



Search & AI: a new era

David Pilato | [@dadoonet](https://twitter.com/dadoonet)

Agenda

- "Classic" search and its limitations
- ML model and usage
- Vector search or hybrid search in Elasticsearch
- OpenAI's ChatGPT or LLMs with Elasticsearch

Elasticsearch

You Know, for Search



Elasticsearch

lucene

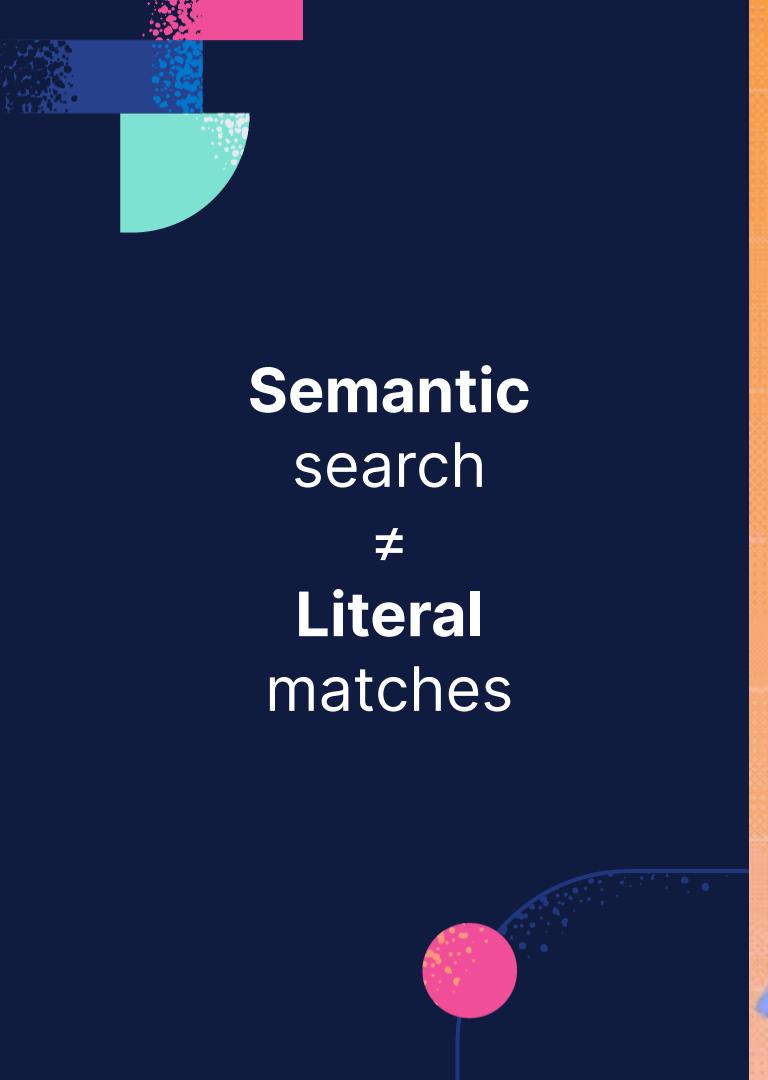
66

These are not the droids
you are looking for.

```
GET /_analyze
{
  "char_filter": [ "html_strip" ],
  "tokenizer": "standard",
  "filter": [ "lowercase", "stop", "snowball" ],
  "text": "These are <em>not</em> the droids
          you are looking for."
}
```

These are *not* the **droids you** are **looking** for.

```
{ "tokens": [ {  
    "token": "droid",  
    "start_offset": 27, "end_offset": 33,  
    "type": "<ALPHANUM>", "position": 4  
} , {  
    "token": "you",  
    "start_offset": 34, "end_offset": 37,  
    "type": "<ALPHANUM>", "position": 5  
} , {  
    "token": "look",  
    "start_offset": 42, "end_offset": 49,  
    "type": "<ALPHANUM>", "position": 7  
} ] }
```



Semantic
search
≠
Literal
matches

similarweb

YOU'RE COMPARING APPLES TO NECTARINES





TODAY



X-wing starfighter squadron

TOMORROW



*What ships and crews do I need to destroy an almost finished death star?
Or is there a secret weakness?*

Elasticsearch

You Know, for **Vector** Search

What is a **Vector** ?



