



Aditya Chaudhary

M.Sc. Artificial Intelligence
AI Agent & Full-Stack Engineer

CONTACT

Location: Berlin, Germany
E-Mail: adicadi158@gmail.com
LinkedIn: linkedin.com/in/adicadi
GitHub: github.com/adicadi

EDUCATION

M.Sc. Artificial Intelligence

BTU Cottbus, Germany
Apr 2023 – Present
Machine Learning, AI Systems,
Distributed Systems, Data Analysis,
Modelling of Perception & Action

B.Tech Computer Science

Amity University, India
Apr 2018 – Jun 2022
Data Structures & Algorithms,
Software Engineering,
Machine Learning

SKILLS

Agents: LangGraph, LangChain,
Claude Agent SDK, MCP, Tool
Calling, Streaming, Evals,
Multi-Agent Orchestration
GenAI: RAG Pipelines, LLM
Workflows, Prompt Engineering,
Structured Outputs
Languages: Python,
TypeScript / JavaScript, Dart, SQL
Backend: FastAPI, REST APIs,
PostgreSQL, PDF & Document
Parsing
Frontend: React, Next.js, Flutter
Infra: Docker, Git, Linux, AWS
Data & ML: pandas, NumPy,
scikit-learn, PyTorch

AWARDS

2026: Winner – Build with AI
Hackathon (BTU Cottbus & Google
Developer Groups)

LANGUAGES

English: Fluent
German: B1 (actively improving)

PROFILE

M.Sc. AI student at BTU Cottbus with an experience in **building and shipping AI agent systems end-to-end** – from **LangGraph multi-agent workflows and MCP-powered tool pipelines** to **production RAG and document-ingestion systems**, backed by **FastAPI / Next.js services** and Dockerised deployments. Most recently **won the 2026 Build with AI Hackathon** shipping a full-stack AI hospital management platform with 8+ live LLM workflows. Excited to bring that same **high-agency, ship-first mentality** to automating manual ops work and making existing agents smarter, cheaper, and more reliable.

EXPERIENCE

Full-Stack GenAI Engineer (Intern) | Belzir GmbH *May 2026 – Present*
Berlin, Germany (Hybrid)

- Owned the entire **end-to-end delivery of EmailGuard**, a production multi-tenant SaaS platform (FastAPI + PostgreSQL + React/TypeScript) – from architecture to Docker deployment, including a DMARC ingest pipeline, multi-tenant alert engine, DNS manager, and Keycloak-based IAM with webhook-driven user sync.
- Built and maintained a 7-VM private cloud** (VMware vSphere) covering PKI, IAM, SIEM, and monitoring – owning infra decisions end-to-end with no hand-holding.
- Debugged and resolved **complex integration failures** (Keycloak webhook sync, OAuth callback edge-cases, async pipeline race conditions) by reading traces, isolating root causes, and shipping fixes the same day.

Machine Learning Intern | Pixlia Tech *Mar 2024 – Sep 2024*
Remote

- Built and deployed **LLM-powered REST APIs with FastAPI**, integrating generative AI into internal tools with **streaming responses**, structured output parsing, and robust error recovery.
- Designed end-to-end **LLM & ML pipelines (Python, PyTorch, scikit-learn)**; delivered benchmark reports that informed product decisions on model selection and prompt design.
- Developed **modular, reusable AI components and prompt templates** that cut prototype ramp-up time; translated technical findings into clear reports for non-technical stakeholders.

PROJECTS

CliniFlow – AI-Powered Hospital Management Platform *2026*
Winner – Build with AI Hackathon 2026 (BTU Cottbus & Google Developer Groups)

- Engineered **8+ production LLM workflows** (live clinical scribe, structured document extraction, imaging analysis with confidence-scored differentials, discharge letter generation, DRV-compliance peer review) using a **dual-provider abstraction layer** (Claude Sonnet 4.6 + Gemini 2.5 Flash) hot-swappable at runtime.
- Built **PDF & document ingestion pipeline**: pdf-parse → LLM structuring → Prisma + PostgreSQL persistence – same pattern applicable to automating ops document workflows.
- Full-stack delivery** in days: Next.js 14 + TypeScript, Prisma ORM, NextAuth v5, multipart upload for medical imaging.

Stack: Next.js · TypeScript · Anthropic Claude · Gemini · LangChain · Prisma · PostgreSQL · pdf-parse · Tailwind · Zod

Second Brain Agent – Agentic RAG & Knowledge Assistant *2026*

- Shipped a **multi-tool LangChain agent** with document ingestion, vector-based semantic retrieval, and structured output parsing, deployed as a FastAPI service – a personal knowledge base that reasons over ingested content.
- Designed the **retrieval-augmented pipeline** end-to-end: chunking, embedding, similarity search, and context injection into LLM calls.

Stack: LangChain · Python · FastAPI · RAG · Vector DB · Structured Outputs · Docker