

# The Curse of Dense Low-Dimensional Information Retrieval for Large Index Sizes

Nils Reimers and Iryna Gurevych

Ubiquitous Knowledge Processing Lab (UKP-TUDA)

Department of Computer Science, Technical University of Darmstadt

[www.ukp.tu-darmstadt.de](http://www.ukp.tu-darmstadt.de)

## Abstract

Information Retrieval using dense low-dimensional representations recently became popular and showed out-performance to traditional sparse-representations like BM25. However, no previous work investigated how dense representations perform with large index sizes. We show theoretically and empirically that the performance for dense representations decreases quicker than sparse representations for increasing index sizes. In extreme cases, this can even lead to a tipping point where at a certain index size sparse representations outperform dense representations. We show that this behavior is tightly connected to the number of dimensions of the representations: The lower the dimension, the higher the chance for false positives, i.e. returning irrelevant documents.

## 1 Introduction

Information retrieval traditionally used sparse representations like TF-IDF or BM25 to retrieve relevant documents for a given query. However, these approaches suffer from the lexical gap problem (Berger et al., 2000).

To overcome this issue, dense representations have been proposed (Gillick et al., 2018): Queries and documents are mapped to a dense vector space and relevant documents are retrieved e.g. by using cosine-similarity. Out-performance over sparse lexical approaches has been shown for various datasets (Gillick et al., 2018; Guo et al., 2020; Guu et al., 2020; Gao et al., 2020).

Previous work showed the out-performance for fixed, rather small indexes. The largest dataset where it has been shown is the MS Marco (Bajaj et al., 2018) passage retrieval dataset, where retrieval is done over an index of 8.8 million text passages. However, in production scenarios, index sizes quickly reach 100 millions of documents.

We show in this paper, that the performance for dense representations can decrease quicker for increasing index sizes than for sparse representations. For a small index of e.g. 100k documents, a dense approach might clearly outperform sparse approaches. However, with a larger index of several million documents, the sparse approach can outperform the dense approach.

We show theoretically and empirically that this effect is closely linked to the number of dimensions for the representations: Using fewer dimensions increases the chances for false positives. This effect becomes more severe with increasing index sizes.

## 2 Related Work

A common choice for dense retrieval is to fine-tune a transformer network like BERT (Devlin et al., 2018) on a given training corpus with queries and relevant documents (Guo et al., 2020; Guu et al., 2020; Gao et al., 2020; Karpukhin et al., 2020; Luan et al., 2020). Recent work showed that combining dense approaches with sparse, lexical approaches can further boost the performance (Luan et al., 2020; Gao et al., 2020). While the approaches have been tested on various information and question answering retrieval datasets, the performance was only evaluated on fixed, rather small indexes. Guo et al. (2020) evaluated approaches for eight different datasets having index sizes between 3k and 454k documents.

We are not aware of previous work that compares sparse and dense approaches for increasing index sizes and the connection to the dimensionality. The only work we are aware of that systematically studies the encoding size for dense approaches is (Luan et al., 2020), but they only studied the connection to the document length.

### 3 Theory

Dense retrieval approaches map queries and documents<sup>1</sup> to a fixed size dense vector. The most relevant documents for a given query can then be found using cosine-similarity.

Using as few dimensions as possible is desirable, as it decreases the memory requirement to store (an index) of millions of vectors and leads to faster retrieval. However, as we show, lower-dimensional representations can have issues with large indices.

Given a query vector  $q \in \mathbb{R}^k$ , we search our index of document vectors  $d_1, \dots, d_n \in \mathbb{R}^k$  for the documents that maximizes:

$$\text{cossim}(q, d_i) = \cos(\theta) = \frac{q \cdot d_i}{\|q\| \|d_i\|}$$

Note: In the following we just show the case for cosine similarity. The proof extends to other similarity functions like dot-product and any p-norm (Manhattan, Euclidean) as long as the vector space is finite. A finite  $n$ -dimensional vector space can be mapped to an  $n + 1$ -dimensional vectors space with vectors of unit length. In that case, dot-product in  $n$  dimensions is equivalent to cosine-similarity in  $n + 1$  dimensions. Similar, any  $p$ -norm in  $n$  dimensions can be re-written as cosine-similarity in  $n + 1$  dimensions.

