

SQL AI in a Day

Duration: 08 hours

Course Overview

SQL AI in a Day is an intensive, hands-on workshop designed for SQL developers who want to integrate artificial intelligence capabilities directly into their existing SQL workflows — without migrating to a new platform or learning a new language. Across seven guided labs and accompanying presentations, participants explore the landscape of AI-powered search: starting with the limitations of traditional and full-text search, progressing through the deployment of Azure OpenAI embedding models, and culminating in a fully functional Retrieval Augmented Generation (RAG) pipeline inside Azure SQL. By the end of the day, attendees will have built real vector indexes, executed hybrid searches using Reciprocal Rank Fusion (RRF), and generated intelligent, contextual summaries using GPT-4o — all from T-SQL.

Course Agenda

Module 0— Introduction to SQL and AI Integration

- Why Vector Searching?
 - Searching Capabilities
 - Traditional Text Search
 - Full-Text Index Search
 - Vector Search
 - Vector Searching in SQL Database
 - What is AI
 - What are you trying to Solve with AI?
-

Module 1 – Retrieval Augmented Generation (RAG)

- What is Retrieval Augmented Generation (RAG)
- RAG Implementation
- What is RAG & why evaluate it?
- The 6 RAG evaluators

- Precision vs. recall in RAG quality
 - Document retrieval quality metrics
 - RAG - Search Patterns
 - Exploring Text Based Searching
 - Types of Text Based Searching
-

Module 2 – Microsoft Foundry and OpenAI

- What is Azure OpenAI
- What types of Models are there?
- Microsoft Foundry Embedding Models
- Microsoft Foundry Embedding Models w/ Comparison
- Where can I get models from?
- Embeddings Phase
- Deploy Embedding Model

Module 3 – Database Setup and Configuration

- What is an External Model?
- Why Use External Models?
- Creating an External Model
- Limitations & Considerations
- `sp_invoke_external_rest_endpoint`

Module 4 – Generating Embeddings

- Generate Embeddings
- Generate Embeddings – Best Practices
- Keep Embeddings up to Date
- VECTOR INDEX

Module 5 – Vector + Text Searching

- VECTOR_SEARCH

- DiskANN Indexes
- Approximate Nearest Neighbor
- Hybrid Search Methods

Module 6 – Reciprocal Rank Fusion (RRF)

- What is Reciprocal Rank Fusion (RRF)
- RAG – RRF Hybrid Search
- Reciprocal Rank Fusion (RRF)
- Why Reciprocal Rank Fusion (RRF)
- Ranking the Results

Module 7 – Microsoft Foundry – GPT Model

- What is Chat GPT-4o
- Prompts and prompt engineering
- `sp_invoke_external_rest_endpoint`

Module 8 – Security and Governance

- Responsible AI
- Security mechanisms to consider