

Oracle AI Vector Search Deep Dive ELS

Oracle Database

DURATION

2 Days

MODULES

11 Lectures

COURSE CODE

—

Course Overview

Discover how Oracle Database 23aitransforms natural language questions into secure, actionable insights directly from your data. Master the integration of OCI Generative AI with Autonomous Database to unlock advanced AI-driven capabilities for your organization. You'll also learn to use natural language queries and vector search to streamline data exploration and decision-making processes. By the end of this course, you'll build expertise in retrieval-augmented generation (RAG) and embedding models, enhancing your ability to implement AI solutions. You'll also be able to equip yourself with future-proof skills in AI-powered data management, making you a valuable asset in the evolving tech landscape.

What You Will Learn

Module 1: Course Overview

- Course Overview
- Agenda
- Audience
- Learning Objectives
- Objectives

Module 2: Refresher - Vector & Vector Embeddings

- Vector Data Type
- Vector Embeddings
- Generate Vector Embeddings
- Vector Data Type Review

Module 3: Refresher - Exact Similarity Search

- Similarity Search
- Exact Similarity Search
- Euclidean and Euclidean Squared Distances
- Cosine Similarity

- Dot Product Similarity
- Manhattan Distance
- Hamming Similarity
- Vector Distance Operand

Module 4: Refresher - Approximate & Multi-Vector Similarity Search

- Approximate Similarity Search
- Multi-Vector Similarity Search
- Comparison with Exact Search

Module 5: Refresher - AI Vector Search Fundamentals

- Vector Index Overview
- Why Do We Need Vector Indexes?
- 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

Module 6: RAG (Retrieval Augmented Generation) Overview

- RAG Agenda Topics
- Vector Data Workflow
- RAG Workflow Explanation
- Interacting with LLMs | Complete the RAG Pipeline
- RAG with Oracle AI Vector Search

Module 7: Using Embedding Models with Oracle AI Vector Search

- Embedding Models Overview
- VECTOR_EMBEDDING() Function
- Create Table & Vectorize a Table
- Similarity Search
- Changing Embedding Models

Module 8: Oracle Vector Search & OCI Generative AI Service (Python)

- OCI Gen Service Overview
- Summary of Steps
- Load Sources & Text Chunks
- Vectorize Data
- Create and Call Functions

Module 9: RAG with Oracle AI Vector Search & OCI Gen AI Service (PL/SQL)

- Process Overview
- Step 1: Text Extraction and Preparation
- Step 2: Embedding Models and Vectors
- Step 3: Similarity Search and Response Generation
- Step 4: Build the Prompt
- Step 5: Invoke the Chain

Module 10: Oracle AI Vector Search Supporting Features

- Exadata AI Storage
- GoldenGate Microservices 23ai
- Distributed AI Processing with Vector Replication
- Generative AI with Your Own Business Data
- Real-Time Vector Hub for GenAI
- Actionable AI/ML from Streaming Pipelines
- SQL Loader & Data Pump
- Summary: Why GoldenGate 23ai for AI?

Module 11: Course Wrap-Up

- Agenda Review
- Learning Outcomes