CORSEN AI Academy

Embeddings, RAG and Semantic Search

Course syllabus

Version 2026-11

Overview

This course covers OpenAI embeddings, vector databases, semantic search, a complete RAG pipeline, and production applications.

Learning objectives

- Generate embeddings and measure semantic proximity between texts using cosine similarity
- Index a corpus in a vector database and set up hybrid semantic and lexical search
- Build a complete RAG pipeline: chunking, retrieval, reranking, then generation
- Evaluate the relevance of a RAG pipeline and choose between File Search and a custom architecture
- Use embeddings to classify, cluster, recommend, and detect anomalies and duplicates

Audience and prerequisites

Developers, engineers, and integrators who want to use ChatGPT via API in their applications.

Ability to read Python and JSON, command-line basics. An OpenAI platform API key to try the examples.

Duration and format

Estimated duration : 2 h 15 min. 100% online self-paced course, accessible 24/7 from a browser. Progress is saved, and you can resume at any time. Estimated duration: 2 h 15 min.

Content

- OpenAI embeddings (5 lessons)
- Semantic search (5 lessons)
- RAG — Retrieval-Augmented Generation (6 lessons)
- Applications (6 lessons)

Assessment

Self-assessment based on the Key takeaways that conclude each lesson, and direct practice of the examples on your own examples. A final 5-question quiz with answers, at the end of the last lesson, validates what you have learned.

Accessibility

Self-paced course accessible to people with disabilities: text content compatible with screen readers, keyboard navigation, and compliant color contrast. For any adaptation needs, email hello@corsen.ai.

Price and recognition

Free

Certificate of completion issued by Corsen AI.

Access time

Immediate after account creation and email confirmation.

Corsen AI — Paris, France — SIRET 10297040700015
academy@corsen.ai — https://academy.corsen.ai