

On Single and Multiple Representations in Dense Passage Retrieval

Craig Macdonald¹, Nicola Tonellotto² and Iadh Ounis¹

¹University of Glasgow, UK

²University of Pisa, Italy

Abstract

The advent of contextualised language models has brought gains in search effectiveness, not just when applied for re-ranking the output of classical weighting models such as BM25, but also when used directly for passage indexing and retrieval, a technique which is called dense retrieval. In the existing literature in neural ranking, two dense retrieval families have become apparent: single representation, where entire passages are represented by a single embedding (usually BERT's [CLS] token, as exemplified by the recent ANCE approach), or multiple representations, where each token in a passage is represented by its own embedding (as exemplified by the recent ColBERT approach). These two families have not been directly compared. However, because of the likely importance of dense retrieval moving forward, a clear understanding of their advantages and disadvantages is paramount. To this end, this paper contributes a direct study on their comparative effectiveness, noting situations where each method under/over performs w.r.t. each other, and w.r.t. a BM25 baseline. We observe that, while ANCE is more efficient than ColBERT in terms of response time and memory usage, multiple representations are statistically more effective than the single representations for MAP and MRR@10. We also show that multiple representations get better improvements than single representations for queries being the hardest for BM25, as well as for definitional queries, and those with complex information needs.

1. Introduction

Pre-trained contextualised language models such as BERT have been shown to greatly improve retrieval effectiveness over the previous state-of-the-art methods in many information retrieval (IR) tasks [1]. These contextualised language models are able to learn semantic representations called *embeddings* from the contexts of words and, therefore, better capture the relevance of a document w.r.t. a query, with substantial improvements over the classical approach in the ranking and re-ranking of documents [2]. Most BERT-based models are computationally expensive for estimating query-document similarities in ranking, due to the complexity of the underlying transformer neural network [3, 4, 5]. As such, BERT-based ranking models have been used as second-stage rankers in retrieval cascades, in particular to re-rank candidate documents generated by classical relevance models such as BM25 [6, 7, 8]. BERT-based models are also limited in the length of text that they can process, and hence are often applied on passages rather than full documents (which we focus on in this paper); entire document rankings can be obtained by estimating relevance at a passage level, then aggregating [9].

IIR 2021: The 11th Italian Information Retrieval Workshop, September 13–15, 2021, Bari, Italy



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

Recently, several works have proposed investigating whether BERT-based systems are able to identify the relevant passages among all passages in a collection, rather than just among a query-dependent sample; these systems represent a new type of retrieval approaches called *dense retrieval*. In dense retrieval, passages are represented by real-valued vectors, while the query-document similarity is computed by deploying efficient nearest neighbour techniques over specialised indexes, such as those provided by the FAISS toolkit [10]. Thus far, two different families of dense retrieval approaches have emerged recently, based on *single representation* and *multiple representation*. In particular, DPR [11] and ANCE [12] use a single representation, by indexing only the embedding of BERT’s [CLS] token, and therefore this is assumed to represent the meaning of an entire passage within that single embedding. At retrieval time, the [CLS] embedding of the query is then used to retrieve passages by identifying nearest neighbours using a FAISS index. In contrast, ColBERT [3], which uses multiple representation, indexes an embedding for each token in each document. At retrieval time, a set of the nearest document embeddings to each query embedding is retrieved, by identifying the approximate nearest neighbours from a FAISS index. These passages must then be exactly scored, based on the maximal similarity between the query and the passage embeddings, to obtain the final ranking.

These are two markedly different families of dense retrieval approaches. Indeed, as ColBERT records one embedding for every token, this makes for a large index of embeddings, which may allow a richer semantic representation of the content. On the other hand, DPR and ANCE rely on a single embedding sufficiently representing the content of each passage. However, at the time of writing, no systematic study has compared these two families of dense retrieval approaches. For this reason, this work contributes a first investigation into the effectiveness of single and multiple representation embeddings for dense retrieval, as exemplified by ANCE and ColBERT, respectively. We perform experiments in a controlled environment using the same collection and query sets, and we report several effectiveness metrics, together with a detailed comparison of the results obtained for the two representation families w.r.t. a common baseline, namely BM25. To derive further insights, we also provide a per-query analysis of the effectiveness of single and multiple representations. We observe that, while ANCE is more efficient than ColBERT in terms of response time and memory usage, multiple representations are statistically more effective than the single representations for MAP and MRR@10. We also show that multiple representations obtain better improvements than single representations for queries that are the hardest for BM25, as well as for definitional queries, and those with complex information needs.

