



Superlinked

# One Mixture of encoders retrieval vs many rerankers: A modern approach to search and recommendations



**Filip Makraduli** 

ML, Gen AI, LLMs | DevRel | ex-Imperial | Global talent visa UK  
London, England, United Kingdom · [Contact info](#)



Superlinked

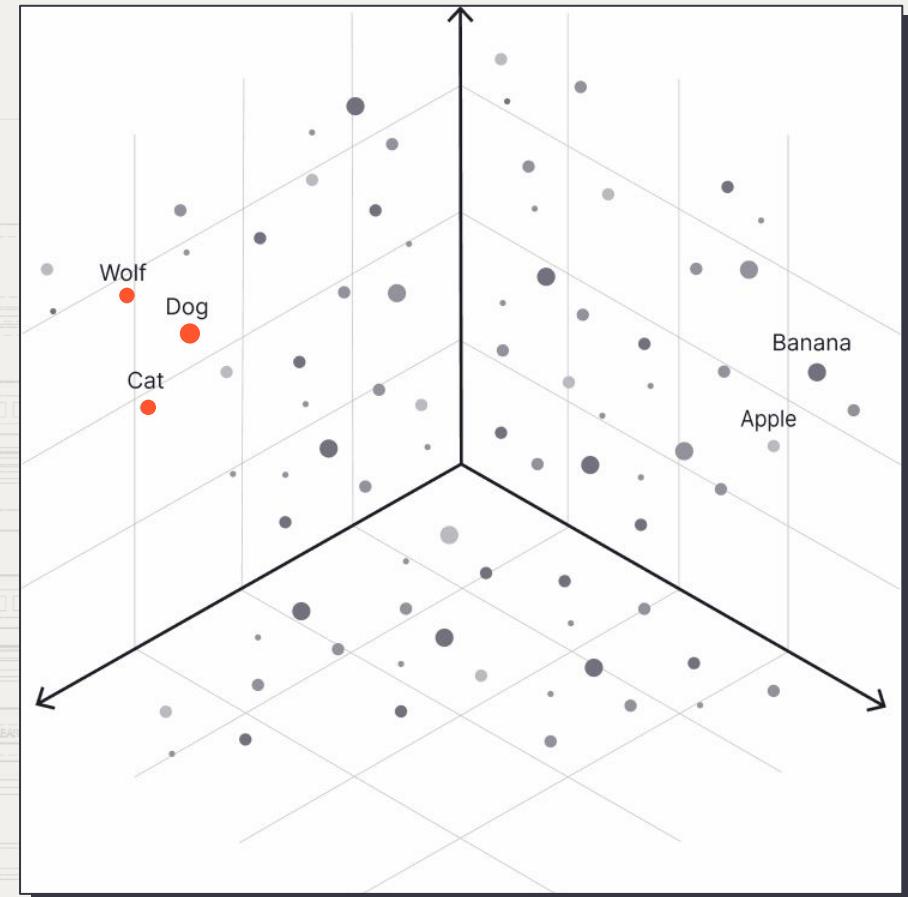
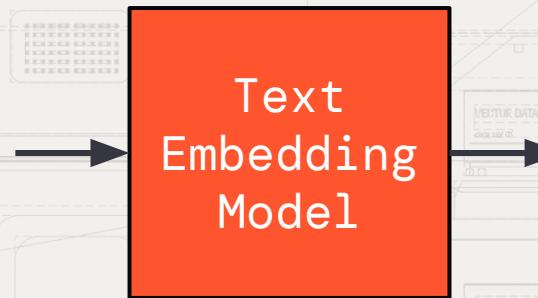


Imperial College London

# Vector based search

## If the data was just text

wolf  
cat  
dog  
banana  
apple





Superlinked

Well ... there are  
problems

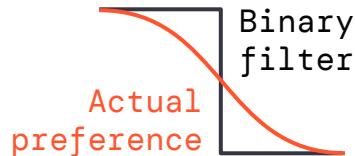
# Real world data is not just text!

```
{  
  rating: 3.7,  
  description: "toothbrush with small head for hard-to-reach areas"  
  category_id: 47  
  price: $29.99  
  total_availability: 10  
  locations: "399 Geary St.", "1059 Hyde St."  
  image: s3://path/to/image.png  
  ... 50 other properties ...  
}
```

With **text-only embeddings**, you have to handle text separately from your other data and/or use **slow, low quality & high cost workarounds\***

:(**Metadata filters**

When you convert a fuzzy preference like "recent", "near" or "popular" into a filter, you model a sigmoid with a binary step function = not enough resolution.