**Theorem:** The probability for false positives (I) increases with the index size  $n$  and (II) with the decreasing dimensionality  $k$ .

**Proof (I):** Given a query  $q$  and the relevant document  $d_r$ . For simplicity, we assume only a single relevant document. If multiple documents are relevant, we consider only the one with the highest cosine similarity. In order that no false positive is returned,  $\text{cossim}(q, d_r)$  must be greater than  $\text{cossim}(q, d_i)$  for all  $i \neq r$ . Assume the possible vectors are independent. Then, the probability for a false positive is

$$P(\text{false positive}) = 1 - (1 - P(\text{false positive}_i))^{n-1}$$

for an index with  $n - 1$  negative elements and  $P(\text{false positive}_i)$  the probability that a single element is a false positive, i.e.  $\text{cossim}(q, d_i) > \text{cossim}(q, d_r)$ .

**Proof (II):** While the previous proof is straightforward, that the chance of false positives increases with larger index sizes, the more interesting aspect is the relation to the dimensionality, i.e., what is the probability  $P(\text{false positive}_i)$

$= P(\text{cossim}(q, d_i) > \text{cossim}(q, d_r))$  for a random  $d_i$ ? We show that this probability decreases with more dimensions.

Without loss of generality, we assume that the vectors are of unit length. The vectors are then on an  $k$ -dimensional sphere with radius 1. A false positive happens if  $\text{cossim}(q, d_i) > \text{cossim}(q, d_r)$ , or, equivalent if  $1 - \text{cossim}(q, d_i) < 1 - \text{cossim}(q, d_r)$ . I.e., we intersect the sphere in  $k$  dimensions with a hyperplane in  $k - 1$  dimensions. The area of the cut-off portion is defined by  $1 - \text{cossim}(q, d_r)$ . All vectors within the cut-off portion (i.e. spherical cap) are false positives. The probability that a random vector will be returned as false positive is:

$$P(\text{false positive}_i) = A_{cap}/A_{sphere}$$

with  $A_{cap}$  the surface area of the spherical cap and  $A_{sphere}$  the surface area of the sphere in  $k$  dimensions. Define the surface area of the sphere in  $k$  dimensions as  $A_k$ , then the surface area of  $A_{cap}$  is (Li, 2011):

$$A_{cap} = \frac{1}{2} A_k I_{\sin^2 \theta} \left( \frac{k-1}{2}, \frac{1}{2} \right)$$

with  $I_x(a, b)$  the regularized incomplete beta function and  $\theta$  the polar angle, i.e. the angle between  $q$  and the relevant document  $d_r$ . Hence:

$$P(\text{false positive}_i) = \frac{1}{2} I_{\sin^2 \theta} \left( \frac{k-1}{2}, \frac{1}{2} \right) \quad (1)$$

For constant cosine similarity between query  $q$  and relevant document  $d_r$ ,  $I_{\sin^2 \theta} \left( \frac{k-1}{2}, \frac{1}{2} \right)$  is a monotonically decreasing function with increasing dimension  $k$ . In conclusion, more dimensions decrease the probability for false positives.

Combining (I) and (II) shows that a low dimensional representation might work well for small index sizes. However, with more indexed documents, the probability of false positives increases faster for low dimensional representations than for higher dimensional representations. Hence, at some index size, higher dimensional representations might outperform the lower-dimensional representation.

### 4 Empirical Investigation

In the proof, we have assumed that vectors are independent and uniformly distributed over the space, which gives us a lower bound on the false positive rate. However, in practice, dense representations are neither independent nor uniformly

<sup>1</sup>We use *document* as a cover-term for text of any length.

distributed. As shown in (Ethayarajh, 2019; Li et al., 2020), dense representations derived from pre-trained Transformers like BERT map to an anisotropic space, i.e., the vectors occupy only a narrow cone in the vector space. This drastically increases the chance that an irrelevant document is closer to the query embedding than the relevant document. Hence, we study how actual dense models are impacted by increasing index sizes and lower-dimensional representations.

