

Vector Databases Training

Duration: 2 Days

Description

The purpose of this training is to teach how to generate embeddings from text and images, store them in vector databases (FAISS, Chroma, PgVector, Milvus, Weaviate), perform efficient similarity search, and integrate them with RAG scenarios in production environments. Participants will gain hands-on experience with chunking strategies, metadata schemas, index and distance metric selection, hybrid search, and performance evaluation.

Learning Outcomes

- Generate embeddings and store them in Oracle 23ai, Chroma, PgVector, and Milvus
- Design hybrid search and re-ranking pipelines
- Build and optimize RAG workflows for quality and performance

Audience

- Backend and data engineers
- Search and RAG application teams
- Teams building semantic search and retrieval systems

Outline

Vector & Retrieval Fundamentals

- Embedding models and normalization
- Chunking strategies
- Distance metrics and indexing
- Vector database types
- Metadata schema design
- Hybrid search and evaluation
- Hands-on: PgVector similarity search

RAG Integration & Best Practices

- RAG architecture
- Advanced chunking strategies
- Re-ranking techniques

- Performance and cost optimization
- Observability and governance
- Hands-on: End-to-end RAG pipeline

Prerequisites

Basic programming knowledge and familiarity with data structures.