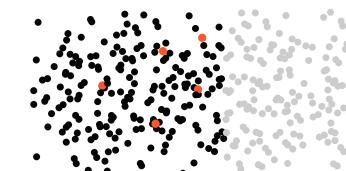


:(**Candidate re-ranking**

Semantic ranking & ColBERT only use text, learn2rank models need ML Engineers.

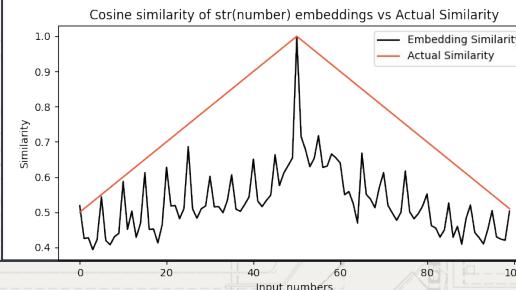
Broad queries eg "popular pants" can't be handled by re-ranking at all, due to poor candidate recall.

Other products  
Pants  
Popular pants



:(**Stringify & vectorize**

It's extremely noisy to stringify & text-embed whole objects like users, SKUs or content+metadata. See below for `cosine(textembed(str(50)), textembed(str(1..100)))`:



\* additionally making more advanced topics like feedback loops, personalization, explainability and control inaccessible



# "Sparkly dress for a Christmas party under 100"

## Without Superlinked

Search results for sparkly dress for a christmas party under 100

Brand  Colour  Size  Product Type  Categories

1 - 48 of 433 items

Relevance

£45 AUTOPHAP Pure Extra Fine Merino Wool Half Zip Jumper  
4.5★★★★★ 508 reviews

£20 AUTOPHAP Pure Supima® Cotton Crew Neck T-shirt  
4.5★★★★★ 703 reviews

£20 AUTOPHAP Pure Supima® Cotton Crew Neck T-shirt  
4.5★★★★★ 703 reviews

£20 AUTOPHAP Pure Supima® Cotton Crew Neck T-shirt  
4.5★★★★★ 703 reviews

Black Asymmetric Sequin Dress  
Mango \$49

Ivory Sequin Shoulder Pad Dress  
Mango \$49

Purple Velvet Crossed Dress  
Mango \$35

Rose Gold Sequin Strap Top  
Mango \$29

Khaki Velvet Knot Dress  
Mango \$45

## With Superlinked

Superlinked Store

Women Men Sports & Leisure Home Beauty Kids Christmas Outdoor

Search Results for "sparkly dress for a christmas party under 100"

