layout fidelity.
- Built **LLM evaluation and observability infrastructure** (LangSmith, red-team datasets) to enable rapid testing of retrieval and safety, deploying the pipeline as a **FastAPI microservice**.

Data Scientist | Kraft Heinz Company

May 2023 – Mar 2025

- Designed a custom **genetic-algorithm-based ensemble optimization engine**, improving forecast stability by 18% and contributing to an estimated \$100M annual profit uplift.
- Engineered **distributed time-series forecasting pipelines** (Pandas, PySpark, Azure ML), scaling processing across thousands of SKUs to optimize inventory planning.
- Built LLM-based SQL generation tools to automate data access workflows, significantly **increasing analyst productivity** and reducing manual query-writing load.
- Led a rapid-prototype project to analyze 100K+ nutrition labels for market trends, **winning the company AI Hackathon (2023)** and influencing portfolio strategy.

Research Data Scientist | Texas A&M University

Sep 2018 – May 2023

- Built forecasting, optimization, and computer vision models (LSTM, XGBoost, Faster R-CNN) for deployment in real-world transportation systems; earned multiple research awards.
- Processed and modeled 1,000+ large-scale infrastructure datasets for strategic planning and operational optimization.

RELATED PROJECTS & CONSULTING

Independent AI Architect & Consultant | Self-Employed April 2025 – June 2025

- Designed and piloted a **voice-enabled, multi-agent workflow** (Retell AI, Twilio) for dental clinics, focusing on appointment intake, triage, and tool-based routing.
- Architected a **privacy-preserving AI prototype (SoulShield)**, defining system design, guardrails, and zero-retention LLM policies using Azure-based models.

AI Engineer (Part-Time) | AIHubAgents Dec 2024 – Apr 2025

- Fine-tuned Mistral and GPT-4o models using **LoRA/QLoRA** on Vast.ai to adapt agent behavior across tone, memory, and safety dimensions.
- Developed adaptive microagents with dynamic prompting and routing, contributing to a **12% improvement in user engagement** during pilot testing.

EDUCATION

Ph.D., Multidisciplinary Engineering (AI-focused) | Texas A&M University 2023

M.S., Project Management | Sharif University of Technology 2014

B.S., Engineering | University of Tehran 2012

LEADERSHIP

Founder & President, Austin AI Community (Nonprofit) 2024 – Present

Lead AI literacy and hands-on workshops for 200+ learners. Delivered AI training for PCSI executives (AI 101/102, Copilot, workflow automation). Designed trauma-aware AI education programs and supported development of SoulShield AI.

AWARDS

Best Paper Award, ASCE Journal 2024

Kraft Heinz AI Hackathon Winner 2023

Outstanding Graduate Scholarship, Texas A&M University 2020–2022

High Value Research Award, Montana DOT 2021