

Hermann Ndamen

AI ENGINEER

ndamenh@gmail.com · +90 552 856 82 04 · Remote (worldwide)

CONNECT  LinkedIn  GitHub  Portfolio

SUMMARY

I'm an AI Engineer who ships production **Large Language Model (LLM)** systems end to end. I design and deploy AI systems like **RAG pipelines**, **AI agents**, and automation, built with **Python**, **Claude Code**, **LangChain / LangGraph**, **n8n**, **Make**, and the latest LLM APIs (Claude, GPT-5.2, Gemini, and more), and take them all the way to full-stack, cloud deployment. My public, running systems back it up, several of them open-source on GitHub. Discipline and consistency drive how I work: I practice personal development and believe in getting a little better every day.

TECHNICAL SKILLS

Languages	Python, TypeScript, SQL, R
LLMs & GenAI	Claude, Claude Code, GPT (OpenAI), Gemini, OpenAI Whisper; prompt engineering; tool / function calling; structured outputs (JSON Schema); AI agents
RAG & Retrieval	Retrieval-Augmented Generation (RAG); vector databases (pgvector / PostgreSQL); embeddings; semantic search; dense + lexical (BM25) retrieval; citation grounding
Orchestration & Eval	LangGraph; LangChain; LLM evaluation (recall / precision@k, faithfulness, LLM-as-judge); RAGAS-style eval harnesses
Backend & Infra	FastAPI, Next.js; REST APIs; Docker; serverless (Modal, Vercel); workflow automation (n8n, Make); PostgreSQL; GitHub Actions (CI); Git
Data & ML	scikit-learn, XGBoost, pandas, NumPy; hypothesis testing, regression, multivariate analysis (MANOVA, PCA, factor analysis); Claude Vision

PROJECTS

Website RAG Chatbot: live, lead-capturing assistant github.com/hermann-ndamen/website-rag-chatbot

- Built and shipped a production **RAG** chatbot on a live site: embeds each query, runs **semantic search** over a vector store, and answers grounded in retrieved documents with clickable **source citations** to reduce hallucination.
- Engineered a pluggable **vector database** (in-memory to **pgvector / PostgreSQL**) and a **tool-calling agent** that captures leads and books calls; wrote an **evaluation harness** scoring retrieval **recall / precision@k** and answer **faithfulness (LLM-as-judge)**. Stack: Claude, Gemini embeddings, pgvector, Next.js.

WhatsApp Planning Agent github.com/hermann-ndamen/whatsapp-day-planner

- Built an **LLM agent** on serverless cron (Modal) that holds conversational check-ins over the WhatsApp Cloud API and writes a Notion schedule each night from the day's real outcomes. Used **structured outputs** and tool-use; containerized with **Docker**.

Multimodal Content Pipeline

github.com/hermann-ndamen/content-pipeline

- Built an agentic pipeline that ingests media from Drive, understands it (**Whisper** transcription + **Claude Vision**), writes per-platform captions, and schedules to a Notion calendar behind one human approval, on a twice-daily serverless job.

Branded Content Agent: LangGraph orchestration

github.com/hermann-ndamen/carousel-producer

- Orchestrated a content-generation agent as a **LangGraph** state machine with a human-in-the-loop interrupt and a self-check to remake loop (conditional edges); pluggable image providers; Dockerized. Stack: Claude, LangGraph, Python.

Machine Learning & Statistics Research

github.com/hermann-ndamen

- Loan-approval prediction: MICE imputation and five tuned models (logistic regression, SVM, random forest, **XGBoost**, neural network) compared on accuracy vs minority-class recall (81.4% accuracy, 53% minority recall), SMOTE for class imbalance. Plus a multivariate study (MANOVA, PCA, factor analysis, 0.80 balanced accuracy). Python + R, **scikit-learn**.

PROFESSIONAL EXPERIENCE

AI Engineer & Content Creator · Independent

Remote · 2024 – Present

- Shipping **production AI systems** end to end (RAG pipelines, AI agents, automation), each open-source and metrics-backed, with one running live on my own site.
- Growing an audience by creating technical content that teaches founders and businesses how to apply production AI systems, building in public.

AI Content Producer & Video Editor · Easy Finance Academy

Remote · 2024 – 2025

- Cut content production turnaround time significantly by building reusable **AI production workflows** for an AI-generated video series (EFA Teens Academy), owning the full pipeline from concept to output.
- Earned a promotion from Video Editor to full-time production lead on the strength of those systems; worked hands-on across a wide AI toolset for image, video, and voice (Midjourney, Kling AI, Veo 3.1, Higgsfield, ElevenLabs, Pictory, InVideo, and more).
- Trained the content team on **prompt engineering** for AI image, video, and voice generation, standardizing quality and turnaround across the series.

Performance Marketing Analyst Intern · Enhancer (AI Ads for Meta & Google)

Ankara · 2025

- Improved campaign ROI by analyzing Meta and Google ad-campaign **data** alongside an AI ad-optimization platform to surface optimization opportunities; refined B2B audience targeting to lift conversion.

Marketing Content Producer · FutureOfVR

Ankara · 2023 – 2024

- Grew the company's YouTube channel from **13K to 21K subscribers** in about six months (+8K), producing platform-specific video and campaign copy that built the brand's organic audience.
- That audience drove new players to the company's game, Egg VR, which reached roughly **\$40K per month** in revenue at its peak.

EDUCATION & LANGUAGES

B.Sc. Statistics, Middle East Technical University (METU), Ankara, Türkiye. Coursework: machine learning, regression, multivariate analysis, statistical inference and hypothesis testing.

Languages: English (native) · French (native) · Turkish (conversational)