Vector Similarity Search

Users expect search functionality in every application and website they encounter. Yet more than 80% of business data is unstructured, stored as text, images, audio, video, or other formats. AI-powered search capabilities are built into Redis Enterprise to enable Vector Similarity Search.

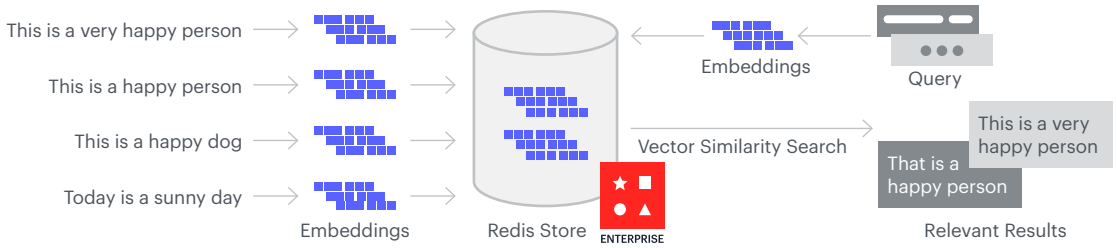

1 **Transform raw data into embeddings.**

Embeddings are numeric representations of unstructured data that capture semantic information, typically created by an AI model.

Image, text, video or even an audio. Where each can be converted into their respective embeddings with the help of Huggingface model Hub.

2 **Store in Redis**

To learn more about embeddings Huggingface model Hub

3 **Redis supports two vector indexing methods.**

FLAT

A brute force approach that searches through all possible vectors.

HNSW

An approximate search that yields faster results with lower accuracy.

4 **Both have the same mandatory parameters.**

- Type: FLOATT32
- DIM: 128
- Distance metric: Cosine

5 **Indexes only need to be created once. They automatically re-index as new hashes are stored in Redis.**

Redis exposes search functionality, combining full text, tag and numeric pre-filters with K Nearest Neighbors (KNN) vector search.

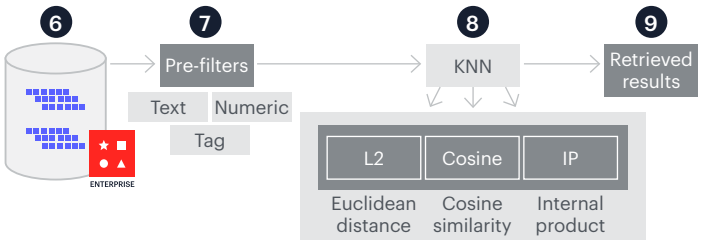

Achieve unmatched performance and scale

- 12x more queries per second (QPS)
- 13x lower latency
- 99.999% available
- And much more

Use cases

- Document retrieval
- Question answering
- Recommendation systems
- Visual search

Use case: Recommendation systems

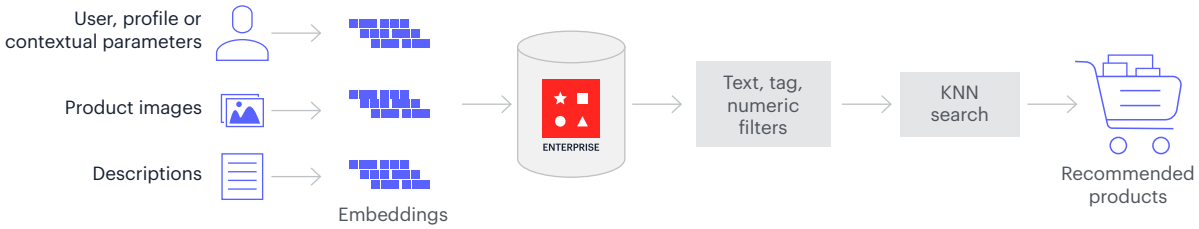

Use case: Question answering

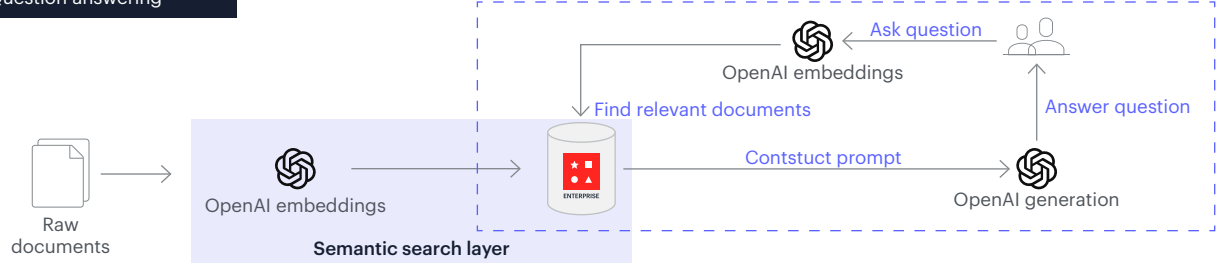

© 2023 Redis