



Retrieval-Augmented Generation (RAG) Course

—
Build AI that answers from your own knowledge base — accurately.

Certificate Course | A26

Duration: 24 hours

Level: Intermediate–Advanced



AMSTERDAM
DATA ACADEMY



WHAT YOU'LL LEARN

- Build a complete RAG pipeline from document ingestion to LLM response
- Create and query vector embeddings using FAISS or Pinecone
- Implement chunking and retrieval strategies for accuracy
- Evaluate RAG output quality and reduce hallucinations
- Connect your pipeline to a real knowledge base or document set
- Deploy a working knowledge-grounded AI assistant



COURSE CONTENT

Module 1 — RAG Fundamentals

- RAG vs fine-tuning: when to use which approach
- The RAG pipeline: ingest → embed → retrieve → generate

Module 2 — Document Processing & Chunking

- Chunking strategies: fixed, semantic, and hierarchical
- Preprocessing and cleaning text for embeddings

Module 3 — Vector Databases & Embeddings

- What are embeddings and how do they encode meaning?
- Similarity search: cosine distance and ANN algorithms



AMSTERDAM
DATA ACADEMY



AFTER THIS COURSE

You'll have built a working RAG application grounded in a real document set — a portfolio piece and a practical foundation for AI engineering roles.

OFFICIAL CERTIFICATE

This course leads to an official ADA certificate at NLQF level 6, the nationally recognised qualification framework equivalent to a bachelor's degree. Your certificate proves a verified, professional level of knowledge and skills — valued by employers across the Netherlands.

Ready for your next step?

Book a free study-advice call

amsterdamdataacademy.com/study-advice

+31 85 004 9642