2. Problem Statement

Embeddings. Contextualized language models such as BERT have been trained on a large corpus for language understanding, and then fine-tuned on smaller, more specific textual collections targeting a particular IR task. Through this fine-tuning, BERT learns how to map texts, either queries or documents, into a multi-dimensional space of one or more vectors, called *embeddings*. Both queries and documents are tokenised into terms according to a predefined vocabulary; BERT learns a function mapping tokens in a query into multiple query embeddings, one per query term, and another potentially different function mapping tokens in a document

into a document embedding per term. BERT and its derived models also make use of special tokens, such as the [CLS] (classification) token, the [SEP] (separator) token, and the [MASK] (masked) token. In particular, [CLS] is always placed at the beginning of any text given as input to BERT, both at training and inference time, and is used to let BERT learn a global representation of the input text as a single embedding. In more detail, a text composed by m terms given as input to BERT will produce $m + 1$ embeddings, one per input term plus one additional embedding for [CLS]. In *single representation* models, such as ANCE, the embedding corresponding to [CLS] is assumed to encode all possible information about the input text, including the possible semantic context of the composing terms. In contrast, in *multiple representation* models, such as ColBERT, each input term’s embedding encodes its specific semantic information within the context of the entire input text.

Dense Retrieval. The embeddings produced by the BERT models have recently demonstrated their promise in being a suitable basis for *dense retrieval*. In dense retrieval, documents and queries are represented using embeddings. The embeddings from the documents in a collection can be pre-computed through the application of the BERT learned mapping and stored into an index data structure for embeddings supporting nearest neighbour similarity search, as exemplified by the FAISS toolkit [10]. Depending on the number and dimensions of the embeddings stored into the index, advanced compression strategies, together with suitable nearest neighbour search algorithms, can be employed. In order to reduce the time required to identify the most similar document embeddings to a given input embedding, it is possible to shift from exact nearest neighbour search to approximate nearest neighbour search. While ANCE stores embeddings in an uncompressed format supporting exact search, ColBERT, given the larger number of document embeddings it has to store, resorts to compressed and quantised embeddings supporting approximate search. However, the approximate similarity scores produced by approximate search are not used by the ColBERT implementation to compute the final top documents to return for a given query [13]¹. Hence, ColBERT uses approximate search over compressed embeddings to identify a candidate *set* of documents, which are then re-scored using an index with direct lookup for retrieving the candidate documents’ embeddings, to obtain the final ranking of documents returned to the user.

Research Questions. In this work, we aim to compare the single and the multiple embedding representations leveraging the ANCE and ColBERT implementations. Indeed, there was not an effectiveness comparison with ColBERT in the ANCE paper [12]. ANCE embodies a recent single representation approach, where we have a single large embedding per query/document, which can be processed with exact similarity search in a single stage. In contrast, in the multiple representation approach (ColBERT), we have a smaller sized embedding for each term in the queries/documents, but due to the large number of embeddings, they must be processed using approximate similarity search. Thereafter the candidate set must be re-ranked to compute the exact similarity scores. The need for ColBERT to re-score all documents in the candidate set necessitates storing all document embeddings in memory. As noted by Lin et al. [2], this presents a significant storage overhead. This underlines the importance of an in-depth analysis

¹Indeed, in [13] we show that these approximate scores can allow a high recall but low precision ranking to be obtained, which can be used to apply rank cutoffs to the the candidate set.

of the pros and cons of both approaches, in particular:

- **RQ1.** What is the effectiveness of single and multiple representations in dense retrieval, in terms of MAP, NDCG@10, MRR@10?
- **RQ2.** What are the relative gains and losses of single and multiple representations w.r.t. a common baseline such as BM25?
- **RQ3.** For which queries are single representations better than multiple representations, and vice-versa?

In Section 3, we perform comparative experiments to address these research questions.

3. Experiments

In the following, we report our experimental setup, followed by analyses for RQs 1-3.