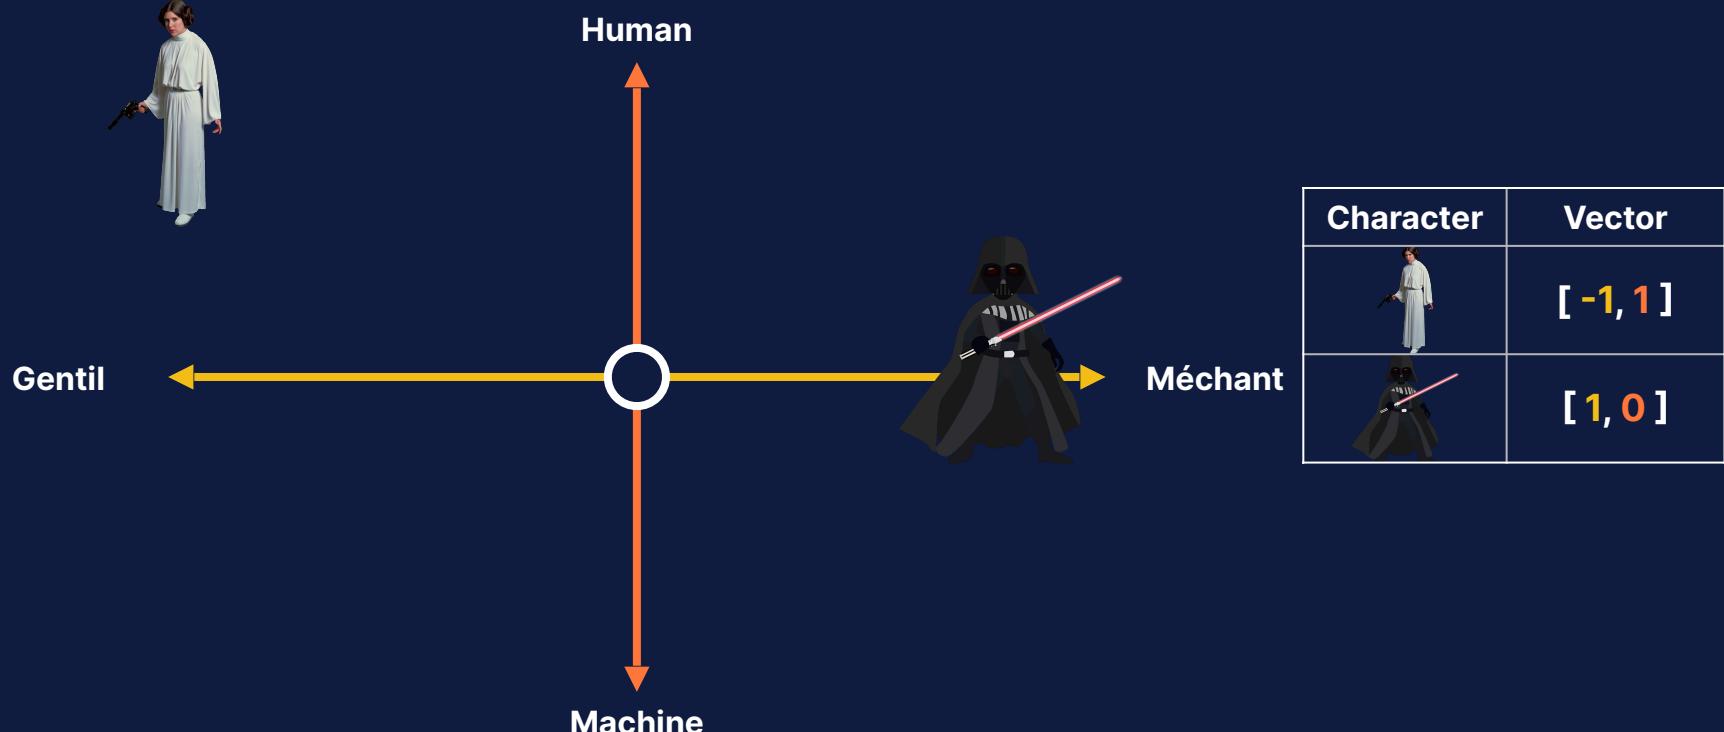
Embeddings represent your data

Example: 1-dimensional vector

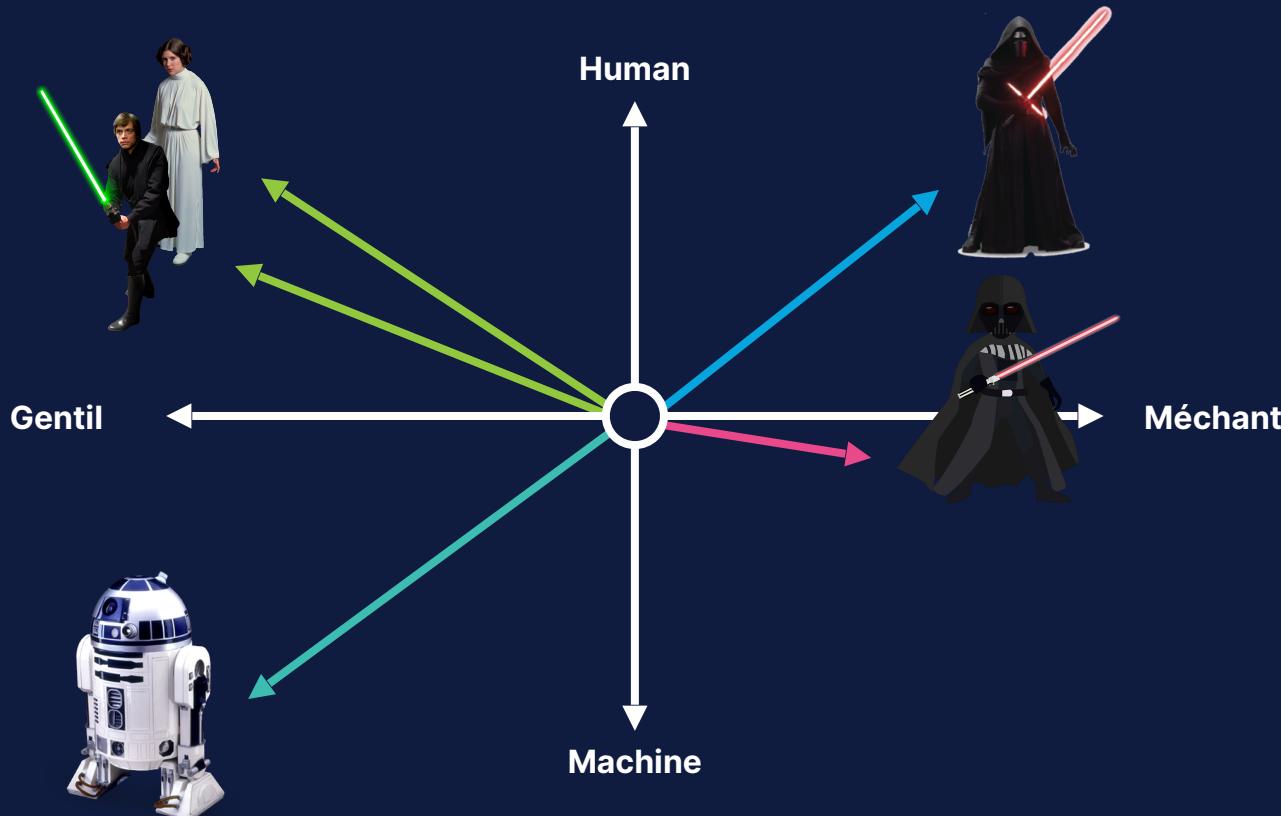


Character	Vector
	$[-1]$
	$[1]$

Multiple dimensions represent different data aspects

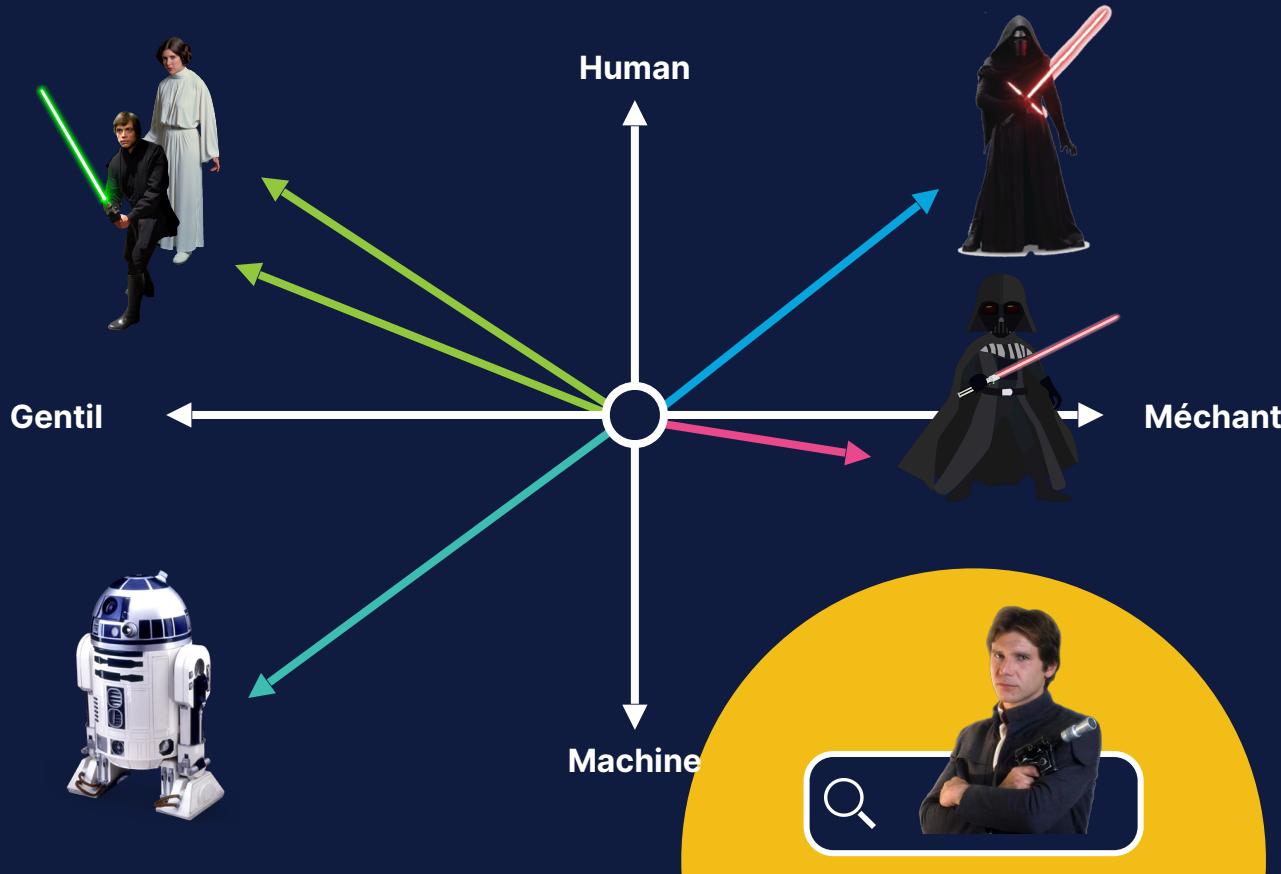


Similar data is grouped together



Character	Vector
Luke Skywalker and Princess Leia	$[-1.0, 1.0]$
Emperor Palpatine	$[1.0, 0.0]$
Darth Vader	$[-1.0, 0.8]$
R2-D2	$[1.0, 1.0]$
	$[-1.0, -1.0]$

Vector search ranks objects by similarity (~relevance) to the query



Choice of Embedding Model

Start with Off-the Shelf Models

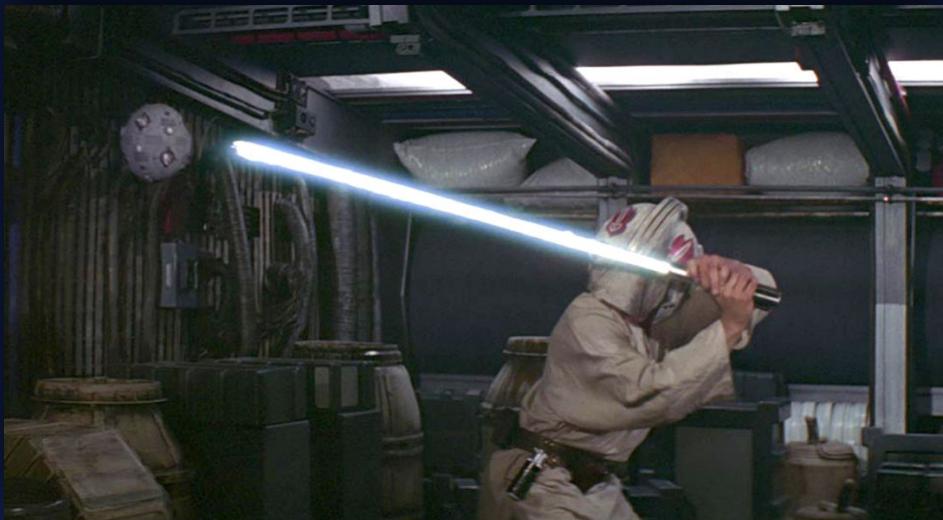
- Text data: Hugging Face (like Microsoft's E5)
- Images: OpenAI's CLIP

Extend to Higher Relevance

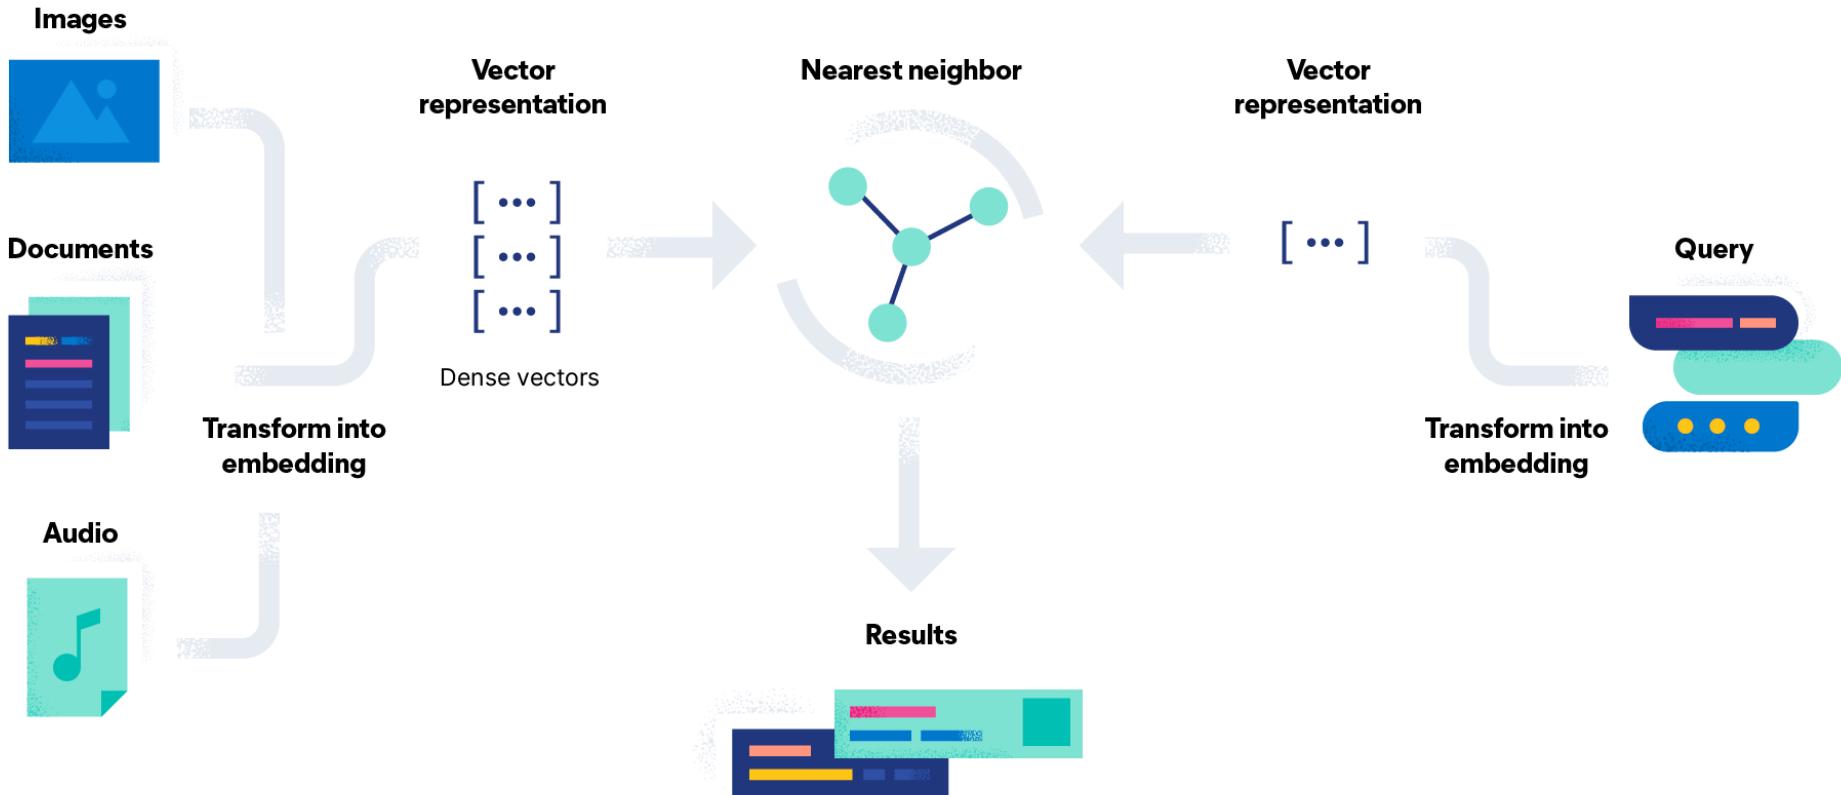
- Apply hybrid scoring
- Bring Your Own Model: requires expertise + labeled data

