Paul Stanley Ganganapalli

+49 162 7220780 |Stuttgart,Germany,70771| **PaulStanleyganganapalli@gmail.com** |
http://www.linkedin.com/in/paul-stanley-218002195 |**https://github.com/PaulStanley0211**

PROFESSIONAL SUMMARY

AI Application Engineer building production LLM systems with hybrid retrieval, agentic patterns, and reproducible evaluation. 6 shipped AI projects spanning RAG, multi-agent orchestration, document AI, and conversational systems for the German market. Mechanical engineering foundation and quantitative trading discipline applied to AI product development.

PROFESSIONAL PROJECTS

Production RAG System

Production-grade RAG system with hybrid retrieval, agentic CRAG, self-correction, three layered security and reproducible evaluation harness.

- Built end-to-end RAG pipeline: hybrid retrieval (BGE-small dense + BM25 sparse with Reciprocal Rank Fusion), cross-encoder reranking, CRAG self-correction.
- Implemented three security layers: prompt-injection input guard, content filter on retrieved chunks, and post-stream PII redaction with Luhn-validated credit card detection, email/phone/SSN pattern matching.
- Achieved Hit Rate 5 of 0.96, MRR of 0.96, faithfulness of 0.944, and 100% routing accuracy across answer/partial/refuse paths on a reproducible evaluation harness with LLM-as-judge scoring (30 retrieval cases, 15 generation cases).
- Architected SSE streaming API with token/citation/status events; tiered model strategy (Sonnet for generation, Haiku for routing/grading) keeping per-query cost in the low single cents.

Tech Stack: Python, FastAPI, Qdrant, Redis, Anthropic Claude, React, TypeScript, Docker
GitHub link: https://github.com/PaulStanley0211/Production_RAG_System.git

QuantFlow – Multi-Agent Trading Workflow for German DAX 40

8-agent orchestrated system for pre-market analysis, strategy back testing and human in the loop trade approval**.**

- Architected 8 specialized agents (morning briefing, setup scanner, risk monitor, trade executor, security guardrails, strategy builder) with scheduled triggers, email-based human approval workflow via IMAP, and file-based audit trail for full traceability.
- Built full-stack Strategy Builder: FastAPI backend with Server-Sent Events streaming, React plus Recharts dashboard, plain-English-to-JSON strategy parser via Claude, multi-instrument backtest engine across 54 instruments (DAX 40, indexes, crypto, commodities).
- Implemented 7-rule pre-trade security guardrail (stop loss verification, position sizing, R/R minimum, daily loss limit, trade count cap, blocklist, duplicate detection) with hard-block on rule violations.
- Designed paper-trading-only with no broker integration every trade requires human YES/NO email approval before logging, ensuring full auditability and no automated execution risk.

Tech Stack: Python, Anthropic Claude, FastAPI, React, Recharts, yfinance, IMAP/SMTP
GitHub link: https://github.com/PaulStanley0211/Quantflow.git

MIA-German Retail AI Agent

Production grade conversational AI for German e-commerce with bilingual support and DSGVO-compliant memory

- Architected full-stack AI system managing 25+ orders and 25+ products across 6 categories using PostgreSQL persistence and LangChain orchestration with Claude as the reasoning engine.
- Automated 5 core customer service functions (order tracking, returns, product recommendations, FAQ, complaints), reducing support workload by 60-70%.
- Implemented persisted vector memory in ChromaDB enabling personalized customer context across sessions, with automatic DSGVO-compliant cleanup of expired user data.

- Engineered bilingual support (German + English) with automatic language detection, deployed via Docker with FastAPI REST API (3 endpoints) and Streamlit web interface for the German e-commerce market.

Tech Stack: Python, Anthropic Claude, LangChain, PostgreSQL, FastAPI, ChromaDB, Streamlit, Docker
GitHub Link: https://github.com/PaulStanley0211/MIA---German-Retail-Agent.git

PROFESSIONAL EXPERIENCE

Independent Quantitative Trader **Oct 2022 - Sep 2025**
India (Remote)

- Active trader in Indian equity, derivatives, and commodities markets, developing rule-based systems for entry/exit decisions, position sizing, and post-trade review across 3 years of consistent practice.
- Built personal frameworks for risk management daily loss limits, position sizing constraints, and trade journaling the same patterns later codified into the QuantFlow project's 7-rule security guardrail and audit-trail architecture.
- Applied systems thinking from disciplined trading clear rules, audit trails, defined risk limits to AI engineering practice, particularly in evaluation harness design and security guardrail systems.

Mechanical Design Engineer (Part time) **Mar 2021 - Feb 2022**
Balaji Transporters and handlers, India

- Designed mechanical components and assemblies for material-handling equipment using SolidWorks and AutoCAD, producing 3D models, technical drawings, and assembly documentation for production.
- Conducted design reviews and revisions across the product lifecycle, optimizing for manufacturability and cost collaborating with manufacturing and quality teams to validate feasibility before production.
- Applied engineering judgment to troubleshoot design issues during prototyping and pre-production, contributing to reliable, production-ready outputs.

EDUCATION

Riga Technical University, Riga, Latvia **Feb 2019 - Feb 2021**
Master of Engineering Science in Mechanical Engineering

Karunya Institute of Technology and Sciences, Coimbatore, Tamil Nadu, India **June 2014 - June 2018**
Bachelor of Technology in Mechanical Engineering

SKILLS

- **Languages & Frameworks: Python, SQL, TypeScript, FastAPI, React, LangChain, LangGraph, Streamlit**
- **AI Engineering: RAG, Hybrid Retrieval, Agentic Workflows, Vector Search, Cross-encoder Reranking, LLM Evaluation, Prompt Engineering, LLM-as-Judge, Multi-agent Orchestration**
- **LLMs & APIs: Anthropic Claude (Sonnet, Haiku, Opus), Claude Vision, OpenAI API**
- **Data & Storage: PostgreSQL, SQLite, Qdrant, ChromaDB, FAISS, Redis**
- **DevOps & Infrastructure: Docker, Docker Compose, Git/GitHub, Server-Sent Events (SSE), Azure (AZ-900)**
- **Certifications: Microsoft Azure Fundamentals (AZ-900), Microsoft Azure AI Fundamentals (AI-900)**

LANGUAGES

- German (B1 – Intermediate, working towards B2)
- English (C1 – Advanced)