#### 4.1 Dataset

We conduct our experiments on the MS MARCO passage dataset (Bajaj et al., 2018). It consists of over 1 million unique real queries from the Bing search engine, together with 8.8 million paragraphs from heterogeneous web sources. Most of the queries have only 1 passage judged as relevant, even though more can exist. The development set consists of 6980 queries and the performance is evaluated using mean reciprocal rank MRR@10.

To better compare the relative performance differences, we compute a rank-aware error rate:

$$\text{Err} = \frac{1}{n} \sum_{i=1}^n \left( 1 - \frac{1}{\text{rank}_i} \right)$$

with  $\text{rank}_i$  being the rank of the relevant document for the  $i$ -th query. To be compatible with MRR@10, we set  $\text{rank}_i = \infty$  for  $\text{rank}_i > 10$ . We then define the relative error rate as  $\text{Err}_{\text{Dense}}/\text{Err}_{\text{BM25}}$ . A relative error rate of 50% indicates that the dense approach makes only 50% of the errors compared to BM25 retrieval.

#### 4.2 Model

For sparse, lexical retrieval, we use ElasticSearch, which is based on BM25. For dense retrieval, we use a DistilRoBERTa-base model (Sanh et al., 2020) as a bi-encoder: The query and the passage are passed independently to the transformer model and the output is averaged to create fixed-sized representations. We train this using InfoNCE loss (van den Oord et al., 2018):

$$L = -\log \frac{\exp(\tau \cdot \text{cossim}(q, p_+))}{\sum_i \exp(\tau \cdot \text{cossim}(q, p_i))}$$

with  $q$  the query,  $p_+$  the relevant passage. We use in-batch negative sampling and use the other passages in a batch as negative examples. We found that  $\tau = 20$  performs well. We train the model in

two setups: 1) only with random (in-batch) negatives, and 2) we provide for each query additionally one hard-negative passage. We use the hard-negative passages provided by the MS MARCO dataset, which were retrieved using lexical search. Models are trained with a batch size of 128 with Adam optimizer and a learning rate of  $2e - 5$ .

DistilRoBERTa produces representations with 768 dimensions. We also experiment with lower-dimensional representations. There, we added a linear projection layer on-top of the mean pooling operation to down-project the representation to either 128 or 256 dimensions. Dense retrieval is performed using cosine similarity with exact search.

Models were trained using the SBERT framework (Reimers and Gurevych, 2019).<sup>2</sup>

## 5 Experiments

First, we study the impact of increasing index sizes with real text passages. Then, we study the performance when random noise is added.

### 5.1 Increasing Index Size

In the first experiment, we start with an index that only contains the 7433 relevant passages for the 6980 queries. Then, we add step-wise randomly selected passages from the MS MARCO corpus to the index until all 8.8 million passages are indexed.

Model	10k	100k	1M	8.8M
BM25	79.93	63.88	40.14	17.56
Trained without hard negatives				
128 dim	87.50	68.63	39.76	15.71
256 dim	88.82	70.79	41.74	17.08
768 dim	88.99	71.06	42.24	17.34
Trained with hard negatives				
128 dim	90.32	77.92	54.45	27.34
256 dim	91.10	78.90	55.51	28.16
768 dim	91.48	79.42	56.05	28.55

Table 1: Dev performance (MRR@10  $\times$  100) on MS MARCO passage dataset with different index sizes. Higher score = better.

Table 1 shows the MRR@10 performance for the different systems. Increasing the index naturally decreases the performance for all systems, as retrieving the correct passages from a larger index is more challenging. The dense approach trained without hard negatives clearly outperforms BM25 for an index with 10k - 1M entries, but with all 8.8 million passages it performs worse than BM25.