Problem

training vs actual use-case



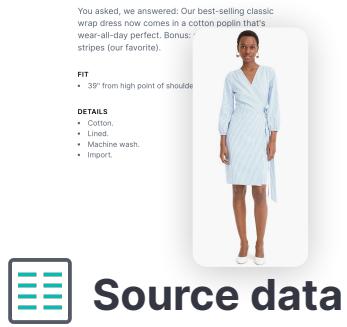
Architecture of Vector Search





How do you index **vectors**?

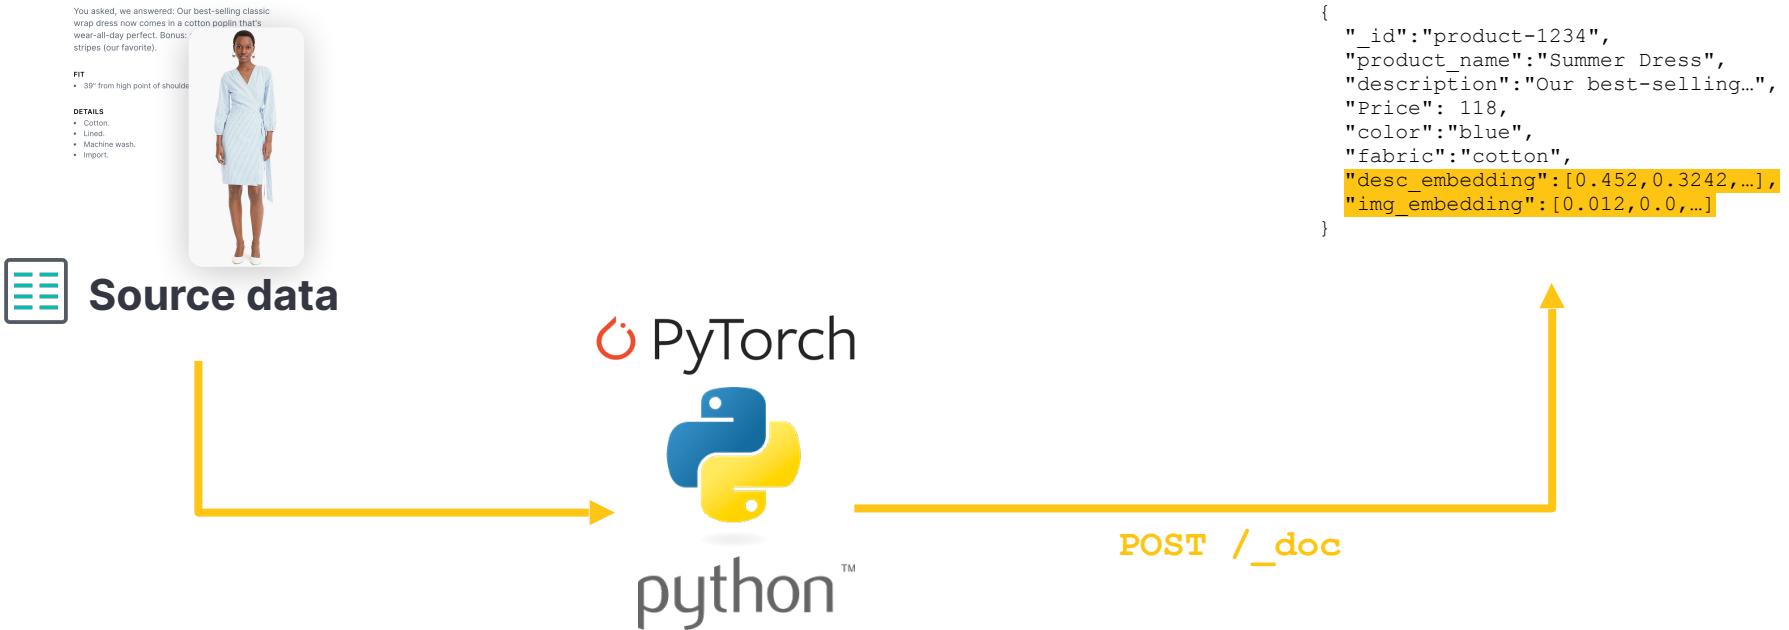
Data Ingestion and Embedding Generation



POST /_doc

```
{  
  "_id": "product-1234",  
  "product_name": "Summer Dress",  
  "description": "Our best-selling...",  
  "Price": 118,  
  "color": "blue",  
  "fabric": "cotton"  
}
```

Data Ingestion and Embedding Generation



With Elastic ML

You asked, we answered. Our best-selling classic wrap dress now comes in a cotton poplin that's wear-all-day perfect. Bonus: stripes (our favorite).

FIT
• 38" from high point of shoulder

DETAILS
• Cotton
• Lined
• Machine wash
• Import



Source data



POST /_doc

ML Inference pipelines + Add inference pipeline

Inference pipelines will be run as processors from the Enterprise Search Ingest Pipeline

ml-inference-embedding-generation	Actions
• Deployed pytorch text_embedding	⋮

ml-inference-emotional-analysis	Actions
• Deployed pytorch text_classification	⋮

Learn more about deploying ML models in Elastic [↗](#)

```
{  
  "_id": "product-1234",  
  "product_name": "Summer Dress",  
  "description": "Our best-selling...",  
  "Price": 118,  
  "color": "blue",  
  "fabric": "cotton",  
  "desc_embedding": [0.452, 0.3242, ...]  
}
```



Elastic's range of supported NLP models

● Fill mask model

Mask some of the words in a sentence and predict words that replace masks

● Named entity recognition model

NLP method that extracts information from text

● Text embedding model

Represent individual words as numerical vectors in a predefined vector space

● Text classification model

Assign a set of predefined categories to open-ended text

● Question answering model

Model that can answer questions given some or no context

● Zero-shot text classification model

Model trained on a set of labeled examples, that is able to classify previously unseen examples

