

Oracle AI Database: AI Vector Search Training

Oracle Database

DURATION

4 Days

MODULES

18 Lectures

COURSE CODE

—

Course Overview

This course will be delivered with a live lab. Master Oracle AI Vector Search and unlock the power of effective data retrieval. In the "Oracle AI Database: AI Vector Search Training" course, you'll learn how to create and query vector indexes, leverage embedding models, and utilize Retrieval Augmented Generation (RAG). You'll also discover how to integrate AI with OCI Generative AI services to enhance search accuracy and efficiency. By the end of this course, with valuable feedback on AI-driven workflows, you'll be equipped to optimize your approach and effectively apply your skills. After completing this course, you should be able to:

- Create and configure vector indexes for efficient similarity searches
- Apply vector distance functions and similarity metrics to effectively query vector data
- Use embedding models to convert unstructured data into vector embeddings within Oracle databases
- Integrate Oracle AI Vector Search with OCI Generative AI services to enhance natural language processing capabilities
- Analyze the performance and accuracy of vector indexes, and implement best practices to optimize data retrieval
- Evaluate the impact of vector search techniques on the efficiency of AI-driven applications

What You Will Learn

Module 1: Overview of Oracle AI Vector Search

- Objectives
- VECTOR Data Type
- Vector Embeddings
- Similarity Search
- Vector Embedding Models
- Summary

Module 2: Why Use Oracle AI Vector Search

- Objectives
- Benefits of Oracle AI Vector Search
- Examples
- Summary

Module 3: Oracle AI Vector Search Workflow

- Objectives
- Generate Vector Embeddings
- Embedding Generation Examples
- Load ONNX Model into the DB
- Generate the Embedding Example
- Store Vector Embeddings
- Vector Indexes
- Query Data with Similarity Searches
- Complete Workflow
- Summary

Module 4: Running Basic Queries on Vectors

- Objectives
- Basic Queries & Comparison Operations
- Demo
- Summary

Module 5: Vector Indexes and Memory

- Objectives
- Vector Indexes Overview
- Why Vector Indexes are Needed
- Vector Pool
- In-Memory Neighbor Graph Vector Index (HNSW)
- Neighbor Partition Vector Index (IVF)
- Creating a Basic Vector Index & Important Parameters
- Using Vector Indexes & Monitoring Accuracy
- Important Limitations & Best Practices
- Memory Considerations: Vector Storage & In-Memory Indexes
- Sample Calculation
- Summary

Module 6: DML Operations on Vectors

- Objectives
- Create a Table with a Vector Column
- VECTOR Data Type Declaration Formats
- Vector DML Operations
- Vector DML Using SQL*Loader
- Demo
- Summary

Module 7: Vector DDL

- Objectives
- Tables with Different Vector Formats
- Insert Examples
- DDL Operations on Vectors
- Prohibited Operations & VECTOR Data Type Restrictions
- Demo
- Summary

Module 8: Creating and Finding the Nearest Vectors

- Objectives
- Vector Constructor & Examples
- Vector Distance & Operands

Distance Metrics: Euclidean, Euclidean Squared, Cosine, Dot Product, Manhattan, Hamming

- Vector Distance Examples & Shorthand Operators
- Demo
- Summary

Module 9: Finding the Closest Vectors

- Objectives
- Exact Similarity Search (Euclidean & Euclidean Squared)
- Approximate Similarity Search Overview
- HNSW & IVF Indexes
- Multi-Vector Similarity Search
- Demo
- Summary

Module 10: Narrowing Search Results

- Objectives
- Attribute Filtering
- Demo
- Summary

Module 11: Testing Other Distance Functions

- Objectives
- L1_DISTANCE, L2_DISTANCE, COSINE_DISTANCE, INNER_PRODUCT
- Demo
- Summary

Module 12: Testing Other Vector Functions

- Objectives
- Vector Constructors

- Vector Serializers & Norm
- Vector Dimension Count & Format
- Demo
- Summary

Module 13: RAG Overview

- Agenda & Topics
- Vector Data Workflow
- Retrieval Augmented Generation (RAG) Workflow
- Interacting with LLMs
- RAG with Oracle AI Vector Search
- Summary

Module 14: Using Embedding Models with Oracle AI Vector Search

- Agenda & Topics
- Embedding Models Overview
- VECTOR_EMBEDDING()
- Create Table & Vectorize a Table
- Similarity Search
- Change Embedding Models
- Summary

Module 15: Oracle Vector Search & OCI Generative AI Service (Python)

- Agenda & Topics
- OCI Gen Service Overview
- Summary of Steps: Load Sources, Text Chunks, Vectorize, Create and Call
- Summary

Module 16: RAG with Oracle AI Vector Search & OCI Gen AI Service (PL/SQL)

- Agenda & Topics
- Process Overview
- Step-by-Step: Steps 1-5
- Invoke the Chain
- Summary

Module 17: Oracle AI Vector Search Supporting Features

- Agenda & Topics
- Exadata AI Storage
- GoldenGate Microservices for AI
- Distributed AI Processing & Real-Time Vector Hub
- SQL Loader & Data Pump
- Summary

- Agenda & Topics
- Select AI Overview
- Benefits & Challenges
- Using SQL for Select AI Queries
- Summary