Table 2 shows the relative error rate in comparison to BM25 retrieval. For small index sizes, we

<sup>2</sup><https://www.SBERT.net>

Model	10k	100k	1M	8.8M
Trained without hard negatives				
128 dim	62.3	86.8	100.6	102.2
256 dim	55.7	80.9	97.3	100.6
768 dim	54.9	80.1	96.5	100.3
Trained with hard negatives				
128 dim	48.2	61.1	76.1	88.1
256 dim	44.3	58.4	74.3	87.1
768 dim	42.5	57.0	73.4	86.7

Table 2: Relative error rate (%) of dense approaches in comparison to BM25 retrieval. Lower score = better.

observe that dense approaches drastically reduce the error rate compared to BM25 retrieval. With increasing index sizes, the gap closes.

## 5.2 Index with Random Noise

MS MARCO is sparsely labeled, i.e., there is usually only a single passage labeled as relevant even though multiple passages would be considered as relevant by humans (Craswell et al., 2020). To avoid that the drop in performance is due to the retrieval of relevant, but unlabeled passages, we perform an experiment where we add random irrelevant noise to the index. Our index consists only of the relevant passages and a large fraction of irrelevant, randomly generated strings.<sup>3</sup>

We also evaluate the popular DPR system by Karpukhin et al. (2020), which is a BERT-based dense retriever trained on the Natural Questions (NQ) dataset (Kwiatkowski et al., 2019). We chose the NQ dev set, consisting of 1772 questions from Google search logs. DPR encodes the passage as Title [SEP] Paragraph. We create a random string for the paragraph and combine it with 1) a randomly generated string as title, 2) selecting randomly one of the over 6 Million real Wikipedia article titles, 3) selecting randomly one of the 1772 article titles found in the NQ dev set.

We count for how many queries a random string is ranked higher than the relevant passage. The results are shown in Table 3. We observe that BM25 does not rank any randomly generated passage higher than the relevant passage for the MS MARCO dataset. The chance that a random passage contains words matching the query is small.

For the dense retrieval models, we observe for quite a large number of queries that a random string passage is ranked higher than the relevant passage. As proven in Section 3, the error increases with larger index sizes and fewer dimensions.

<sup>3</sup>Strings are generated randomly using lowercase characters a-z and space.

Model	100k	1M	10M	100M
BM25	0.00%	0.00%	0.00%	0.00%
Dense without hard negatives - MS MARCO				
128 dim	2.71%	4.41%	6.69%	9.73%
256 dim	2.39%	4.03%	6.16%	9.04%
768 dim	2.13%	3.72%	5.77%	8.52%
Dense with hard negatives - MS MARCO				
128 dim	2.87%	4.20%	6.00%	8.11%
256 dim	2.45%	3.72%	5.59%	7.38%
768 dim	2.12%	3.32%	5.09%	7.03%
DPR (Karpukhin et al., 2020) - Natural Questions				
rnd title	0.17%	0.28%	0.34%	0.51%
all titles	2.48%	5.59%	9.31%	12.08%
dev titles	4.18%	5.36%	6.66%	8.01%

Table 3: Percentage of queries for which a random string passage is ranked higher than the relevant passage. 100k/1M/10M/100M indicates the number of random passages in the index.

For DPR, we observe an extreme dependency on the title. Having 100 million entries in the index with a real Wikipedia article title and a random paragraph, results in the retrieval of those for about 12.08% of all questions at the top position.

The error numbers far exceed the estimation from equation (1), confirming that the representations are not uniformly distributed over the complete vector space and are concentrated in a small space. In the appendix (Figure 1), we plot the representations for the queries, the relevant passages, and the random strings.