Third party fill-mask models

- BE
- Dis
- MP
- Ro

- BE
- De
- Dis

Third party text classification models

- BE
- De
- Dis
- Dis
- Fin
- Tw

Third party named entity recognition models

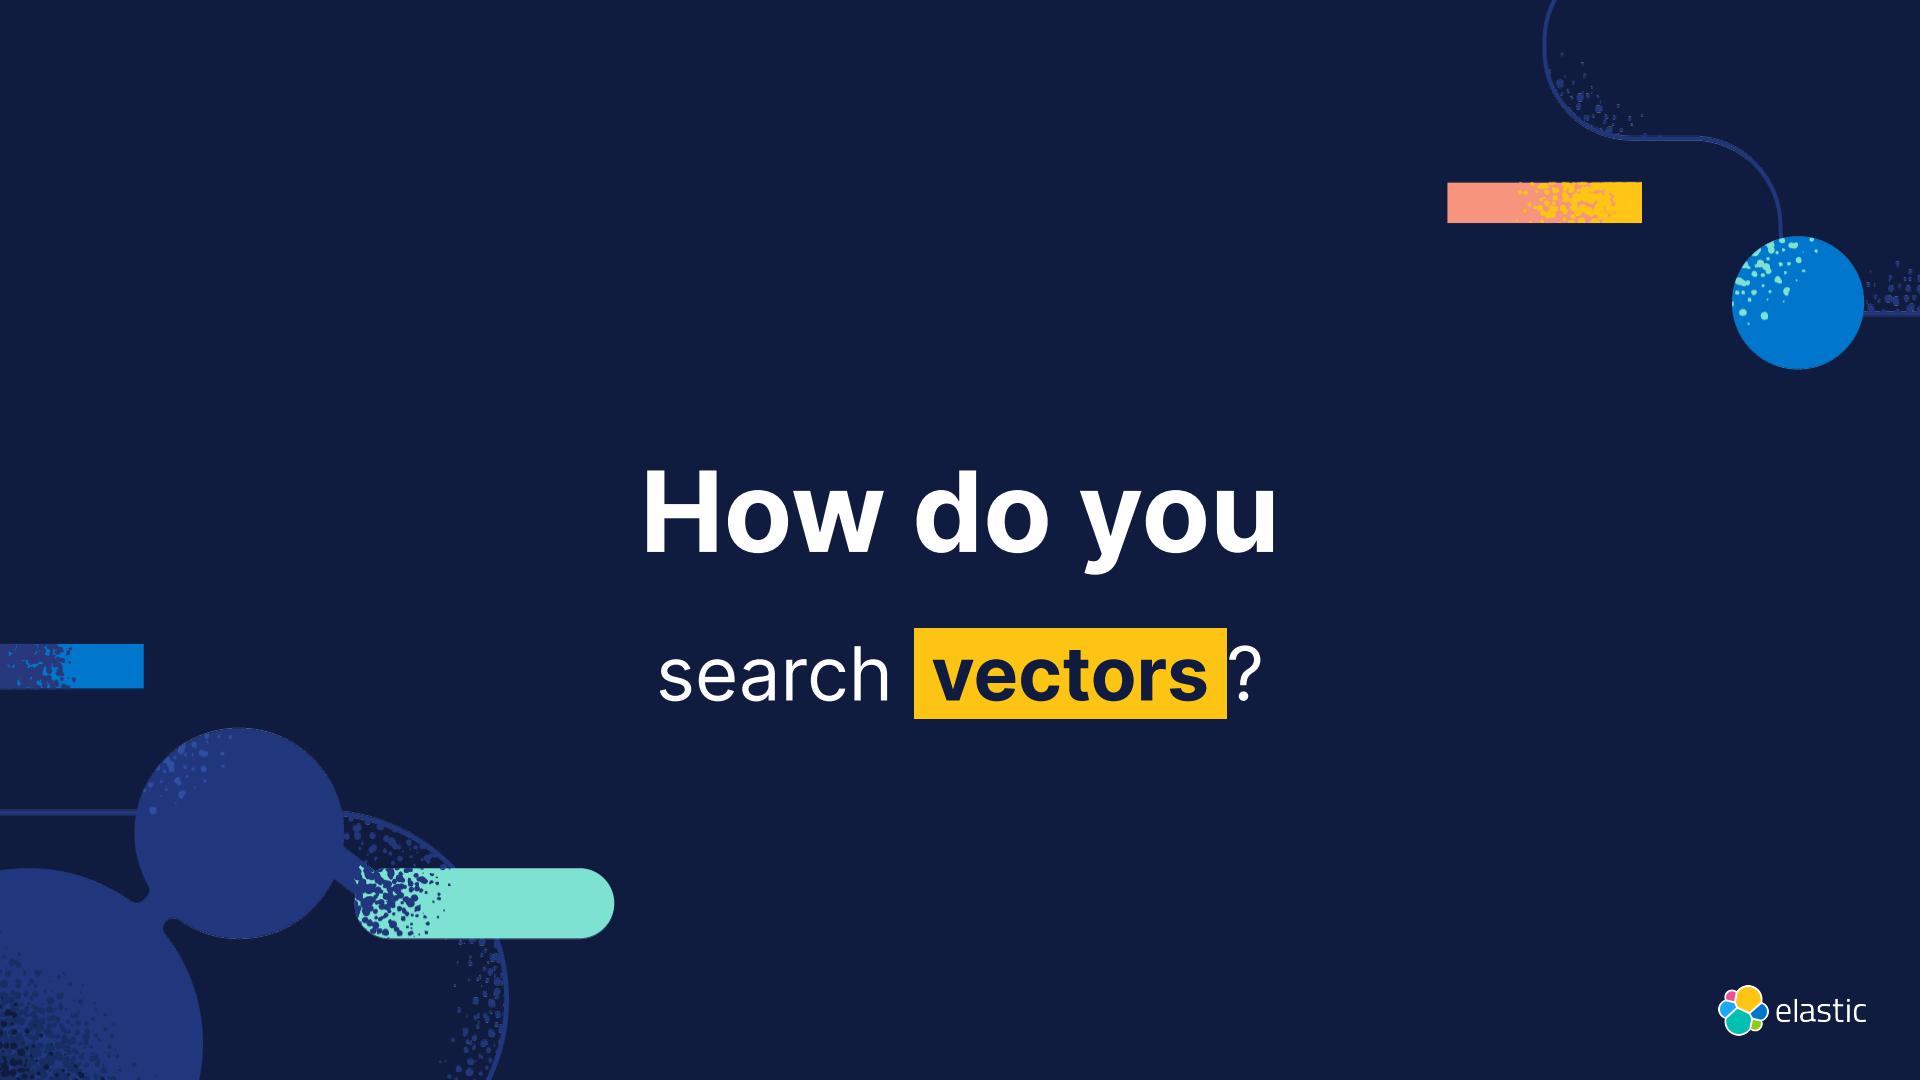
Third party question answering models

Third party text embedding models

Text **Third party zero-shot text classification mode**

hety

- BART large mmlu
- DistilBERT base model (uncased)
- **DistilBart MNLI**
- MobileBERT: a Compact Task-Agnostic BERT for Resource-Limited Devices
- NLI DistilRoBERTa base
- NLI RoBERTa base
- SqueezeBERT



How do you search **vectors**?

Vector Query



```
GET product-catalog/_search
{
  "query" : {
    "bool": {
      "must": [
        {
          "knn": {
            "field": "desc_embedding",
            "num_candidates": 50,
            "query_vector": [0.123, 0.244,...]
          }
        }
      ],
      "filter": {
        "term": {
          "department": "women"
        }
      }
    },
    "size": 10
  }
}
```

Vector Query



Transformer model

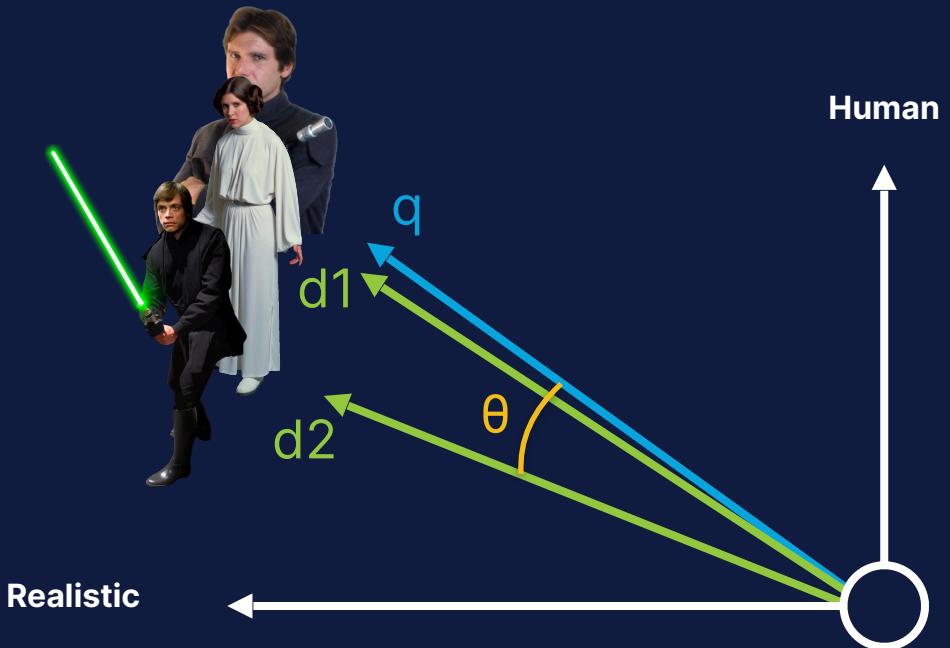


```
GET product-catalog/_search
{
  "query" : {
    "bool": {
      "must": [ {
        "knn": {
          "field": "desc_embedding",
          "num_candidates": 50,
          "query_vector_builder": {
            "text_embedding": {
              "model_text": "summer clothes",
              "model_id": <text-embedding-model>
            }
          }
        }
      } ],
      "filter": {
        "term": {
          "department": "women"
        }
      }
    },
    "size": 10
  }
}
```



But how does it really work?

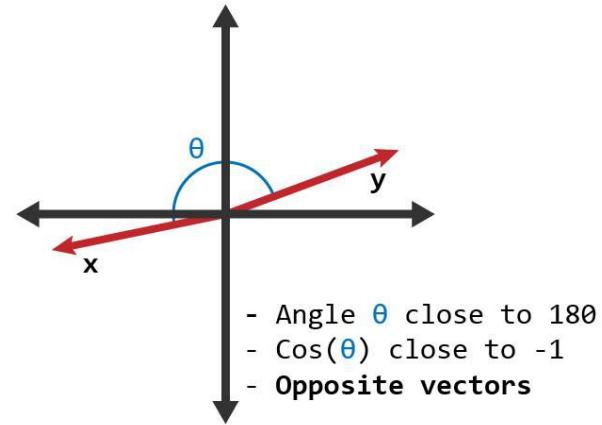
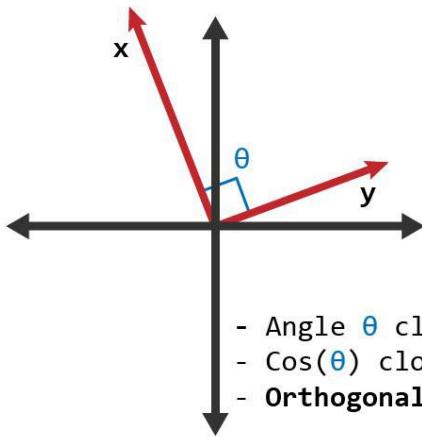
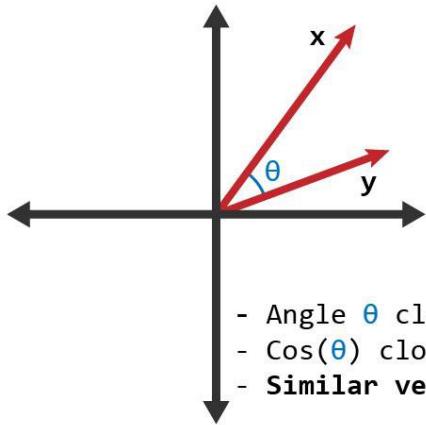
Similarity: cosine (cosine)



$$\cos(\theta) = \frac{\vec{q} \times \vec{d}}{|\vec{q}| \times |\vec{d}|}$$

$$-score = \frac{1 + \cos(\theta)}{2}$$

Similarity: cosine (cosine)

