

Ourbit Vendor-Neutral AI - Day 6: Embeddings and Semantic Search

Day 6 - Embeddings and Semantic Search

Programme: Vendor-Neutral AI - 15-Day Certification Programme

Purpose: Understand how semantic retrieval can find conceptually related information-and why similarity alone is never proof of correctness or authority.

Learning outcomes

- Explain embeddings as numerical representations useful for comparing meaning.
- Contrast semantic search with exact keyword matching.
- Distinguish relevance from authority, freshness and correctness.
- Use metadata and filters to improve retrieval.
- Recognize common semantic-retrieval failure modes.

1. Searching by meaning

Keyword search is strong when the wording is known. Semantic search aims to retrieve material with related meaning even when the words differ. Embeddings support this by representing content numerically so that similarity can be compared mathematically.

2. What similarity does-and does not-mean

A close semantic match means two pieces of content are related according to the representation. It does not prove that a document is authoritative, current, complete or correct. A highly similar obsolete policy can still be the wrong source.

3. Retrieval units matter

Long documents are often divided into smaller retrievable units. Units that are too large may mix unrelated topics; units that are too small may lose necessary context. Titles, section labels and surrounding metadata can help preserve meaning.

4. Combine semantics with metadata

Useful retrieval often combines semantic similarity with filters such as document type, date, jurisdiction, department, language or approval status. This narrows the candidate set before or after similarity ranking and can prevent a relevant-but-invalid document from dominating results.

Worked example - policy search

A user asks, “What is our current refund rule for annual plans?” Semantic search finds a highly similar 2024 policy and a less similar 2026 approved policy. If the task requires the current rule, approval status and effective date are stronger constraints than raw semantic similarity. Retrieval should therefore filter or rerank using those attributes.

Guided practice

Design retrieval for a library containing HR, finance and IT policies across several years. Specify what text would be embedded, which metadata fields you would store, and which filters should apply to a query asking for the current approved remote-work policy.

Knowledge check

- What practical problem do embeddings help solve?
- How does semantic search differ from exact keyword matching?
- Why is similarity not authority?
- What role does chunk size play?
- Name two metadata filters that can improve retrieval quality.

Applied evidence task

A semantic search system repeatedly returns outdated but topically relevant documents. Propose a correction using metadata, filtering or ranking rules. Explain how you would test whether the correction improves retrieval without hiding valid historical material when history is actually requested.

Day 6 takeaway

Semantic retrieval expands search beyond exact wording. Strong systems combine meaning with metadata, authority and time so that “related” does not get mistaken for “right.”

Assessment boundary

This learning evidence is isolated to vendor-neutral-ai-15-day; no vendor-specific embedding service is required.

Learning comes before credentials. Completing this reading alone does not issue a certificate. Programme assessment and practical requirements apply.