elasticsearch

From 101 to 102

Philipp Krenn @xeraa

elastic

Developer 🥑

Everyone's use?

Elasticsearch in 1 minute

- Search Engine (FTS, analytics, geo), near real-time
- Distributed, scalable, highly available, resilient
- HTTP & JSON
- Scale - Speed - Relevance
- 💝 of Elastic Stack & Solutions — Search, Observability, Security

Agenda

Architecture

Full-Text Search

Beyond Search

Architecture

Elasticsearch & Lucene

Structure

Cluster - Node - Index - Shard - Document - ID

Replication

Primary & Replica Shard(s)

Requests

Write - GET - Search

Full-Text Search

Database vs Search Engine

elastic

@xeraa

Example

These are not the droids you are looking for.

`html_strip` Char Filter

These are not the droids you are looking for.

`standard` Tokenizer

These are not the droids you are looking for

`lowercase` Token Filter

these are not the droids you are looking for

`stop` Token Filter

droids you looking

`snowball` Token Filter

droid you look

```
GET /_analyze
{
  "analyzer": "english",
  "text": "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."
}
```

Another Example

Obi-Wan never told you what happened to your father.

Another Example

obi wan never told you what happen your father

Another Example

No. I am your father.

Another Example

i am your father

Inverted Index

	ID 1	ID 2	ID 3
am	0	0	1[2]
droid	1[4]	0	0
father	0	1[9]	1[4]
happen	0	1[6]	0
i	0	0	1[1]
look	1[7]	0	0
never	0	1[2]	0
obi	0	1[0]	0
told	0	1[3]	0
wan	0	1[1]	0
what	0	1[5]	0
you	1[5]	1[4]	0
your	0	1[8]	1[3]

Scoring

Term Frequency / Inverse Document Frequency (TF/IDF)