► Superlinked parameters  
Showing 100 most relevant products for your query

Rose Gold Sequin Strap Top  
Mango \$29

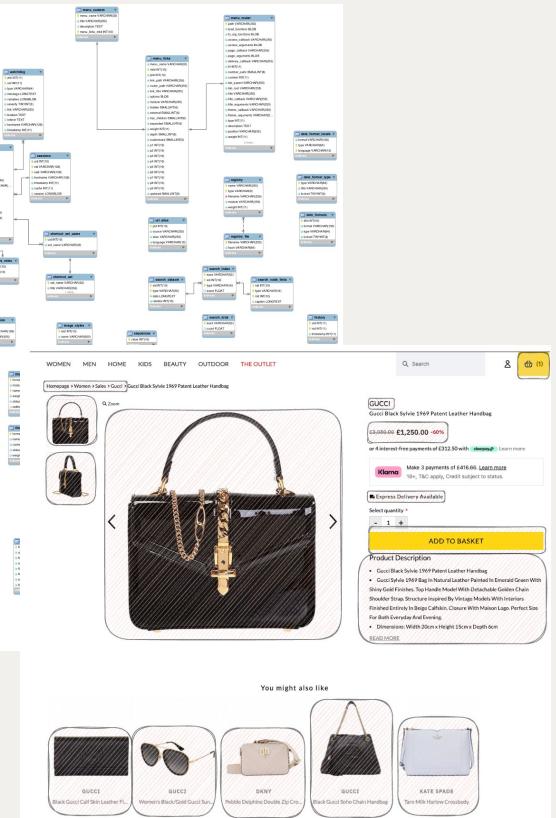
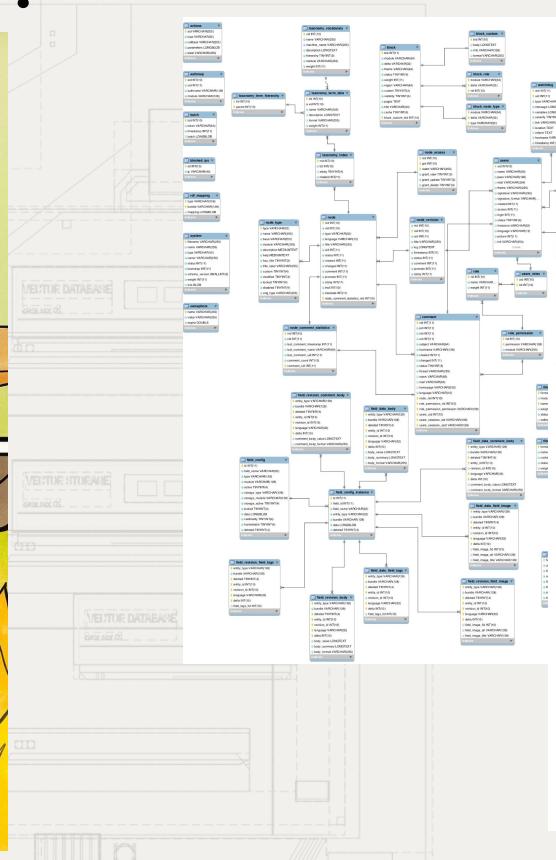
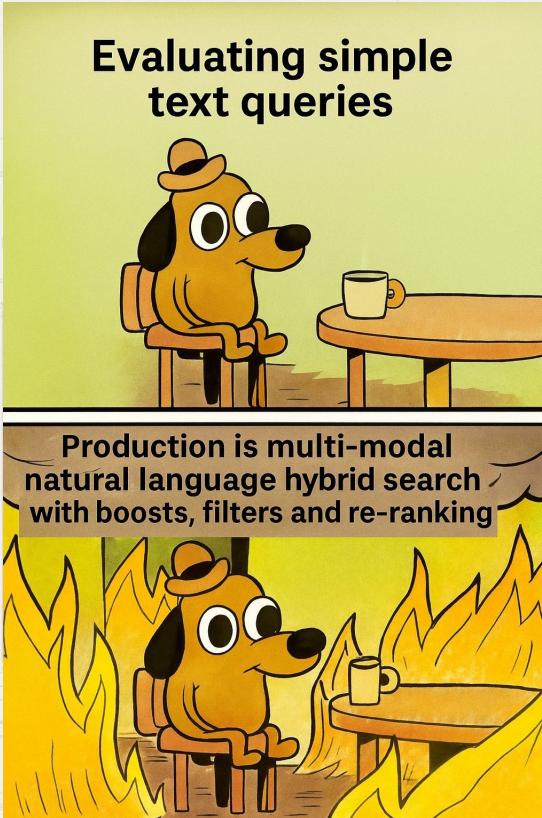
Purple Velvet Crossed Dress  
Mango \$35

Khaki Velvet Knot Dress  
Mango \$45

Ivory Sequin Shoulder Pad Dress  
Mango \$49

Black Asymmetric Sequin Dress  
Mango \$49

Production data is complex - not just a simple text query !





