



The Benefits & Uses of AI Semantic Search: A Guide for CTOs

**ONE
BEYOND**



Contents

Introduction	3
What is Semantic Search?	5
The Technologies Behind Semantic Search	7
How Can AI Semantic Search Be Used In Different Industries?	10
How can organisations implement semantic search capabilities?	13

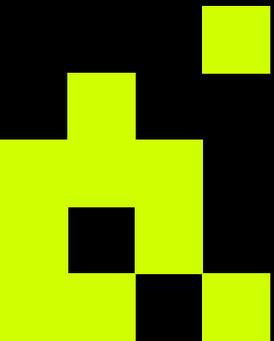
The search box is an always present element of modern software. However, delivering intelligent and relevant search experiences that accurately answer user queries remains a challenge for many organisations.

Traditional keyword-based search engines have their merits, but they often fall short when it comes to understanding the true intent behind user queries. And this is where AI semantic search comes into play.

By leveraging natural language processing (NLP) and machine learning (ML) techniques, semantic search engines aim to decode the underlying meaning and context of queries, serving results that match the user's actual intent, rather than just the literal words used.

In this guide, we'll explore what semantic search really entails, its core technologies, potential business impacts across industries, and architectural considerations for incorporating semantic AI capabilities.

What is Semantic Search?



At its core, semantic search is about mapping user queries to relevant content based on conceptual similarity rather than just lexical matching. Semantic search applies user intent and the meaning of words and phrases to find the right content.

Unlike traditional keyword search that simply looks for exact word matches or predefined synonyms, semantic search takes advantage of AI/ML models to understand the semantics and relationships between words, phrases and concepts. This allows a semantic engine to retrieve highly relevant results even when there is little to no lexical overlap between the query and content. For example, a user searching for 'warm clothing' could be served results about 'sweaters' or 'coats' despite no direct keyword match.

The Technologies Behind Semantic Search

While the concept is powerful, building an AI semantic search engine is complex. It requires multiple interconnected components:

Vector Embeddings

A key technology driving semantic search is vector embeddings - dense vector representations of words/documents in high-dimensional space where semantic similarity is captured by vector proximity.

Modern embedding techniques like Word2Vec, GloVe, BERT are trained on massive text corpora to encode semantics based on co-occurrence patterns and context windows. This allows concepts like 'jumper' and 'knitwear' to have very similar vector representations despite differences in vocabulary.

Retrievers and Scorers

Once texts are embedded as vectors, specialised retriever models like FAISS are used to find the nearest neighbours to the query vector, serving as potential candidates.

These candidates are then scored and re-ranked by more computationally intensive re-rankers or poly-encoders based on deeper semantic understanding of query-document relevance.

Language Models

Large pre-trained language models like GPT-4 have shown remarkable ability to understand and generate human-like text. These can be adapted for tasks like query expansion, document summarisation, improving recall, and even having a conversation to better understand user intents.

While very powerful, actually making semantic search work reliably requires:

- Heavy computational resources for embedding, retrieval and re-ranking
- Careful data preparation for domain-specific model fine-tuning
- Robust model selection, tuning and evaluation processes
- Techniques to interpret, debug and maintain evolving AI models

How Can AI Semantic Search Be Used In Different Industries?

Despite these challenges, the potential upsides of AI semantic search make it extremely beneficial for a majority of industries. Here are some use cases of how AI semantic search can be leveraged:

Ecommerce

For retail & e-commerce brands, AI semantic search can significantly enhance user experience and conversions. By understanding true product intent beyond keywords, semantic AI serves highly relevant product listings even without exact text matches in titles or descriptions. It enables personalised recommendations and search results by mapping user behaviour and preferences. Semantic search also upgrades faceted navigation, comprehending conceptual queries like 'casual summer dresses' rather than just applying rigid filters. Overall, it bridges the gap between natural language and structured data, guiding customers to their desired products through intent-driven experiences.

Healthcare

In healthcare, semantic search maps conversational symptom descriptions to relevant clinical concepts. This allows precisely surfacing cases, treatment plans and research materials to aid doctors during diagnosis and treatment phases. By bridging the gap between natural language and medical terminology, semantic AI streamlines information access, delivering contextualised evidence tailored to each unique patient journey. This intuitive, concept-driven retrieval empowers healthcare professionals to make well-informed decisions efficiently.

Customer Service

Delivering effortless customer service experiences is critical in most organisations. AI semantic search can accurately triage queries and issues by comprehending the underlying context and intent of a customer's support request. This enables intelligent routing of enquiries to the most relevant agent, team or solution by capturing the nuanced semantics, whether it's a billing issue, product question or technical troubleshooting.

Legal & Compliance

AI is transforming legal and compliance work by enabling semantic search across legal documents. This goes beyond keywords, using a conceptual understanding of legal language to find relevant information faster and more accurately. This technology has broader applications, impacting scientific research, content analysis, recruitment, and any field where deep understanding of information is necessary.

How can organisations implement semantic search capabilities?

Semantic search represents a significant shift in how we interact with information. It moves us away from archaic keyword matching and towards a more natural, context-aware interaction model powered by AI. While still in its early stages, this technology is already transforming search experiences across major industries.

But how should organisations integrate semantic search?

Integrating semantic search effectively requires careful consideration. Pre-built cloud solutions and open-source toolkits offer a quick entry point, allowing you to experiment with the technology and assess its potential impact. Open-source options can be particularly valuable for organisations with limited budgets or those wanting to build internal expertise before committing to a larger-scale solution.

However, for organisations with unique data sets or specific search functionalities, custom development might be the optimal choice. While custom solutions require more investment upfront, they offer significant long-term benefits.

A custom-built AI semantic search engine can be integrated with your existing infrastructure and data, leading to superior performance and a search experience tailored to your specific needs.

As a leading software development consultancy, we understand the complexities of implementing semantic search. We work closely with clients to strategically evaluate their needs and guide them towards the most suitable solution, whether it's leveraging pre-built solutions, open-source options, or custom development.

**Get in touch with the team
today let's discuss your
bespoke semantic search
requirements.**

WEBENQUIRIES@ONE-BEYOND.COM

01252 902704

**ONE
BEYOND** 