BM25

Default in Elasticsearch 5.0

https://speakerdeck.com/elastic/improved-text-scoring-with-bm25

Term Frequency

$$tf(t\ in\ d) = \sqrt{frequency}$$

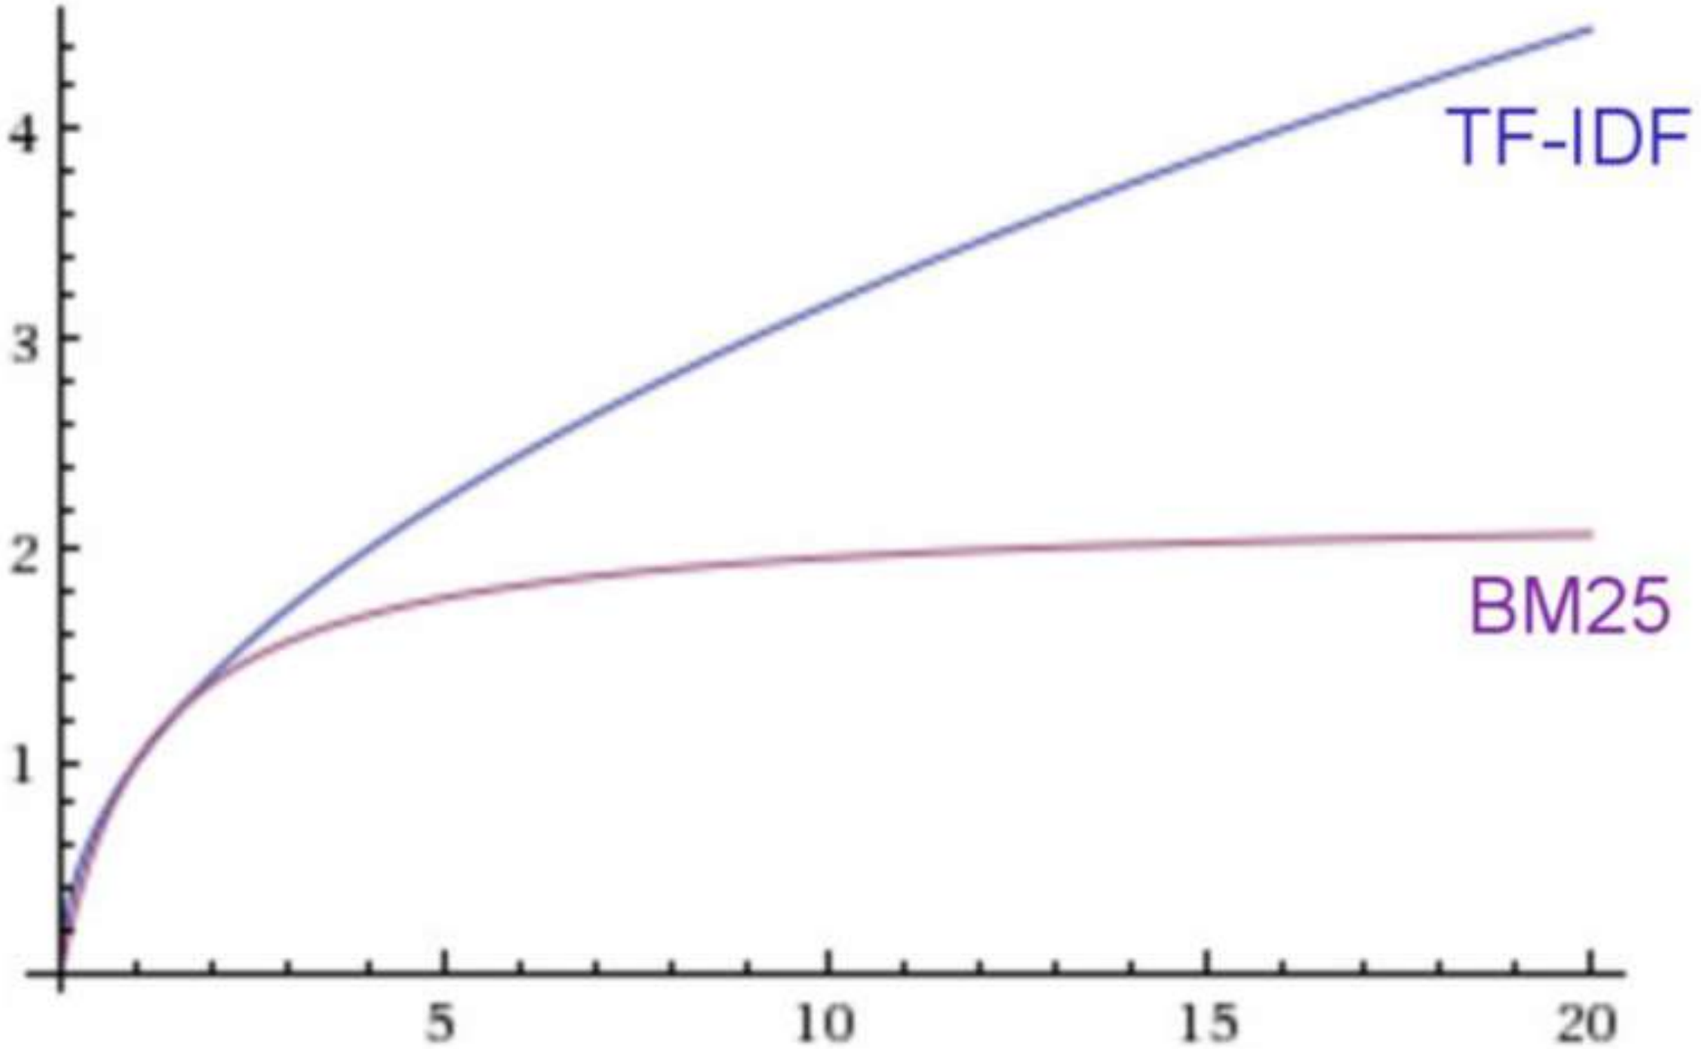

Inverse Document Frequency

$$idf(t) = 1 + log(\frac{numDocs}{docFreq + 1})$$

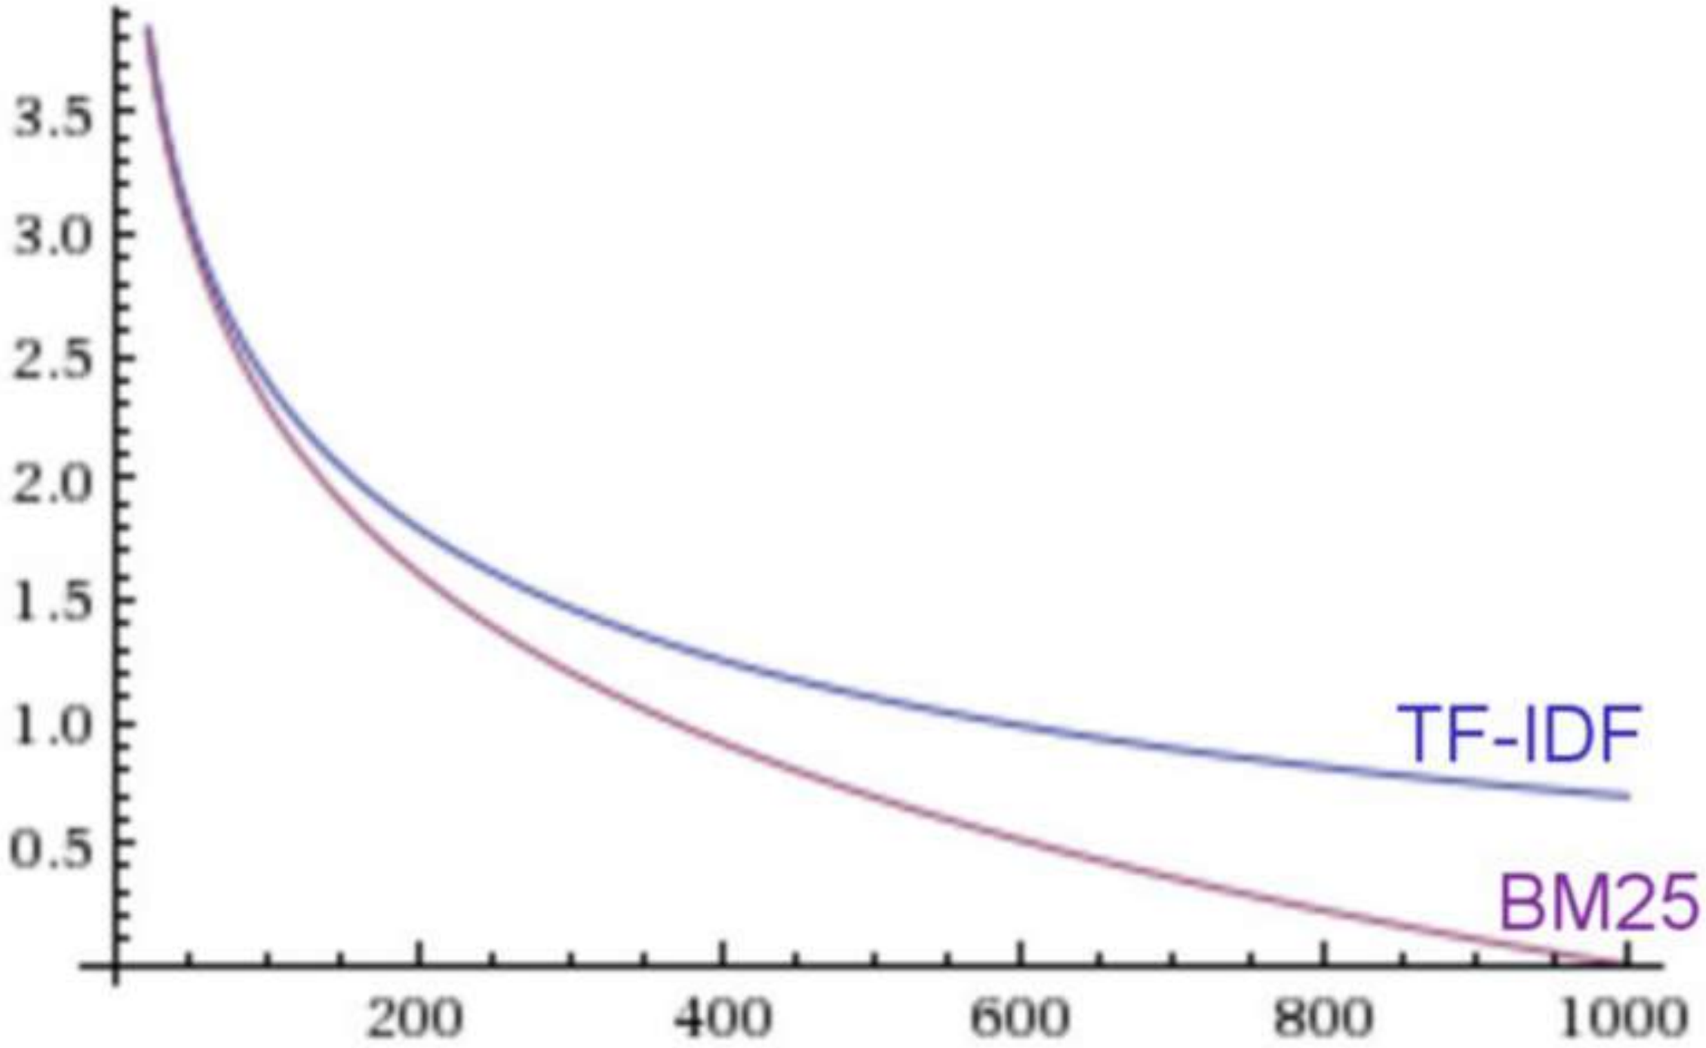

Field-Length Norm

$$norm(d) = \frac{1}{\sqrt{numTerms}}$$

```
POST /starwars/_search?explain=true
{
  "query": {
    "match": {
      "quote": "father"
    }
  }
}
```

Score

0.7209597: i am your father

0.5778280: obi wan never told you what happen your father

Coordination Factor

Reward multiple terms

Search for Three Terms

1 term: $X \cdot \frac{1}{3}$

2 terms: $(X + Y) \cdot \frac{2}{3}$

3 terms: $(X + Y + Z) \cdot \frac{3}{3}$

Practical Scoring Function

```
score(q,d)  =
              queryNorm(q)
            · coord(q,d)
            · ∑ (
                  tf(t in d)
                · idf(t)²
                · t.getBoost()
                · norm(t,d)
              ) (t in q)
```