Superlinked

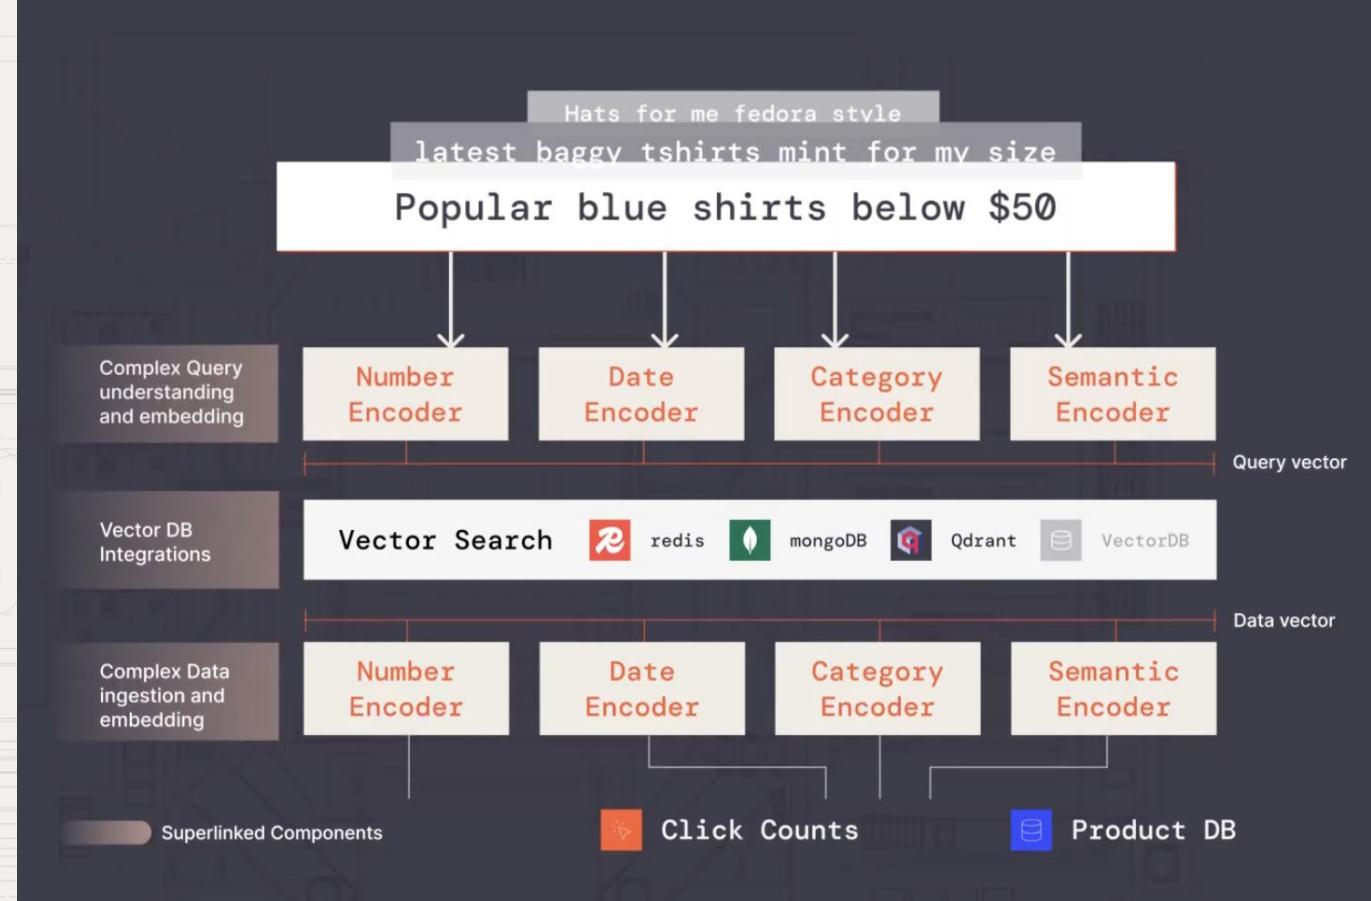
# Introducing mixture of encoders

# Mixture of encoders - a whole lifecycle product

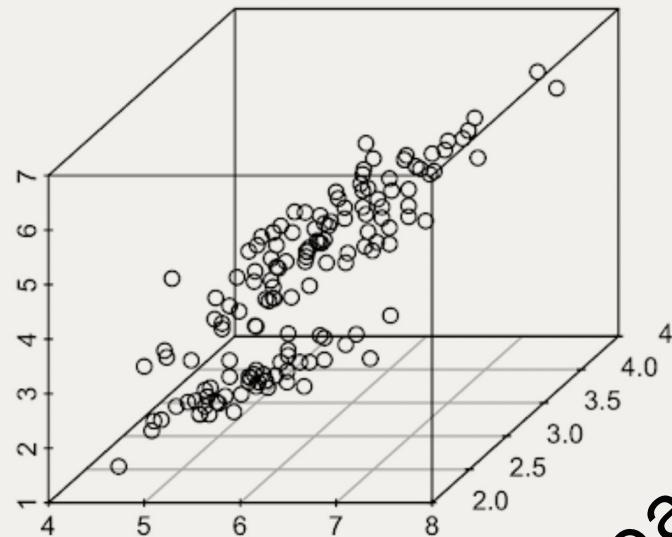
Superlinked's solution

## Benefits:

- Weights
- Explainability
- Business logic



# Map data to a multi dimensional space



Popular

Near to Berlin



Superlinked

Superlinked in  
Production

# 77% Increase in Conversion Rate

## Production Case Study

### Customer

**BRANDALLEY**  
THE NON-STOP DESIGNER SALE

