

Evaluating Embedding Models for Vector Search in E-commerce

Haystack EU

September 24, 2025

Philippe Bouzaglou




AI Foundation Model

Specifically for E-Commerce Search

Embeddings

A decorative graphic in the bottom right corner showing a network of glowing blue nodes connected by thin lines, resembling a data or neural network structure.

Embeddings are...



Embeddings are...

ABSTRACT

[0.461, 0.686, 0.324, ... , 0.891]

[0.639, 0.721, 0.442, ... , 0.278]

A decorative graphic in the bottom right corner consisting of a network of glowing blue nodes connected by thin lines, resembling a neural network or data visualization.

Embeddings Models are...

IMPORTANT

- **nearly all vector search relevance relies on them**

Embedding Models are...

HARD TO EVALUATE

- all-MiniLM-L6-v2
- bge-large-en-v1.5
- e5-large-v2
- ???????

Current evaluation method # 1

- **finger in the wind**
- **use default model**
- **Quick, very un-scientific**
- **Bad results discourage from vector search**

Current evaluation method # 2

- **install a model (GPU?)**
- **embed catalog**
- **run queries**
- **measure**
- **repeat (if time permits)**

- **Very precise, but slooooooow**

We need a better method !!

- **Reasonably good results**
- **Iterate quickly**

Product-Query Similarity Evaluation

- 1 product
- many queries
- similarity score for each query

Step 1: Choose 1 product / document



Step 2: Write or obtain queries



```
"Bohemian midi dress print ruffle hem",  
"Boho print dress",  
"Relaxed fit beach dress",  
"Vintage outfit",  
"Women's casualwear",  
"Black dress with print",  
"Business suit with tailored fit",  
"Men's hiking shoes"
```


Step 3: Calculate cosine similarity scores between each query and the product



Bohemian midi dress print ruffle hem	0.945234
Boho print dress	0.883245
Relaxed fit beach dress	0.692385
Vintage outfit	0.594385
Women's casualwear	0.448571
Black dress with print	0.814219
Business suit with tailored fit	0.098211
Men's hiking shoes	0.004987

✓
vector

Step 4: Rank queries by similarity to product



Bohemian midi dress print ruffle hem	1
Boho print dress	2
Relaxed fit beach dress	4
Vintage outfit	5
Women's casualwear	6
Black dress with print	3
Business suit with tailored fit	7
Men's hiking shoes	8

Step 5: Ground Truth: Use an LLM (or human) to rank queries

Bohemian midi dress print ruffle hem	1
Boho print dress	2
Relaxed fit beach dress	3
Vintage outfit	4
Women's casualwear	5
Black dress with print	6
Business suit with tailored fit	7
Men's hiking shoes	8

Step 6: Compare



	GPT	Model
Bohemian midi dress print ruffle hem	1	1
Boho print dress	2	2
Relaxed fit beach dress	3	4
Vintage outfit	4	5
Women's casualwear	5	6
Black dress with print	6	3
Business suit with tailored fit	7	7
Men's hiking shoes	8	8