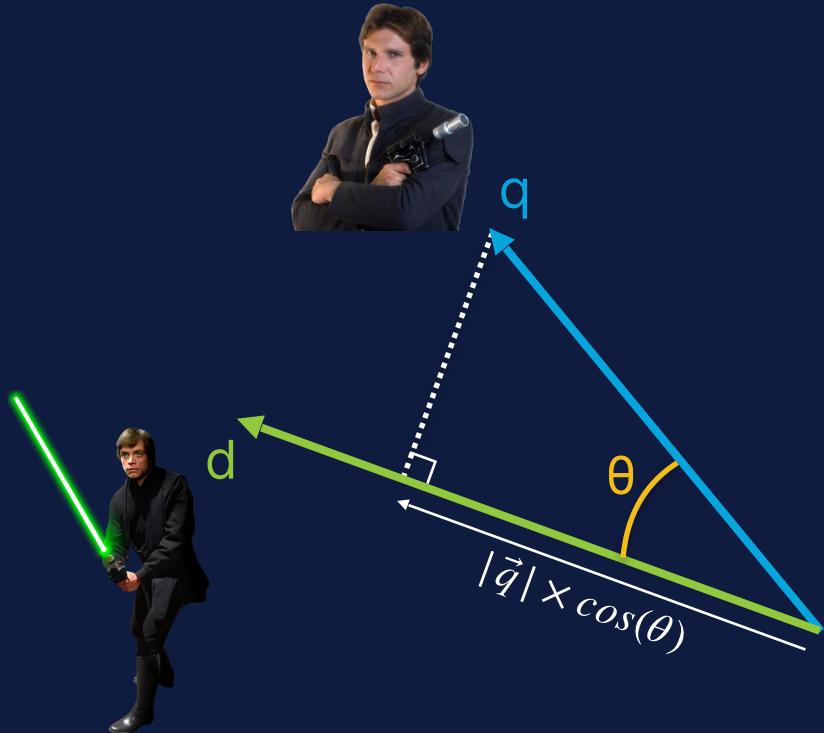


$$\text{_score} = \frac{1 + 1}{2} = 1$$

$$\text{_score} = \frac{1 + 0}{2} = 0.5$$

$$\text{_score} = \frac{1 - 1}{2} = 0$$

Similarity: Dot Product (dot_product)

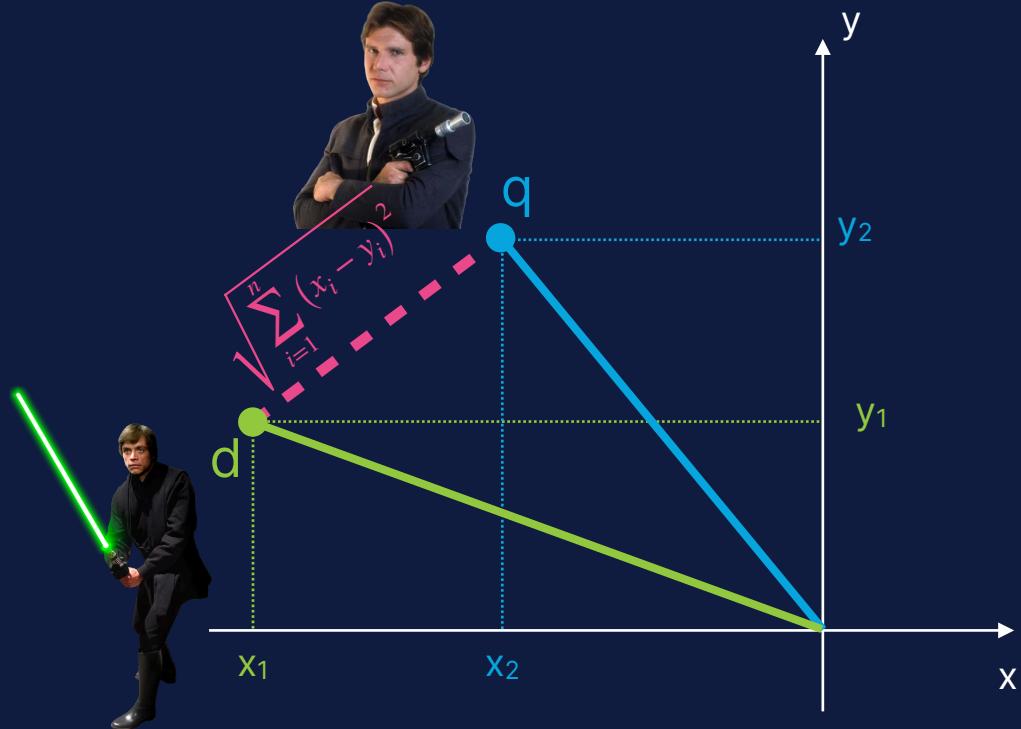


$$\vec{q} \times \vec{d} = |\vec{q}| \times \cos(\theta) \times |\vec{d}|$$

$$-score_{float} = \frac{1 + dot_product(q, d)}{2}$$

$$-score_{byte} = \frac{0.5 + dot_product(q, d)}{32768 \times dims}$$

Similarity: Euclidean distance (l2_norm)



$$-score = \frac{1}{1 + (l2_norm_{q,d})^2}$$

Brute Force



Hierarchical Navigable Small Worlds (HNSW)

One popular approach



HNSW: a layered approach that simplifies access to the nearest neighbor



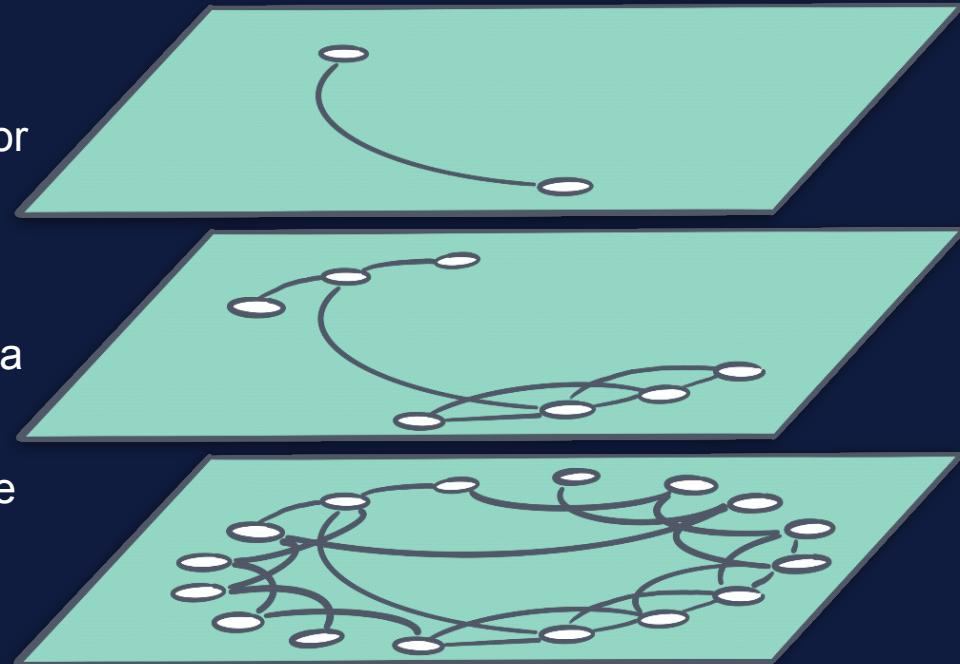
Tiered: from coarse to fine approximation over a few steps



Balance: Bartering a little accuracy for a lot of scalability

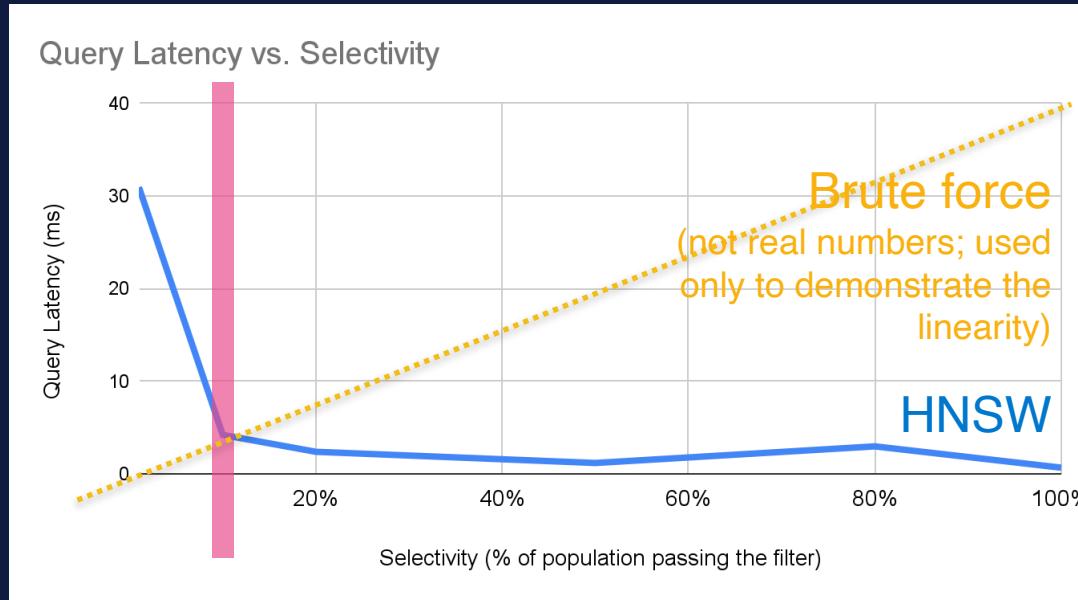


Speed: Excellent query latency on large scale indices



Filtering KNN Vector Similarity

Automatically choose between brute force and HNSW



Bound worst case to 2^* (brute force)

- Brute force scales $O(n)$ of filtered
- HNSW scales $\sim O(\log(n))$ of all docs

Elasticsearch + Lucene = fast progress ❤

Increase max number of vector dims to 2048 #95257

Increase the max vector dims to 4096 #99682

Merged

mayya-sharipova merged 2 commits into `elastic:main` from `mayya-sharipova:increase_vector_dims_4096` ↗ or

Conversation 5

Commits 2

Checks 0

Files changed 8



mayya-sharipova commented on Sep 19

Contributor ...

No description provided.