**\$150M+** in annual sales    **5M+** monthly users

### Results

**+77%**  
Conversion  
rate

**+68%**  
Average order  
value (AOV)

**200ms**  
95th percentile  
latency

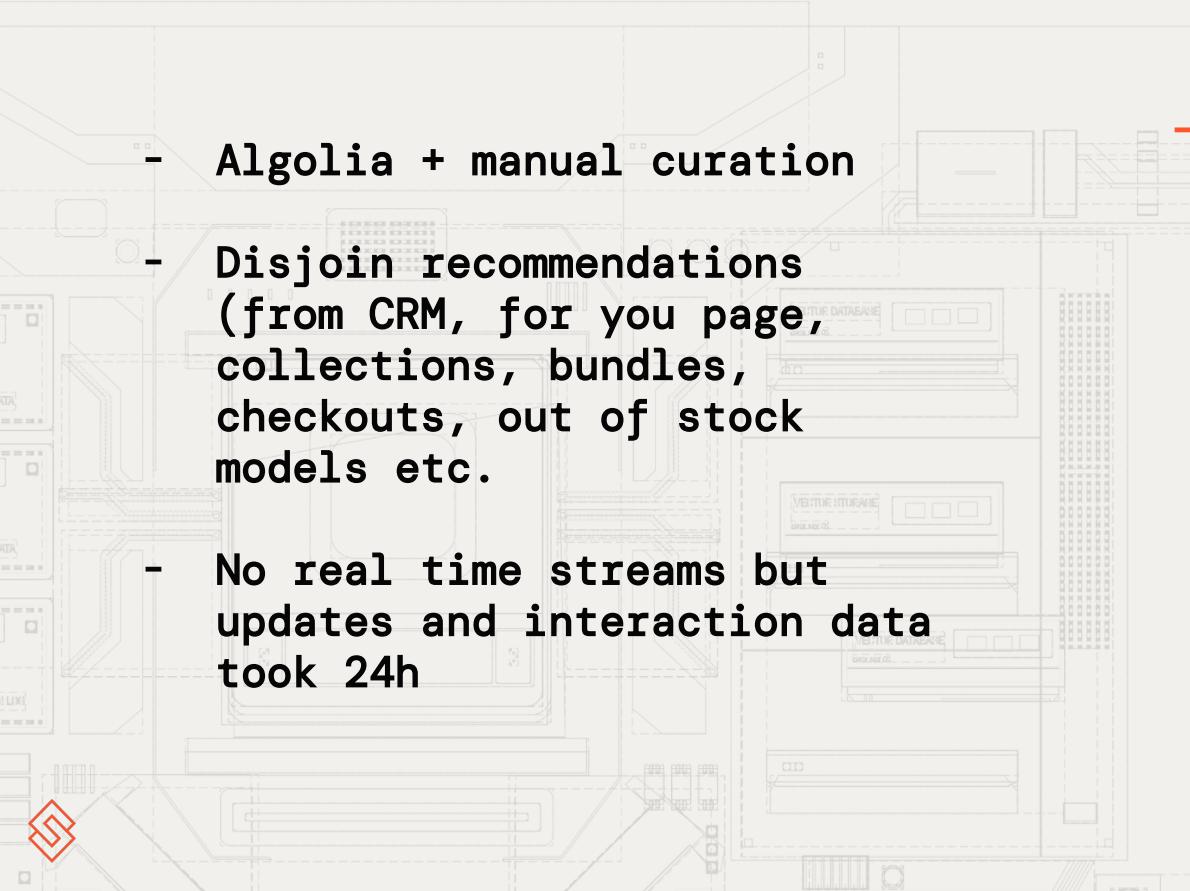
# Before

vs

# After

- Algolia + manual curation
- Disjoin recommendations (from CRM, for you page, collections, bundles, checkouts, out of stock models etc.
- No real time streams but updates and interaction data took 24h

Superlinked mixture of encoders framework





Superlinked

Open source  
Repo + evals

Test our claims, maybe  
you will help us find  
gaps !