Step 6: Compare across models



	Alibaba-NLP_gte-Qwen2-1.5B-instruct																												Alibaba-NLP_gte-Qwen2-7B-instruct																												BAAI_bge-base-en-v1.5																												BAAI_bge-large-en-v1.5																												Qwen_Qwen3-Embedding-0.6B																												Snowflake_snowflake-arctic-embed-1																												ibm-granite_snowflake-arctic-embed-m																												infrad_granite-embedding-m-long																												infloat_stella-base-en-v2																												infloat_multilingual-e5-large																												mixedbread-ai_nomic-embed-large-instruct																												nomic-ai_nomic-embed-large-v1																												sentence-transformers_text-v1.5																												sentence-transformers_all-MiniLM-L6-v2																												average_rank																												4.1 GPT																												4.5 GPT																												4.5 GPT (Judge)																												Vecra																												o3 GPT																												Gemini 2.5 Pro																												Claude 4 Sonnet																												LLM Average																																																																																																																																																																							
query_text	13	3	29	9	3	13	4	6	40	25	29	45	17	3	1	14	2	4	23	7	44	15	4	3	1	2	3	2	1	2	41	52	10	1	14	10	5	5	8	24	1	3	21	10	5	4	1	1	2	25	1	30	12	2	2	2	7	5	3	3	3	12	2	7	4	20	3	11	6	2	2	26	4	19	3	1	11	5	7	14	21	4	32	10	1	1	3	1	1	1	2	1	10	5	36	12	23	5	2	1	21	14	39	7	8	25	4	2	8	3	6	31	38	42	16	14	7	4	17	8	8	10	10	2	4	28	5	15	34	3	11	4	41	32	38	29	11	13	14	7	18	20	38	25	26	19	9	8	5	13	14	9	5	9	4	10	20	6	31	22	21	18	40	13	23	2	15	15	15	7	6	12	10	12	9	25	15	16	4	6	34	25	14	8	15	18	8	34	22	28	30	12	16	11	18	11	23	12	26	11	21	29	16	27	6	33	14	19	13	9	7	9	7	11	11	10	3	14	16	8	9	8	1	2	5	21	22	12	22	9	12	10	15	20	16	17	16	40	14	7	6	8	3	15	6	4	7	24	12	22	20	2	25	28	9	9	34	13	25	17	24	29	23	20	10	8	22	40	19	20	3	11	9	11	11	18	9	10	22	30	1	11	6	6	10	22	12	22	17	6	2	4	14	3	11	5	1	8	30	12	12	6	5	10	16	17	7	14	11	20	22	24	30	1	19	7	4	20	15	18	16	35	12	25	8	38	15	19	45	42	1	20	15	13	11	6	13	10	13	12	14	13	26	27	11	16	15	16	12	12	33	24	28	32	37	10	22	24	32	2	48	22	23	15	12	24	16	24	20	19	14	13	26	27	32	32	20	35	14	7	8	19	30	20	17	31	23	31	38	9	12	15	22	22	14	13	32	10	20	22	19	38	53	3	10	17	28	36	13	26	5	14	37	46	1	39	15	16	25	18	13	17	27	23	26	18	14	29	22	28	25	23	30	36	8	25	40	15	25	31	15	29	24	8	6	6	42	33	13	38	30	11	11	10	22	24	20	15	33	37	31	26	27																																																																																																																																																																																																																																																																																																																																	
Bohemian midi dress ethnic print ruffle hem																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																

Come try it out on your data!

Future of Search Conference

Discover what's next in Search & AI

 **17** September 25, 2025 • Berlin, Germany

berlinsearchweek.com



What makes a good embedding model for e-commerce search?

Embeddings Models are...

TRAINED ON LANGUAGE NOT PRODUCTS

 vectra

 vectra

An AI Foundation Model

Specifically for E-Commerce Search

Embeddings



Pre-Trained Model


Generative
Pre-trained
Transformer

World Knowledge

- has seen most products
- in different contexts
- with comments




Optimized for Search

- **Tiny vectors (256 dim)**
 - **SPEED (5ms median)**
 - **Tiny size (2.8M products in 4GB)**
 - **Separate ultra-fast query embedder**
- 
- An abstract graphic in the bottom right corner consisting of a network of glowing blue nodes connected by thin lines, resembling a data network or a neural network structure.

What is different?

Large Training Data (1 Trillion Tokens)



Quality comparison



Incumbent models

["shoe", "men's", "red"]


```
{
  "product_type": "men's loafers",
  "style": "slip-on dress shoes",
  "color": {
    "primary": "burgundy/wine red",
    "material_appearance": "velvet",
    "finish": "matte with subtle sheen"
  },
  "upper_details": {
    "material": "velvet fabric",
    "pattern": {
      "type": "quilted",
      "description": "diamond-shaped quilting pattern covering entire upper",
      "texture": "raised, padded appearance creating dimensional effect"
    },
    "construction": "slip-on design with no laces or buckles"
  },
  "embellishment": {
    "type": "metal medallion",
    "location": "center front of vamp",
    "design": "ornate crest or shield-shaped emblem",
    "color": "silver/pewter tone",
    "details": "appears to feature intricate engraving or embossed design"
  },
  "sole": {
    "visible_portion": "black",
    "material_appearance": "leather or synthetic",
    "profile": "low-profile dress shoe sole"
  },
  "interior": {
    "lining_color": "black",
    "visible_branding": "partial text visible inside shoe"
  },
  "overall_aesthetic": {
    "style_category": "luxury/formal",
    "occasion": "formal events, evening wear, special occasions",
    "design_influence": "European luxury slipper style"
  },
}
```