3.1. Setup

Our experiments use the MSMARCO passage ranking dataset, a dataset of 8.8M passages and build upon our PyTerrier IR experimentation platform [14, 15]. We adapt the ANCE implementation² and the ColBERT implementation³ provided by their respective authors, using integrations with PyTerrier⁴. We use the provided ANCE model for the MSMARCO passage ranking dataset. We train ColBERT using the same MSMARCO passage ranking training triples file for 44,500 batches. In particular, we follow [12] and [3] for the settings of ANCE and ColBERT, as summarised in Table 1.

Of note, while ColBERT fine-tunes the `bert-base-uncased` BERT model, ANCE fine-tunes a RoBERTa model [16] (specifically `roberta-base`), which is reported to apply more refined pre-training than BERT. To try to eliminate model choice as a confounding factory, we also trained a version of ColBERT by fine-tuning `roberta-base`. We found that even after training for 300k batches ($6\times$ longer than we trained ColBERT using `bert-base-uncased`), this latter model could had relative performance 25% less than the BERT-based ColBERT model (around NDCG@10 of 0.533). Hence we discarded the RoBERTa-based ColBERT model from further consideration. On the other hand, all of the released ANCE models use RoBERTa; training ANCE requires multiple GPUs, e.g., 16, and has not, to the best of our knowledge, been reproduced. Hence, we argue that as the RoBERTa-based ANCE and BERT-based ColBERT are individually shown to be effective by their respective authors, the comparison of these representative models still allows for interesting observations.

We index the corpus using the code provided by the authors. Table 1 reports the statistics of the resulting indices. In particular, the ANCE document index is stored in FAISS using the uncompressed `IndexFlatIP` format. The ColBERT document index is stored in FAISS using the compressed and quantised `IndexIVFPQ` format, which is trained on a random 5% sample of the document embeddings. Mean response times for both ANCE and ColBERT, and their memory consumption, are also shown in Table 1.

²<https://github.com/microsoft/ANCE>

³<https://github.com/stanford-futuredata/ColBERT/tree/v0.2>

⁴See https://github.com/terrierteam/pyterrier_ance and https://github.com/terrierteam/pyterrier_colbert.

Table 1

Salient statistics of the ANCE and ColBERT setups.

	ANCE	ColBERT
Representation	single	multiple
Base model	roberta-base	bert-base-uncased
# emb. per query	1	32
# emb. per passage	1	up to 180
Emb. dimensions	768	128
FAISS index size	26GB	16GB
Embedding index size	–	176GB
Mean Response Time	211ms	635ms

For evaluating effectiveness, we use the publicly available query sets with relevance assessments: 5000 queries sampled from the MSMARCO Dev set – which contain on average 1.1 judgements per query – as well as the TREC 2019 query set, which contains 43 queries with an average of 215.3 judgements per query. To measure effectiveness, we employ MRR@10 for the MSMARCO Dev set⁵, and the MRR@10, NDCG@10 and MAP for the TREC query set.

To examine gains and losses, for each query and each effectiveness metric, we examine the comparative reward (improvement) and risk (degradation) over a BM25 baseline (following [17]), as well as the number of wins & losses (improved and degraded queries).

3.2. Overall Comparison

Table 2 reports the effectiveness metrics of BM25, ANCE and ColBERT computed on the TREC 2019 and the sample of the MSMARCO Dev query sets. As expected, both the ANCE and ColBERT dense retrieval approaches are significantly better than BM25 for the NDCG@10 and MRR@10 metrics on both query sets. Comparing the two dense retrieval approaches, for MAP, ColBERT significantly outperforms ANCE; for NDCG@10, ColBERT enhances ANCE by 6% (0.6537→0.6934), but not significantly so; for MRR@10, ANCE is slightly (but not significantly) better than ColBERT on the TREC2019 query set while ColBERT is statistically better than ANCE on MSMARCO Dev by +7%. Overall, for RQ1, we conclude that multiple representations, employed by ColBERT, experimentally obtain better effectiveness than single representations (as employed by ANCE), exhibiting significant boost in effectiveness for MAP (TREC 2019) and MRR@10 (Dev). Among the most striking differences is that for MAP on TREC 2019, where ColBERT markedly outperforms ANCE (and BM25); this observation suggests that the single representation is not sufficiently good at attaining high recall.