# Information Retrieval benchmarks - How close are they to your business problem ?

## Massive Text Embedding Benchmark (MTEB)

- Text; Image is a recent addition
- A blended score of various metrics
- Retrieval: tops out at  $\approx 0.6$
- **NO METADATA**



## Benchmarking-IR (BEIR)

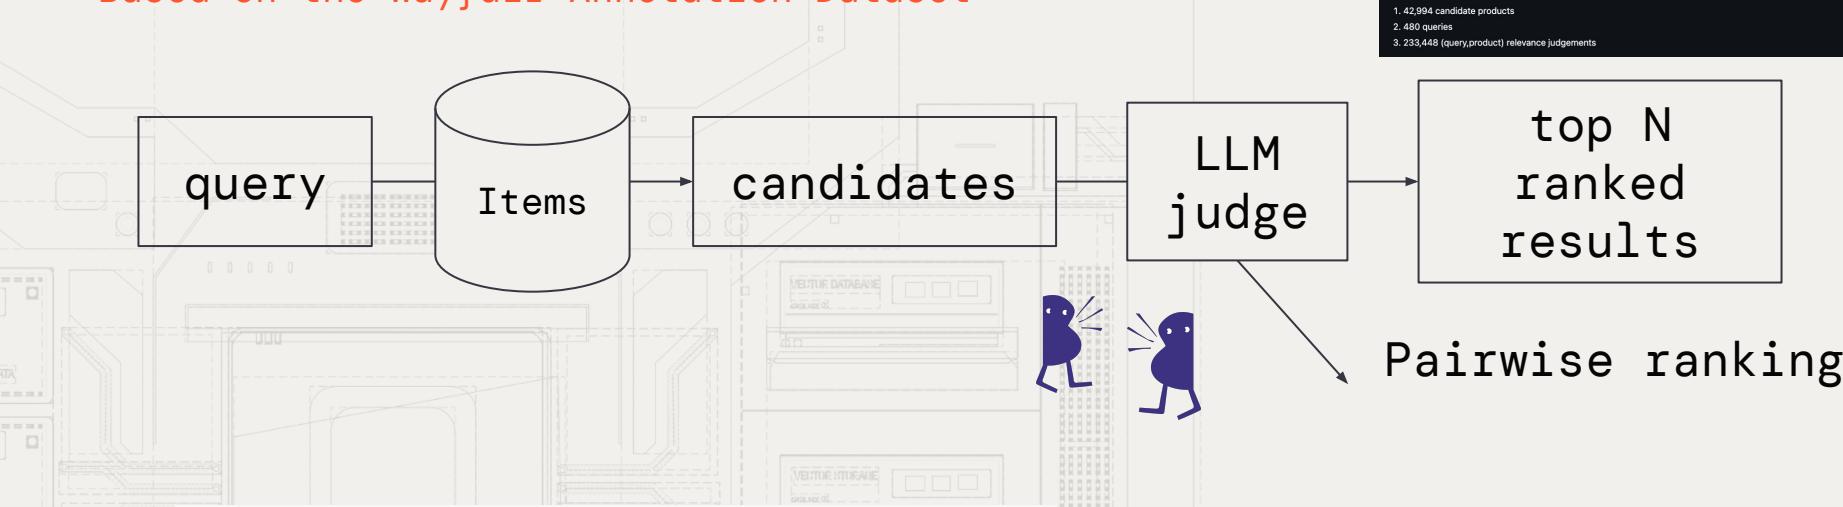
- Text (short- a few ¶!); no images
- Relies heavily on synthetic queries
- No negative examples to teach model what irrelevant results look like
- **NO METADATA**



**Beir**  
Benchmarking IR

# A Superlinked benchmark (coming soon!)

Based on the Wayfair Annotation Dataset



Articles / Airbnb search benchmarking

Publication Date: June 15, 2025 | #Personalized Search #Superlinked

Airbnb Search Benchmarking -  
Comparison of retrieval  
techniques

Superlinked  
Vs  
BM25, ColBERT,  
cross encoders

## WANDS - Wayfair Annotation Dataset

CCS Template 0.3.5 Contributor Covenant 2.0

### About The Project

WANDS is a Wayfair product search relevance dataset that is published as a companion to the paper from ECIR 2022:

WANDS: Dataset for Product Search Relevance Assessment

Yan Chen, Shujian Liu, Zheng Liu, Welyi Sun, Linas Baltrunas and Benjamin Schroeder

The dataset allows objective benchmarking and evaluation of search engines on an E-Commerce dataset. Key features of this dataset includes:

