

CERTIFICATE OF COMPLETION

David Minker

has successfully completed the Program

AI Engineer Intensive Bootcamp

from February 16, 2026 to May 3, 2026

Total Completed Projects: 13 (out of 13 total)

Completed Modules: Intro to LLMs and pipelines, Building Agents & Multi-Agent Systems, Vector Databases, Building RAG Systems, Monitoring & Security, Deploying LLM-based apps

Module Descriptions:

- Intro to LLMs and pipelines: Covers core LLM API interaction patterns including conversation management, token tracking, function calling, and data retrieval with external APIs. Stack: OpenAI API
- Building Agents & Multi-Agent Systems: Covers multi-agent orchestration, memory management strategies, persistent storage, and async workflow automation across complex pipelines. Stack: FastAPI, Pydantic, Groq API, Docker, LlamaIndex, GitHub API.
- Vector Databases: Covers vector search engine design, embedding storage and retrieval, metadata filtering, and performance optimization through quantization and HNSW tuning. Stack: Qdrant, OpenAI API, FastAPI, Docker.
- Building RAG Systems: Covers RAG pipeline construction including document chunking, query rewriting, HyDE retrieval, metadata filtering, re-ranking, and multi-source query routing. Stack: LangChain, Qdrant, Cohere.
- Monitoring & Security: Covers LLM observability, tracing, evaluation metrics, prompt versioning, Redis-backed chat history, guardrails for input validation, and API gateway management. Stack: Langfuse, Ragas, Redis, NeMo Guardrails, LiteLLM.
- Deploying LLM-based apps: Covers productionizing LLM microservices with REST APIs, managed cloud infrastructure migration, multi-stage containerization, and cloud deployment. Stack: FastAPI, Docker, AWS EC2, Amazon ECR.
- Capstone: Developed a personalized study tool for a middle school student that generates curriculum-aligned practice exercises, tracks topic weaknesses over time, and produces PDF worksheets. The app ingests real school curriculum and teacher materials to tailor exercises to what the student actually struggles with. Stack: Ollama (Qwen2.5), OpenAI API, ChromaDB, Sentence Transformers (multilingual-E5), Langfuse, ReportLab, SymPy, TinyDB, Streamlit, OpenAI Whisper, Python