Shirt for kids with a bear wearing sunglasses


AI Model **vectra**

Next Example [Search](#)


Click to see this search on [zalando](#) [ASOS](#) [OTTO](#)




STEIFF
S15249.06S-JC11




COTTON ON
C1Q249.02A-E12



S.OLIVER
SO2269.03N-M11



WATAPPAREL
W00249.04I-JC11



COTTON ON
C1Q249.01X-JC11

Quick Start

Hard to find items

Difficult queries



Bottom half of the purchase intents

- often ignored
- equal size and revenue to top half
- too complex to deal with

Implementation steps

- **Append vector search result to the end**
- **Bubbles up on difficult queries**

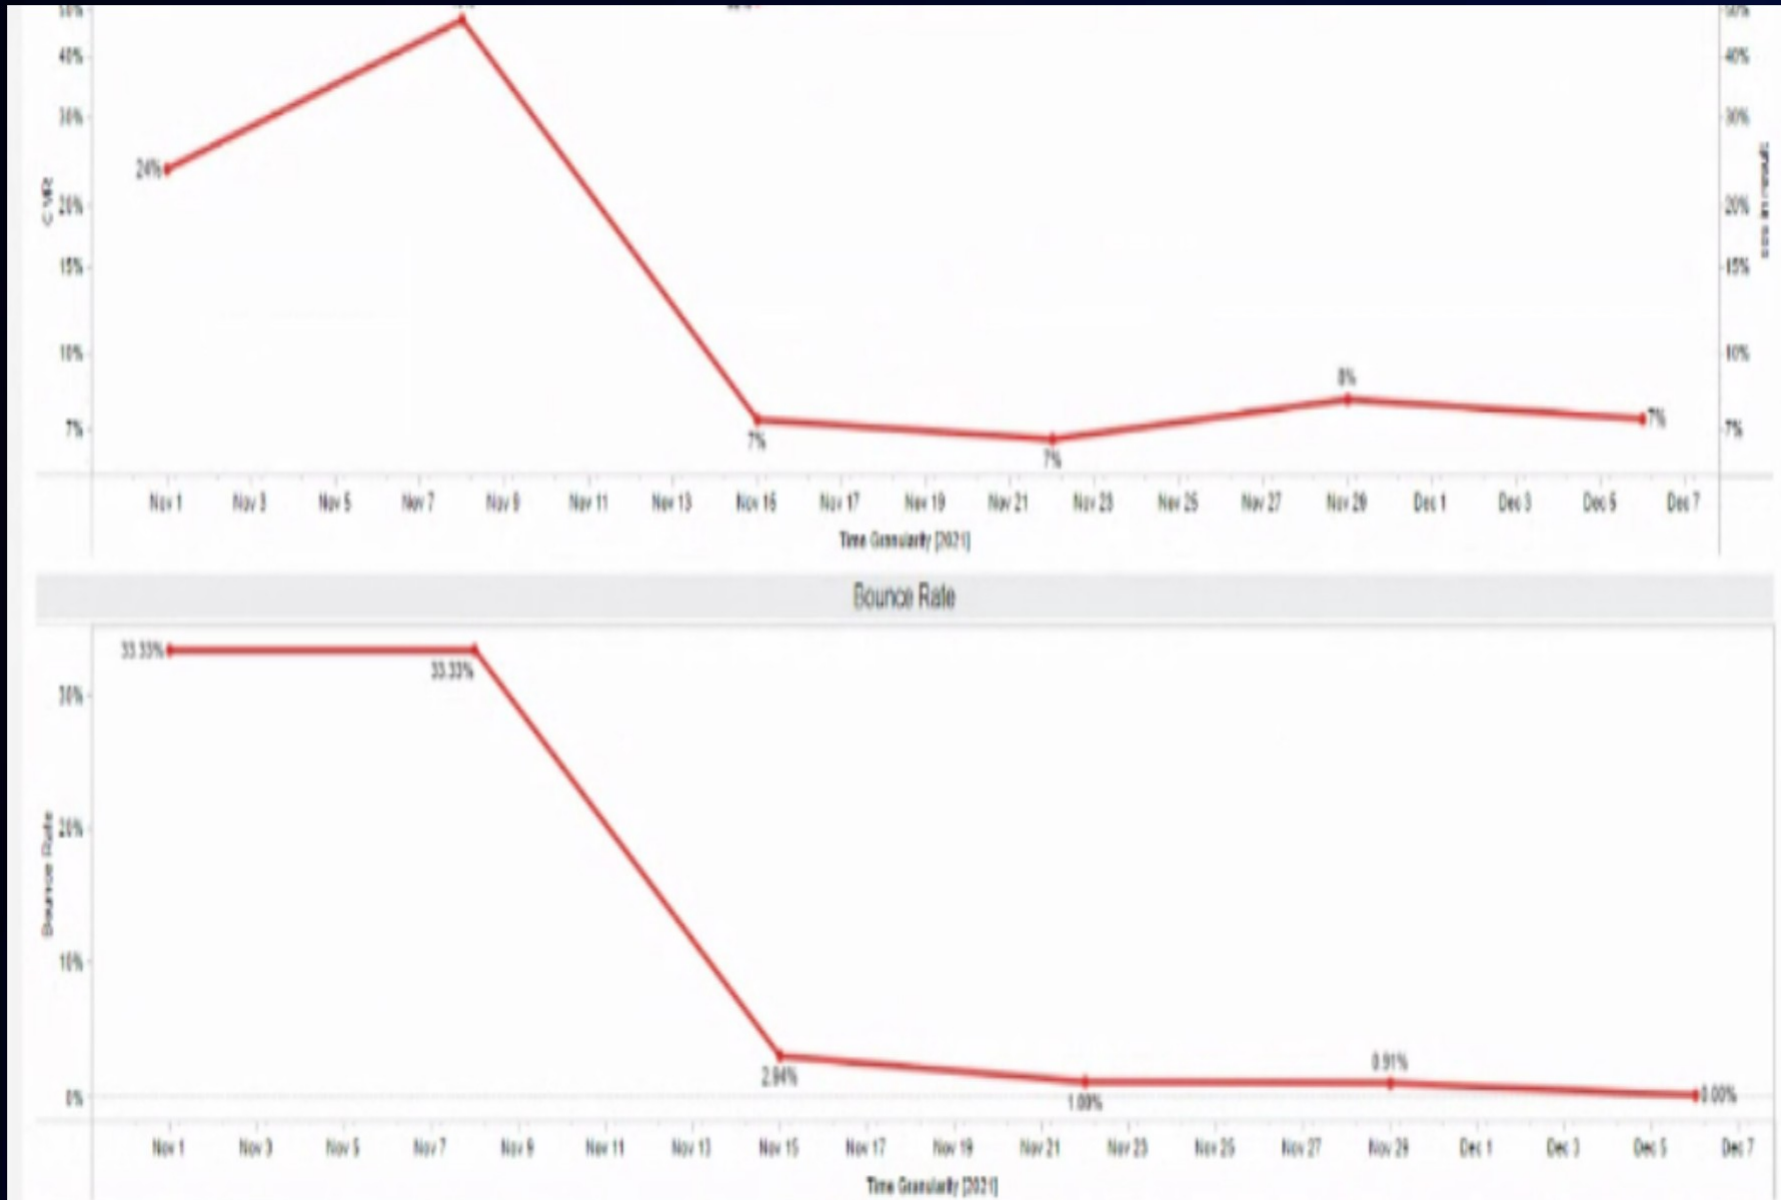
More Advanced Implementation

- **Score threshold**
- **Move threshold until optimum found**

Results Example



Search Abandonment



Get started !

- request access
- 5x cost in incremental sales **GUARANTEED**

Evaluation framework available

- **Join workshop tomorrow**
- **API with 22 models to benchmark against**

Come try it out

- **Workshop tomorrow**
- **FREE Sponsored tickets still available**





Try out our fake shop



- <https://zalando.mariosbox.com>
- User: esattendee
- Pass: letmein

Get In Touch

 Philippe Bouzaglou

