

Combine index and vector based search results

With the rise of the internet, it is clear that having a good search engine is a must. The reason why Google became Google is that they know how people search. People became used to having a good search engine. Therefore, every website and application needs a good search experience. Imagine having Bol.com or Amazon without a search engine. What if Tripadvisor or Booking.com does not have search functionality? Search engines are incredibly important to find your way around the forest of information and content.

In the past decade (open source), index-based search engines have dominated the search market. Think about engines like Elasticsearch and Solr, which are both Apache Lucene based, but there are alternatives. The terms you use in a query match the terms in the resulting documents. You can trick the engine using techniques like: synonyms, stemming, and fuzziness; however, the technology remains matching terms in an index. With the rise of technologies like word2vec, doc2vec, and image2vec, a new way to find the best matching documents started to appear. Vector-based search algorithms like HNSW are used by vector-based search engines like Weaviate and Qdrant. Now you, as a user, can search for concepts instead of exact terms. You can even search using an image and find similar images. That sounds promising, but there is a but; returning relevant results remains a challenge. We see vector-based search engines include index-based search, and we see traditional index-based search add vector-based search capabilities. Combining the two to get the most relevant results is the biggest challenge facing the world of search.

Goal

Determine the best strategy to combine index and vector-based search engines.

Learning objectives

- Advantages and limitations of the index and vector-based search engines;
- Techniques to evaluate search results and determine the quality of the results;
- Algorithms to compare and/or combine index and vector-based results.

Prerequisite knowledge

- Understand the concepts index based search;
- Understand the basics of vectorizing content;
- Enthusiasm for Natural Language Processing.

Over Luminis

Bij Luminis doen we de dingen net even anders. Dit merk je aan wie we zijn en hoe we werken. Met plezier bijvoorbeeld. Wij zijn er namelijk van overtuigd dat een gedeelde passie en plezier in je werk, leiden tot de mooiste resultaten. Resultaten die we behalen met empathie, vakmanschap en ondernemerschap. Zo voegen we waarde toe voor opdrachtgevers, medewerkers en de maatschappij. Dit doen we door heel Nederland.