## 6 Conclusion

We have proven and shown empirically that the probability for false positives in dense information retrieval depends on the index size and on the dimensionality of the used representations. These approaches can even retrieve completely irrelevant, randomly generated passages with high probability. It is important to understand the limitations of dense retrieval:

- 1) Dense approaches work better for smaller, clean indexes. With increasing index size the difference to sparse approaches can decrease.
- 2) Evaluation results with smaller indexes cannot be transferred to larger index sizes. A system that is state-of-the-art for an index of 1 million documents might perform badly on larger indices.
- 3) The false positive rate increases with fewer dimensions.
- 4) The empirically found error rates far exceeded the mathematical lower-bound error rates, indicating that only a small fraction of the available vector space is effectively used.

## Acknowledgments

This work has been supported by the German Research Foundation through the German-Israeli Project Cooperation (DIP, grant DA 1600/1-1 and grant GU 798/17-1) and has been funded by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

## References

- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. 2018. **MS MARCO: A Human Generated MACHine Reading COMprehension Dataset**. *arXiv preprint arXiv:1611.09268 v3*.
- Adam Berger, Rich Caruana, David Cohn, Dayne Freitag, and Vibhu Mittal. 2000. **Bridging the Lexical Chasm: Statistical Approaches to Answer-Finding**. In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '00*, page 192–199, New York, NY, USA. Association for Computing Machinery.
- Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M. Voorhees. 2020. **Overview of the TREC 2019 deep learning track**. *arXiv preprint arXiv:2003.07820 v2*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. **BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding**. *arXiv preprint arXiv:1810.04805*.
- Kawin Ethayarajh. 2019. **How contextual are contextualized word representations? comparing the geometry of BERT, ELMo, and GPT-2 embeddings**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics.
- Luyu Gao, Zhuyun Dai, Tongfei Chen, Zhen Fan, Benjamin Van Durme, and Jamie Callan. 2020. **Completing Lexical Retrieval with Semantic Residual Embedding**. *arXiv preprint arXiv:2004.13969*.
- Daniel Gillick, Alessandro Presta, and Gaurav Singh Tomar. 2018. **End-to-End Retrieval in Continuous Space**. *arXiv preprint arXiv:1811.08008*.
- Mandy Guo, Yinfei Yang, Daniel Cer, Qinlan Shen, and Noah Constant. 2020. **MultiReQA: A Cross-Domain Evaluation for Retrieval Question Answering Models**. *arXiv preprint arXiv:2005.02507*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. **REALM: Retrieval-Augmented Language Model Pre-Training**. *arXiv preprint arXiv:2002.08909*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. **Dense passage retrieval for open-domain question answering**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. **Natural Questions: a Benchmark for Question Answering Research**. *Transactions of the Association of Computational Linguistics*.
- Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020. **On the Sentence Embeddings from Pre-trained Language Models**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9119–9130, Online. Association for Computational Linguistics.
- S. Li. 2011. **Concise formulas for the area and volume of a hyperspherical cap**. *Asian Journal of Mathematics and Statistics*, 4:66–70.
- Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2020. **Sparse, Dense, and Attentional Representations for Text Retrieval**. *arXiv preprint arXiv:2005.00181*.
- Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. 2018. **Umap: Uniform manifold approximation and projection**. *The Journal of Open Source Software*, 3(29):861.
- Aäron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. **Representation Learning with Contrastive Predictive Coding**. *arXiv preprint arXiv:1807.03748*.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. **DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter**. *arXiv preprint arXiv:1910.01108 v4*.

## A Plot of Random Noise Index

Figure 1 shows a two-dimensional plot of the 6980 development queries in the MS MARCO passage dataset, together with the 7433 passages that are marked as relevant and 7433 representations for randomly generated strings (using lowercase characters and space with a random length between 20 and 150 characters). The representation for the random strings are concentrated, but we still observe a significant overlap with the region for queries and relevant documents. This explains why random strings are retrieved for certain queries (Table 3). We use the dense model that was trained with hard negatives with 768 dimensions. UMAP (McInnes et al., 2018) is used for dimensionality reduction to 2 dimensions.




Figure 1: Plot of queries (blue), the relevant document (green) and representations from randomly generated strings (red). Dimensionality reduction via UMAP (McInnes et al., 2018). Model with hard negatives, 768 dimensions.