1. 42,994 candidate products
2. 480 queries
3. 233,448 (query,product) relevance judgements

# Open source repo

For AI engineers



Open source

Github repo



README Apache-2.0 license

 Superlinked

[CHAT WITH OUR AI DOCS](#)

[Docs](#) [Ask DeepWiki](#) [pypi v28.3.1](#) [Stars 1.2k](#) [last commit last friday](#) [license Apache-2.0](#)

*Superlinked is a Python framework for AI Engineers building high-performance search & recommendation applications that combine structured and unstructured data.*

## Quickstart

```
%pip install superlinked
```

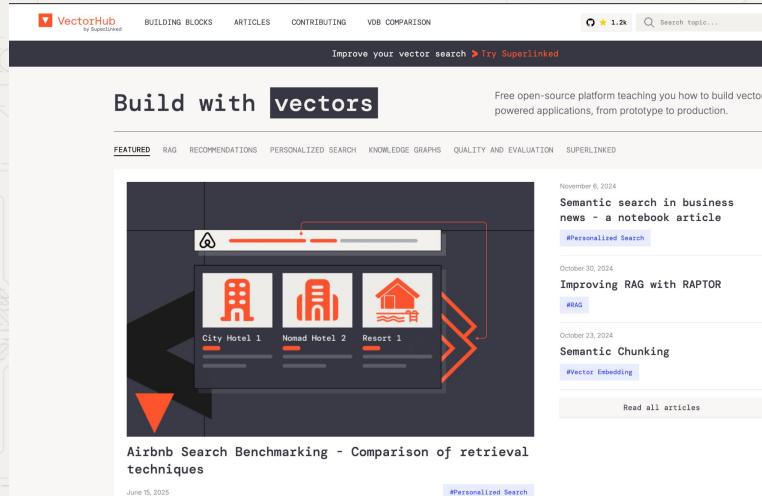
```
from superlinked import framework as sl

# Define schema for movie reviews
class Review(sl.Schema):
    id: sl.IdField
    text: sl.String

review = Review()
space = sl.TextSimilaritySpace(text=review.text, model="all-MiniLM-L6-v2")
index = sl.Index(space)
query = sl.Query(index).find(review).similar(space, sl.Param("search")).select_all()
```

# Even more open source resources

Lots of support in [docs](#) and  
in-depth tech [articles](#)



VectorHub by Superlinked

BUILDING BLOCKS ARTICLES CONTRIBUTING VDB COMPARISON

1.2k Search topic...

Improve your vector search Try Superlinked

### Build with vectors

Free open-source platform teaching you how to build vector-powered applications, from prototype to production.

FEATURED

- RAG RECOMMENDATIONS PERSONALIZED SEARCH KNOWLEDGE GRAPHS QUALITY AND EVALUATION SUPERLINKED

November 2024 Semantic search in business news - a notebook article #Personalized Search

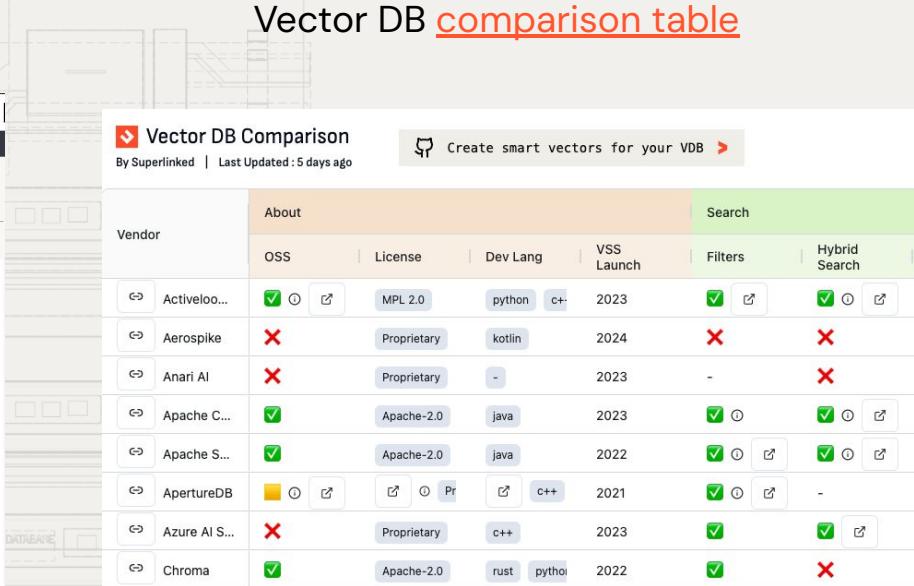
October 30, 2024 Improving RAG with RAPTOR #RAG