Beyond Search

Kibana Sample Data

Dashboard to Aggregation

Geo Search & Aggregation

Performance

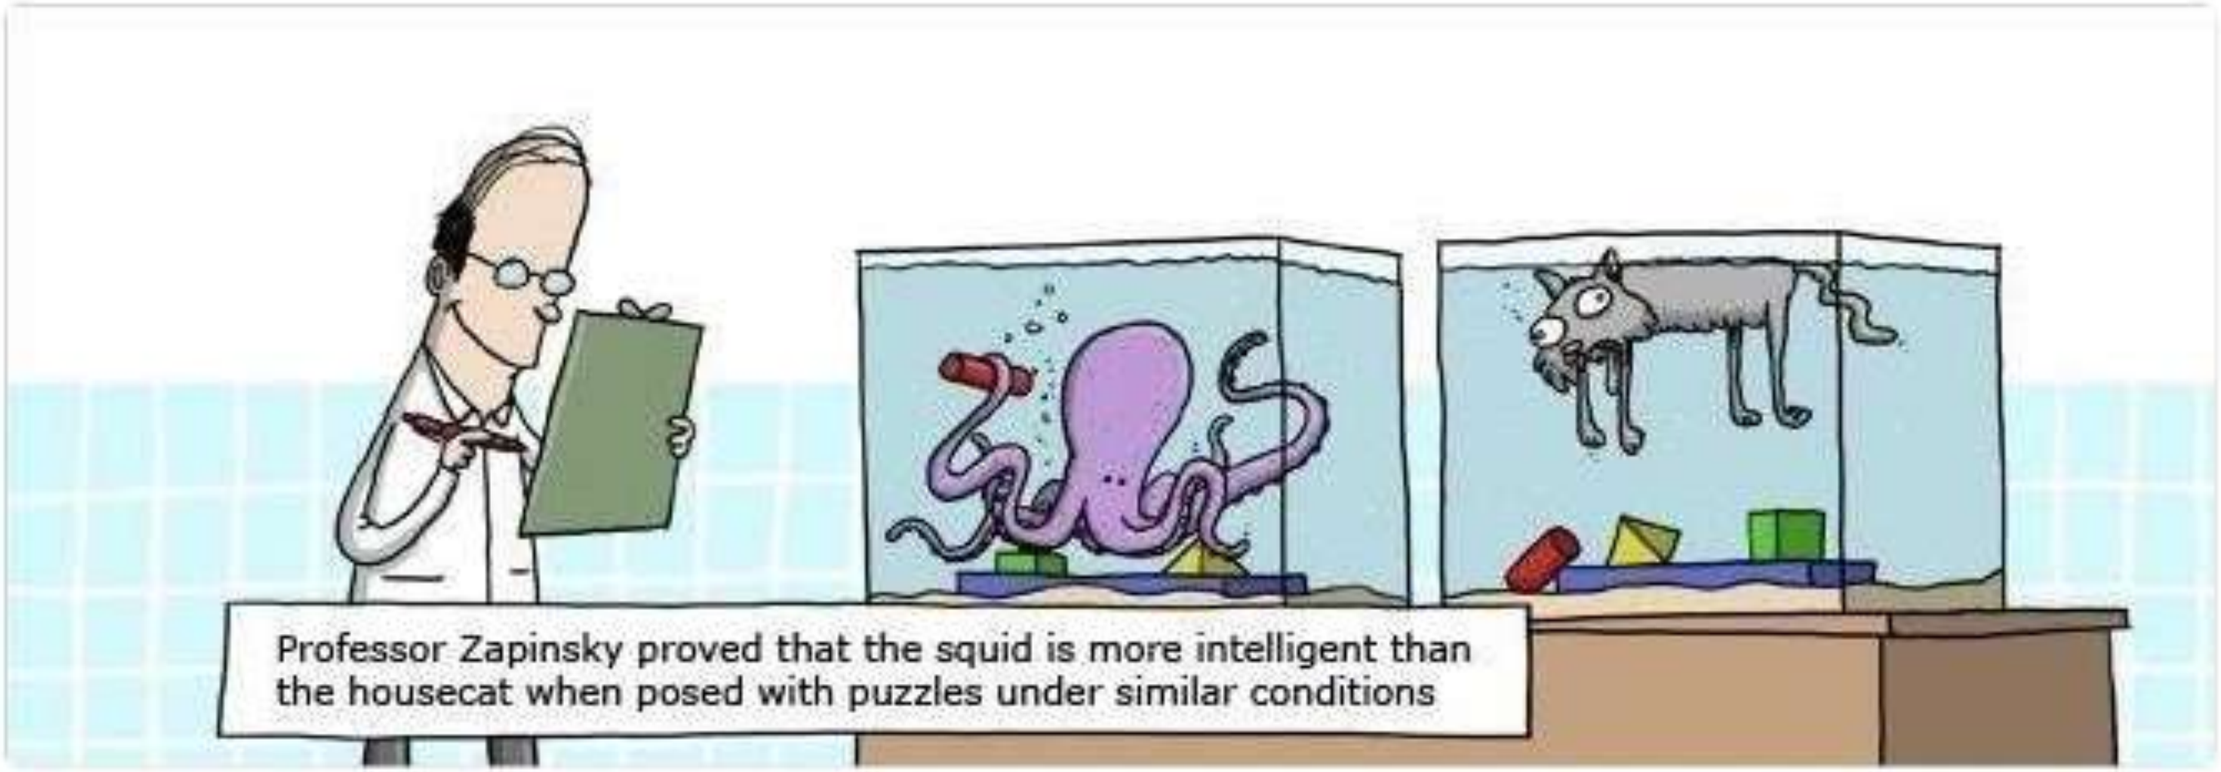
Professor Zapinsky proved that the squid is more intelligent than the housecat when posed with puzzles under similar conditions


ONE MILLION
DOCUMENTS
memegenerator.net

Where to go next?

https://www.elastic.co/guide/

https://www.elastic.co/blog/

https://discuss.elastic.co & https://ela.st/slack

Where to go next?

https://www.elastic.co/elasticon/global
October 13–15

https://www.elastic.co/training/free

Conclusion

Recap

Architecture

Full-Text Search

Beyond Search

Questions?

Philipp Krenn
@xeraa