Commits Increase the max vector dims to 4096

✗ 3f97c5f

Labels mayya-sharipova added `>enhancement` `:Search/Vectors` `v8.11.0` labels on Sep 19

Scaling Vector Search

Vector search

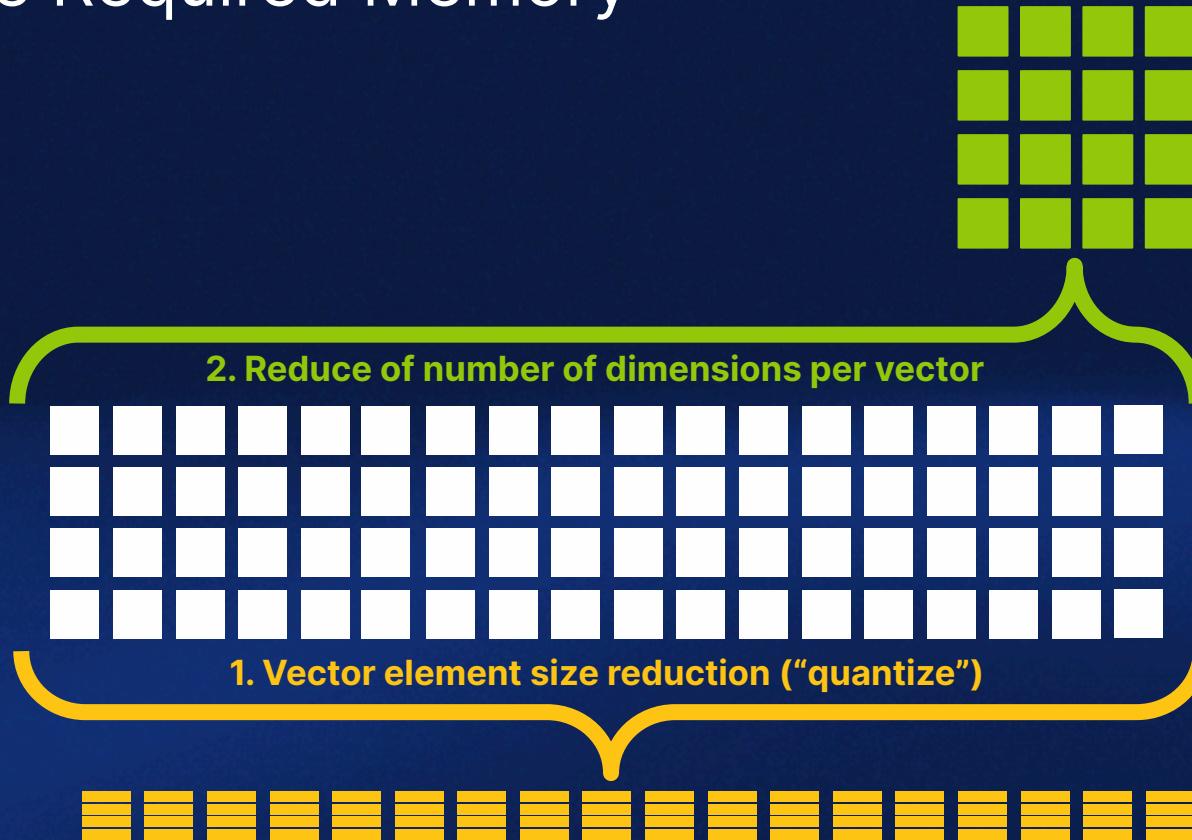
1. Needs lots of memory
2. Indexing is slower
3. Merging is slow

* Continuous improvements in Lucene +
Elasticsearch

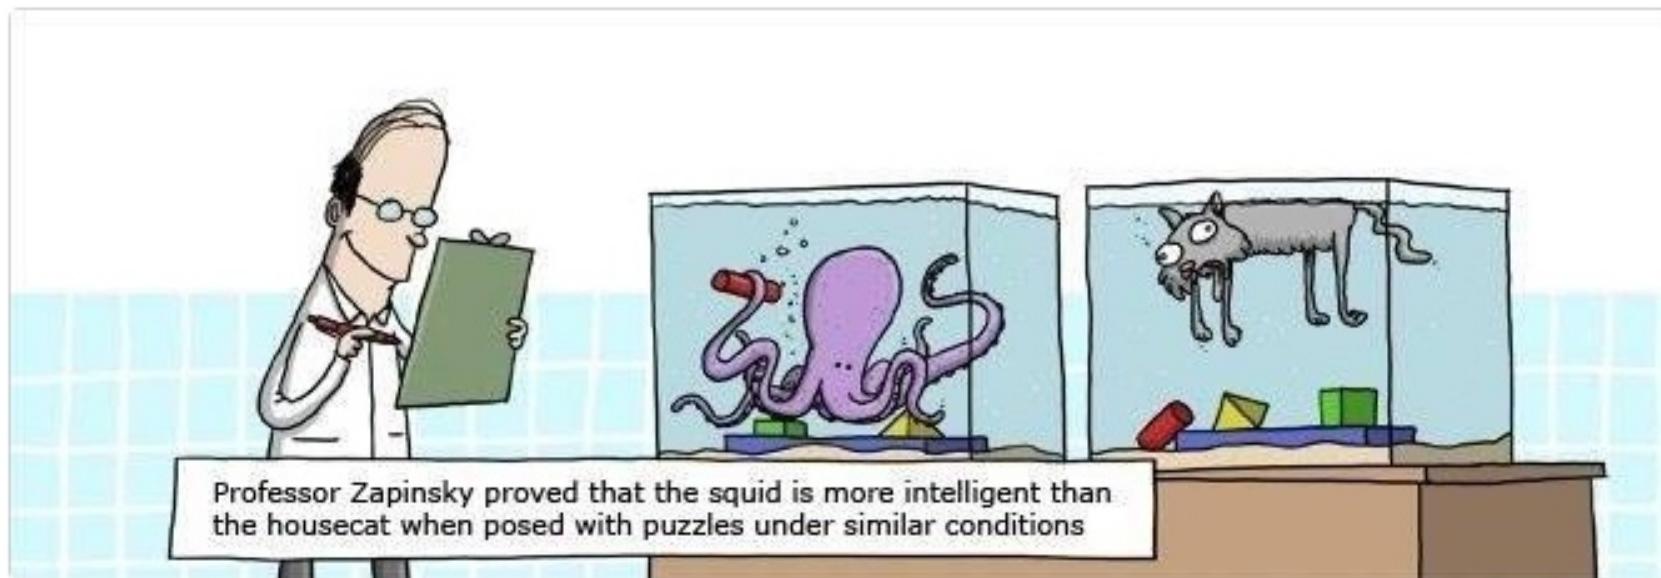
Best practices

1. Avoid searches during indexing
2. Exclude vectors from `_source`
3. Reduce vector dimensionality
4. Use byte rather than float

Reduce Required Memory



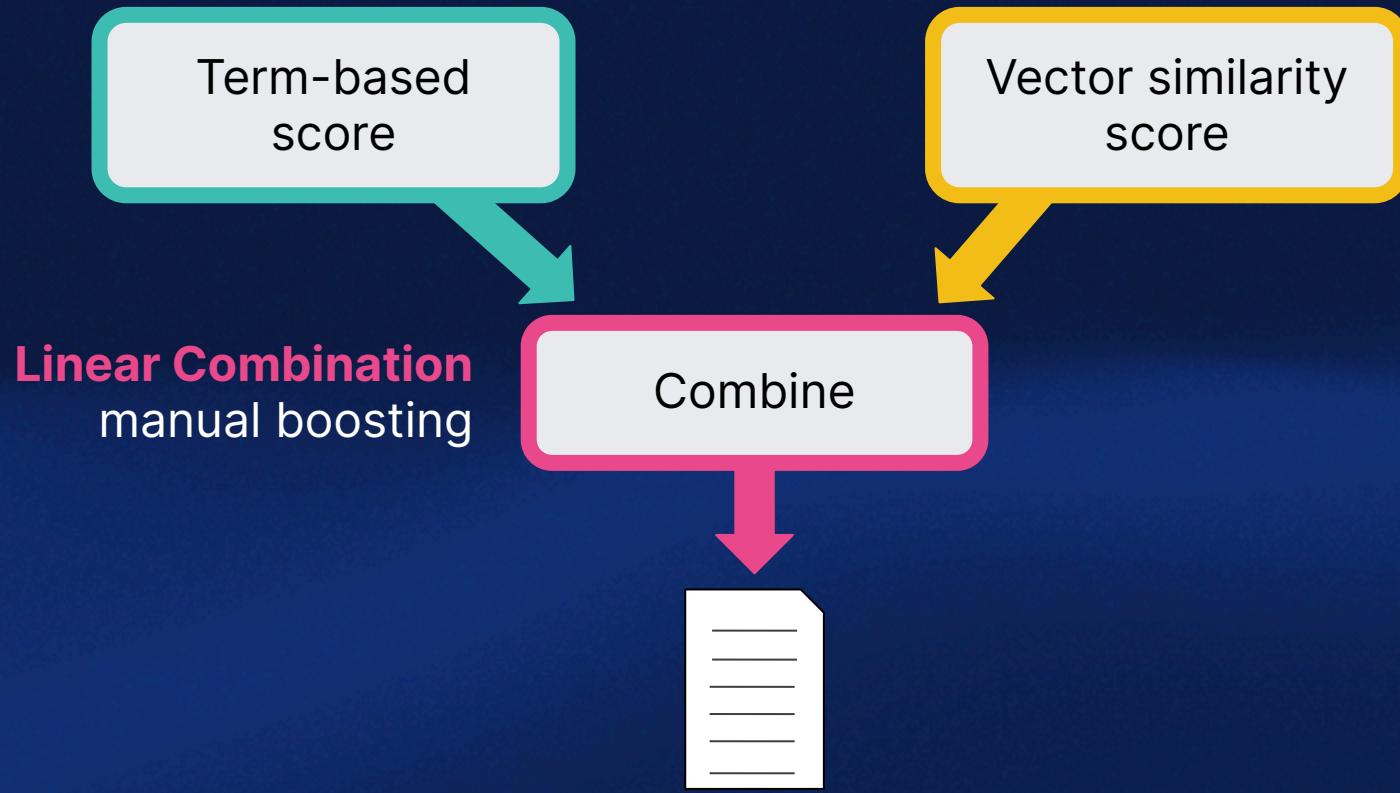
Benchmarketing



Elasticsearch

You Know, for **Hybrid** Search

Hybrid scoring



```
GET product-catalog/_search
{
  "query" : {
    "bool" : {
      "must" : [ {
        "match": {
          "description": {
            "query": "summer clothes",
            "boost": 0.9
          }
        }
      }, {
        "knn": {
          "field": "desc_embedding",
          "query_vector": [0.123, 0.244, ...],
          "num_candidates": 50,
          "boost": 0.1,
          "filter": {
            "term": {
              "department": "women"
            }
          }
        }
      }, {
        "filter" : {
          "range" : { "price": { "lte": 30 } }
        }
      }
    }
  }
}
```

summer clothes

pre-filter

post-filter



```
GET product-catalog/_search
{
  "query" : {
    "bool" : {
      "must" : [ {
        "match": {
          "description": {
            "query": "summer clothes",
            "boost": 0.9
          }
        }
      }, {
        "knn": {
          "field": "image-vector",
          "query_vector": [54, 10, -2],
          "num_candidates": 50,
          "boost": 0.1
        }
      }, {
        "knn": {
          "field": "title-vector",
          "query_vector": [1, 20, -52, 23, 10],
          "num_candidates": 10,
          "boost": 0.5
        }
      }
    }
  }
}
```