October 23, 2024 Semantic Chunking #Vector Embedding

Read all articles

Airbnb Search Benchmarking - Comparison of retrieval techniques

June 15, 2025 #Personalized Search



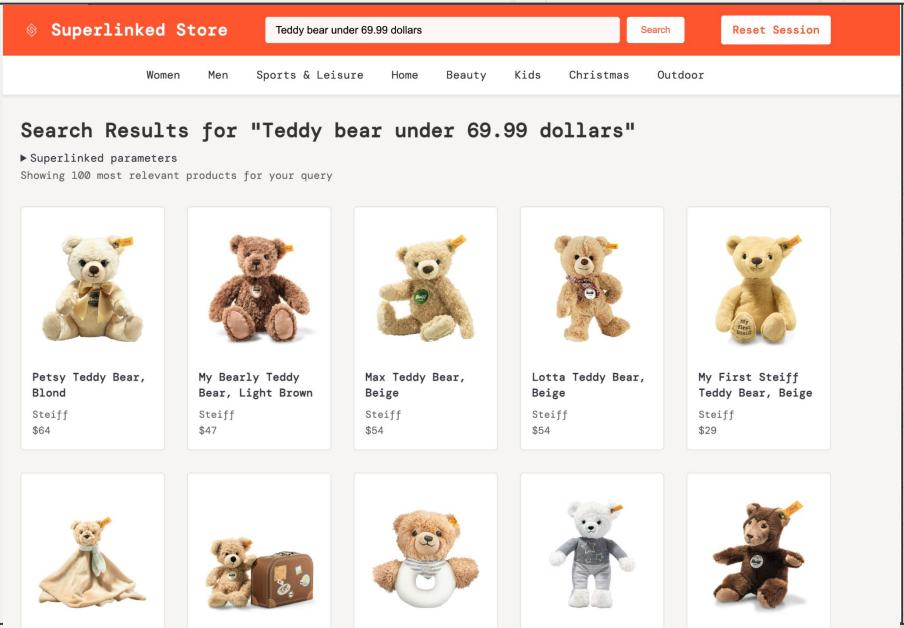
### Vector DB Comparison

By Superlinked | Last Updated: 5 days ago

Create smart vectors for your VDB >

Vendor	About					Search		
	OSS	License	Dev Lang	VSS Launch	Filters	Hybrid Search		
Activeloo...	✓	MPL 2.0	python, c++	2023	✓	✓		
Aerospike	✗	Proprietary	kotlin	2024	✗	✗		
Anari AI	✗	Proprietary	-	2023	-	✗		
Apache C...	✓	Apache-2.0	java	2023	✓	✓		
Apache S...	✓	Apache-2.0	java	2022	✓	✓		
ApertureDB	✗	Pr	c++	2021	✓	✓		
Azure AI S...	✗	Proprietary	c++	2023	✓	✓		
Chroma	✓	Apache-2.0	rust, python	2022	✓	✗		

# Create your own real time recs store - demo



The screenshot shows a search results page for "Teddy bear under 69.99 dollars" on the Superlinked Store. The top navigation bar includes a logo, a search bar with the query, a "Search" button, and a "Reset Session" button. Below the navigation are category links: Women, Men, Sports & Leisure, Home, Beauty, Kids, Christmas, and Outdoor. The main content area is titled "Search Results for 'Teddy bear under 69.99 dollars'" and displays 100 most relevant products. Each product card includes an image, the name, brand, color, and price. The products shown are:

Product	Brand	Color	Price
Petsy Teddy Bear, Blond	Steiff	Blond	\$64
My Bearly Teddy Bear, Light Brown	Steiff	Light Brown	\$47
Max Teddy Bear, Beige	Steiff	Beige	\$54
Lotta Teddy Bear, Beige	Steiff	Beige	\$54
My First Steiff Teddy Bear, Beige	Steiff	Beige	\$29

Below the first row of cards, there are two more rows of five cards each, showing additional teddy bear products.

Check it out:

<https://recsys-nlq-demo.superlinked.io/>

Learn how to build this yourself:

<https://links.superlinked.com/docs-tutorial>



Superlinked

# Thank You

## Questions?