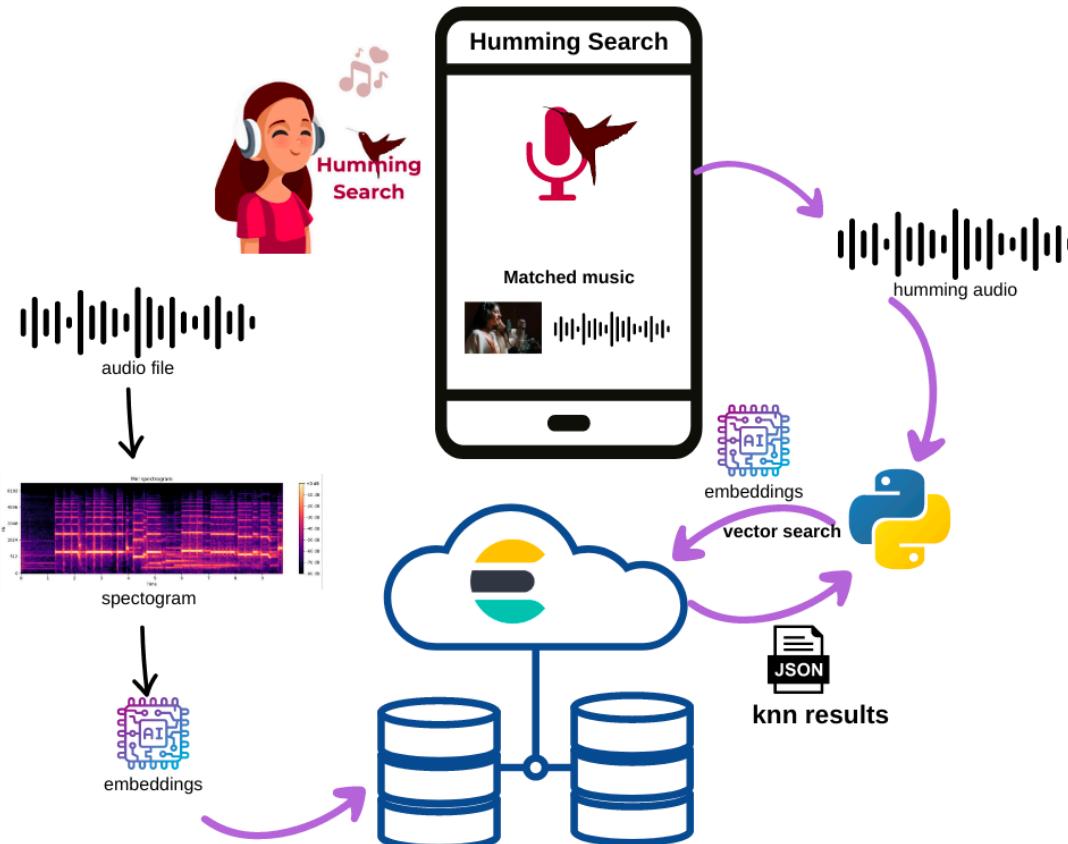


<https://djdadoo.pilato.fr/>



Anniversaire **Lucas** - 25 ans

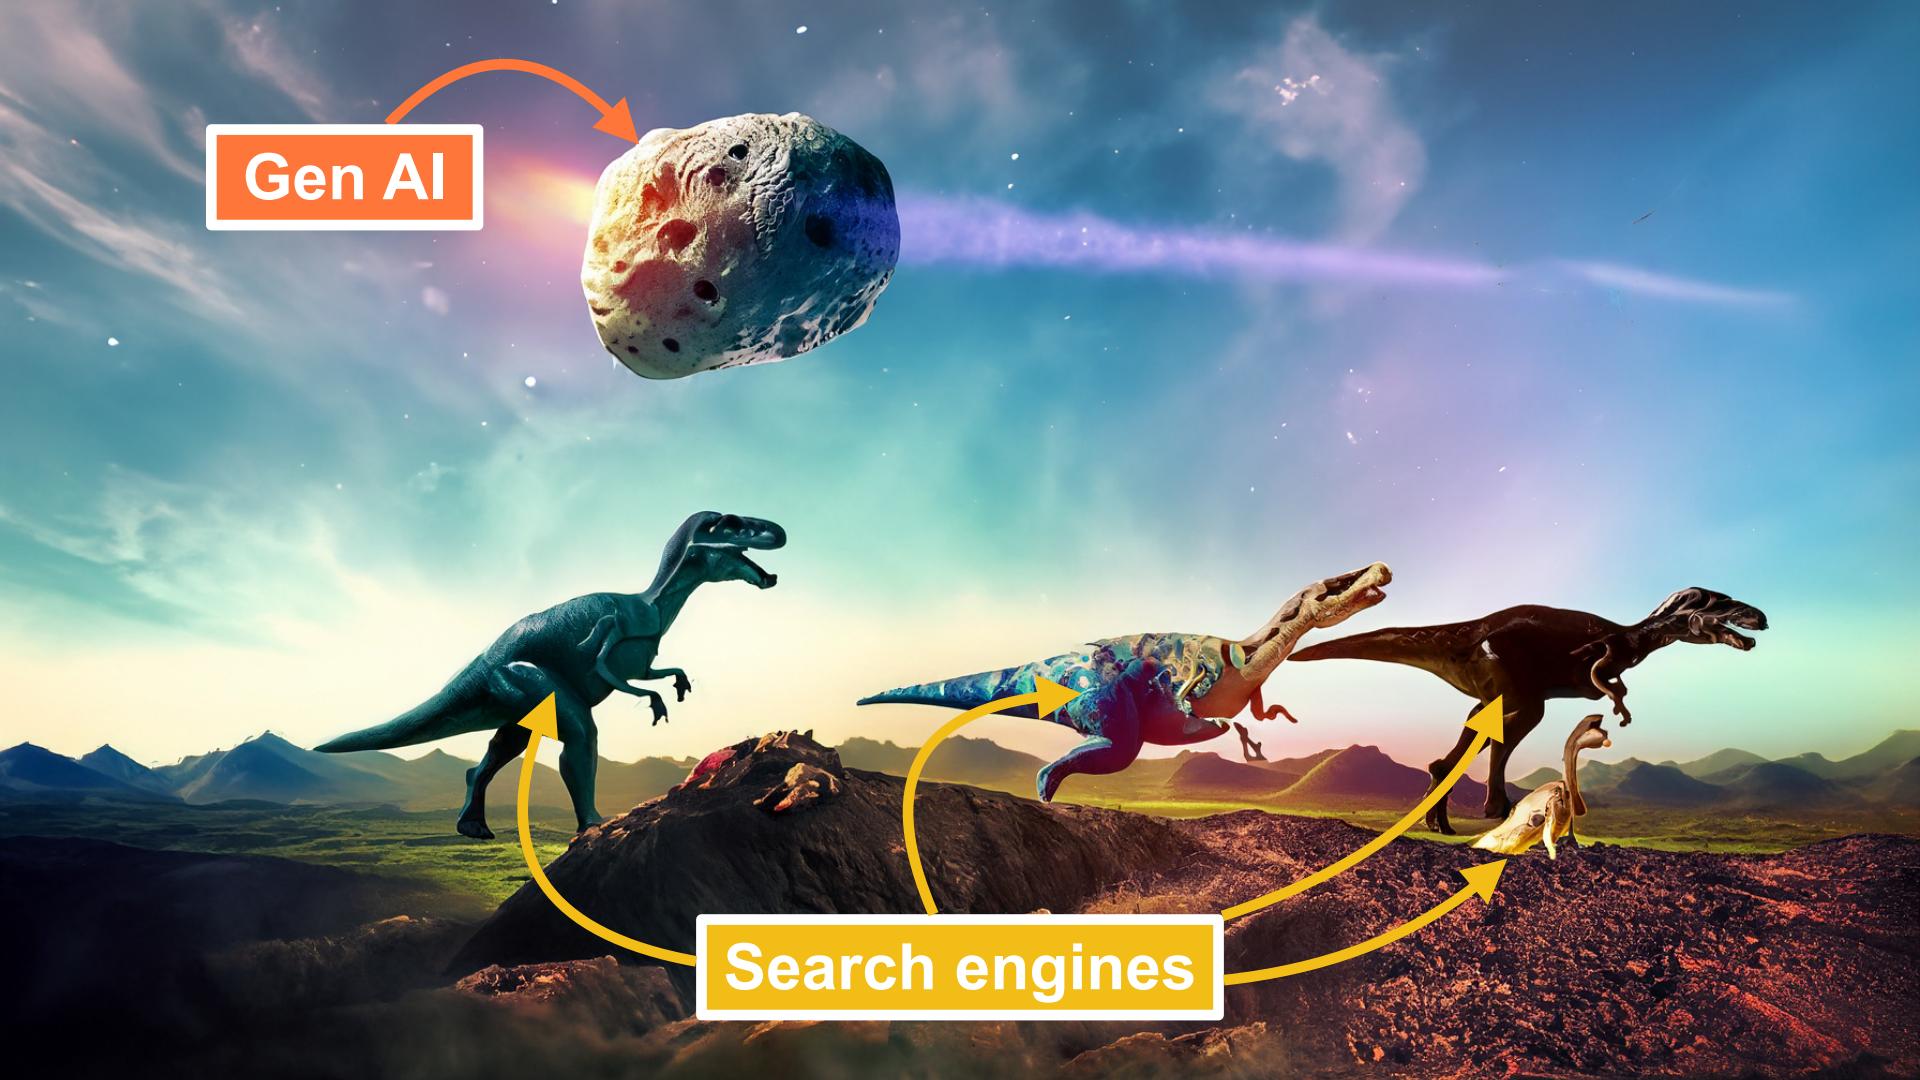




<https://github.com/dadoonet/music-search/>

ChatGPT

Elastic and LLM

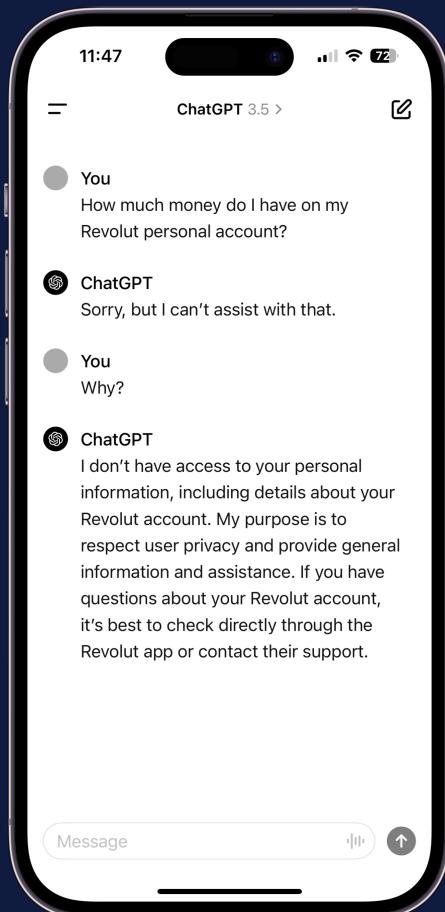
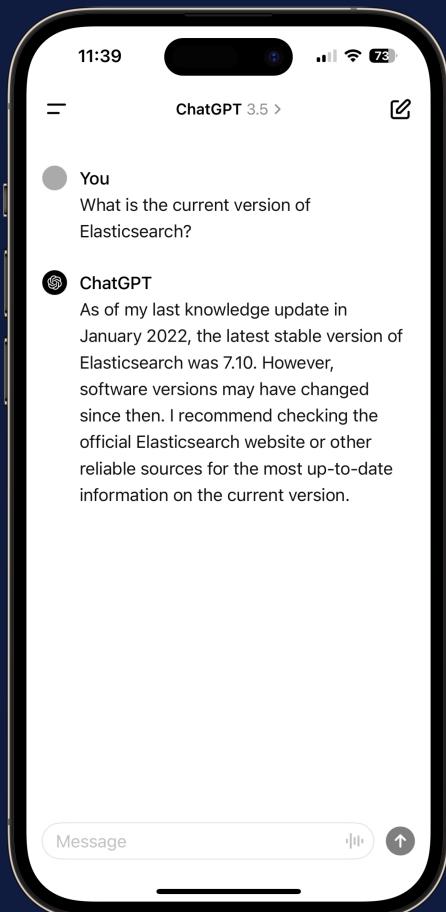


Gen AI

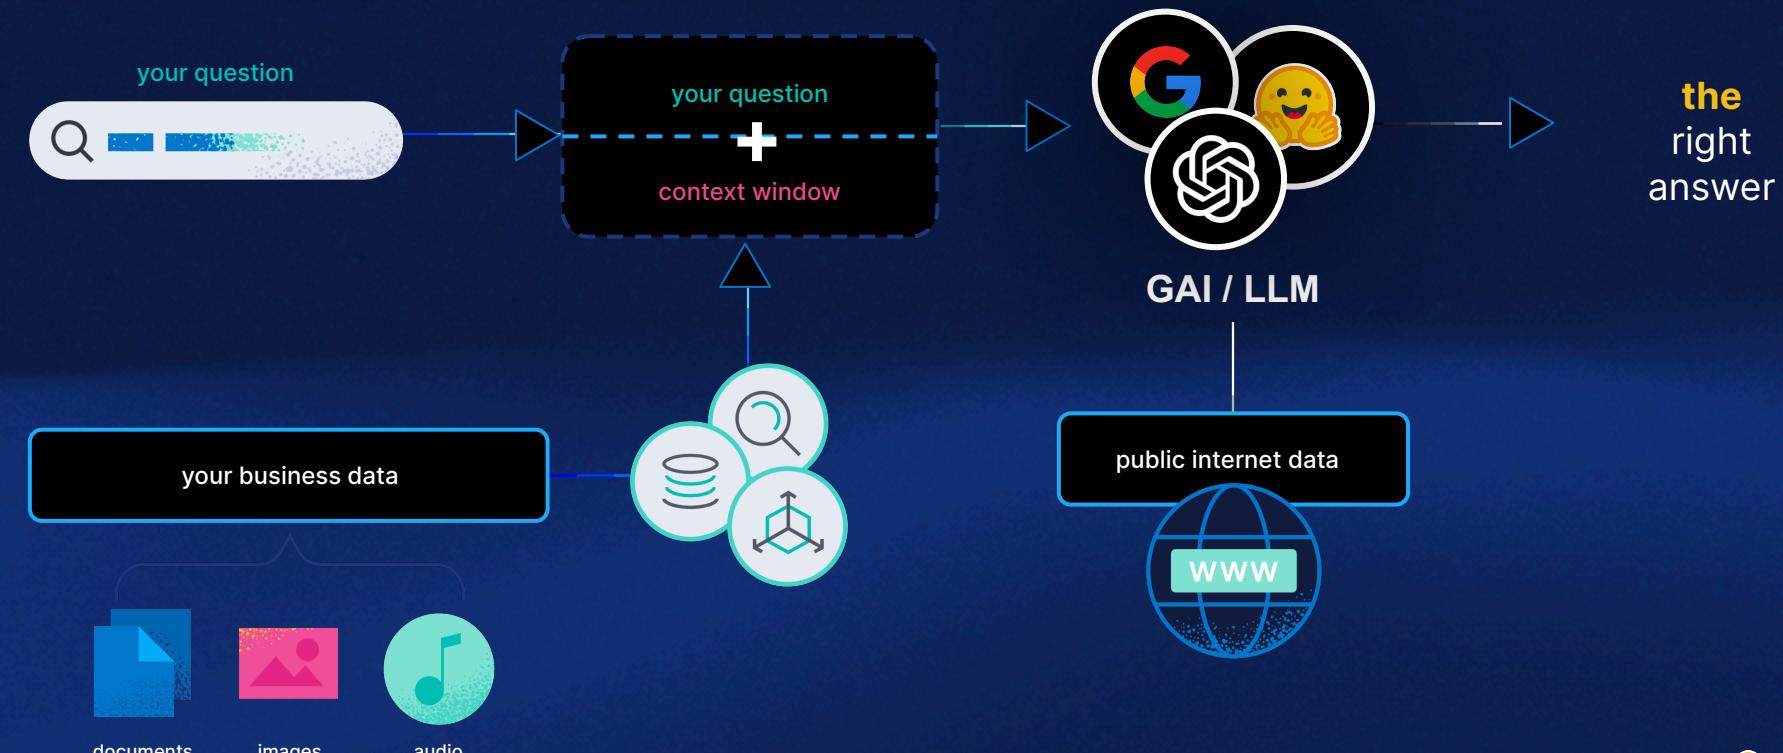
Search engines

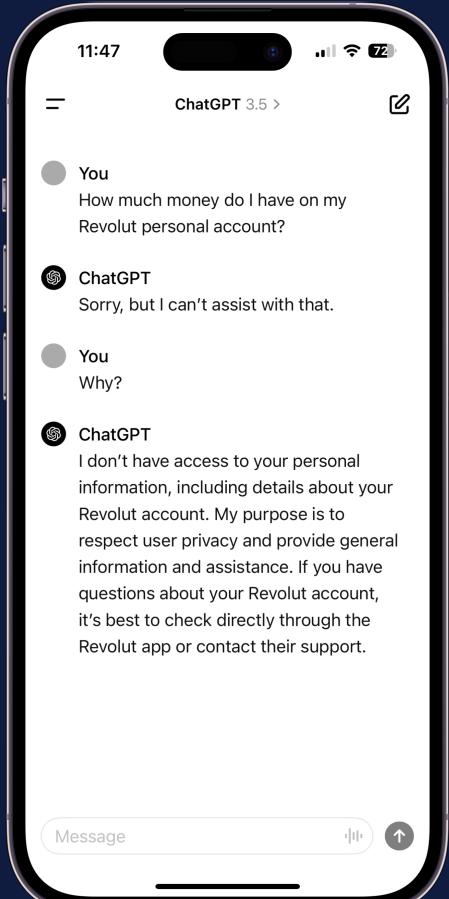
LLM: opportunities and limits





Retrieval Augmented Generation





Home Online banking Enviroment setup

Transaction search Financial summary Customer support

Search your transactions:

This search is not enabled by Elastic and reflects the kind of functionality available to customers today.

Submit

Date	Account	Description	Value	Opening balance	Closing balance
18/06/24	EL03-130981-Transmission	Inbound payment made from EL03-130981-Transmission, St.James's Plac (STJ): 864dce1b-bb95-47d5-87dd-7d02f3b10c3f	7419.0	-825.0	6594.0
18/06/24	EL03-130981-Transmission	Purchase at merchant: Southeastern Grocers, LLC, location: Fayetteville,AR	82.0	6594.0	6512.0
18/06/24	EL03-130981-Transmission	Purchase at merchant: Müller Holding Ltd. & Co. KG, location: Glendale,AZ	188.0	6512.0	6324.0
17/06/24	EL03-130981-Transmission	Payment made from EL03-130981-Transmission to Elwood Erickson, Mitie Grp. (MTO): d37085fc-1382-4593-9cb8-26e5526bd9a0	533.0	20.0	-513.0
17/06/24	EL03-130981-Transmission	Payment made from EL03-130981-Transmission to Classie Johns, Barclays (BARC): 75b603a2-1c1b-45e9-a7ec-4a551bf98a8d	312.0	-513.0	-825.0
16/06/24	EL03-130981-Transmission	Purchase at merchant: E-MART Inc., location: Fayetteville,AR	31.0	51.0	20.0
14/06/24	EL03-130981-Transmission	Purchase at merchant: Dick's Sporting Goods, Inc., location: Montgomery,AL	182.0	329.0	147.0
14/06/24	EL03-130981-Transmission	Purchase at merchant: Valor Holdings Co., Ltd., location: Louisville,KY	96.0	147.0	51.0
13/06/24	EL03-130981-Transmission	Purchase at merchant: The Save Mart Companies, location:	34.0	363.0	329.0



Elasticsearch

You Know, for **Semantic** Search



Search & AI: a new era

David Pilato | [@dadoonet](https://twitter.com/dadoonet)