3.3. Comparison using a Common Baseline

Next, we investigate the comparative effectiveness of ANCE and ColBERT from the perspective of using BM25 as the reference point, going further than reporting average performances over the entire query sets as reported in Table 2. To perform this analysis, we define the difficulty of

⁵This is the metric recommended by the track organisers for this query set.

Table 2

Effectiveness metrics of BM25, ANCE and ColBERT on different query sets. Points marked with \triangle and \blacktriangle denote a significant increase in effectiveness compared to BM25 and ANCE, respectively, according to a paired t-test with Bonferroni correction (p-value < 0.05).

	TREC 2019			MSMARCO Dev
	MAP	NDCG@10	MRR@10	MRR@10
BM25	0.2864	0.4795	0.6410	0.1836
ANCE	0.3715 \triangle	0.6537 \triangle	0.8574 \triangle	0.3292 \triangle
ColBERT	0.4309 $\blacktriangle\triangle$	0.6934 \triangle	0.8527 \triangle	0.3519 $\triangle\blacktriangle$

a query according to an effectiveness metric on the BM25 baseline, following Mothe et al. [18]. Due to the sparsity of the relevance judgements and the official evaluation metrics of the two query sets, we adopt a different query difficulty classification for TREC 2019 and MSMARCO Dev. For the TREC 2019 query set, a query is considered *hard*, resp. *easy*, for the BM25 baseline system if the NDCG@10 (the official TREC metric in [19]) value is in the first quartile, resp. in the fourth quartile, and *medium* otherwise. For the MSMARCO query set, the official metric MRR@10 per query is too sparse to allow percentile computations. Hence we consider a Dev query to be *hard* if its MRR@10 is lesser than or equal to 0.1, and *easy* otherwise.

We partition the queries in each query set according to the corresponding difficulty classification, and compute for how many queries the effectiveness of ANCE and ColBERT is higher (denoted by W(in)) or lower (denoted with L(oss)) than BM25. For each partition, we also compute the average reward and risk associated with the W and L queries, following [17].

Table 3 reports the observed results. For the TREC 2019 queries, both ANCE and ColBERT exhibit approx. the same number of wins/losses for each query difficulty level. However, ANCE obtains higher rewards and higher risks on the class of easy queries than ColBERT (+0.1930 vs. +0.1827 and -0.1976 vs. -0.1380). On the medium difficulty class, the situation is reversed, and ColBERT obtains both higher rewards and higher risks than ANCE (+0.3053 vs. +0.2978 and -0.1521 vs. 0.1366). On the hard difficulty class, ColBERT is markedly superior to ANCE in terms of reward (+0.4114 vs. +0.3750), and risk, even if such risk is computed over a single query. For the MSMARCO Dev queries, ColBERT is able to improve the MRR@10 of both easy and hard queries better than ANCE, and the losses are smaller for ColBERT than for ANCE.

To conclude on RQ2, we have presented experimental evidence that both single and multiple representations are approximately as effective on easy queries. In contrast, for hard queries, the adoption of multiple embeddings helps w.r.t the usage of a single embedding. We explain this by noting that a single representation is learned to compress all semantic information and dependencies of the different tokens composing a query in a single embedding. On the other hand, multiple representations – using one embedding per query token together with additional masked tokens – can encode more diverse semantic information in the different embeddings, allowing to retrieve more relevant documents for queries that are hard to answer.

Table 3

Comparative performances w.r.t. BM25; queries are classified based on BM25 performance (easy/medium/hard); Wins and Losses as well as Reward and Risk are calculated w.r.t. BM25 performance.

Type	Num	ANCE		ColBERT	
		W/L	reward/risk	W/L	reward/risk
TREC 2019 – NDCG@10					
Easy	11	5/6	+0.1930/-0.1976	5/6	+0.1827/-0.1380
Medium	21	17/4	+0.2978/-0.1366	18/3	+0.3053/-0.1521
Hard	11	9/1	+0.3750/-0.1826	10/1	+0.4114/-0.0415
MSMARCO Dev – MRR@10					
Easy	1954	828/712	+0.4735/-0.4001	854/673	+0.4778/-0.3887
Hard	3076	1372/24	+0.4543/-0.1	1455/21	+0.4793/-0.1




Figure 1: Difference in NDCG@10 for queries in the TREC 2019 query set; Differences smaller than 0.15 absolute are omitted. +ve differences are where ColBERT exceeds ANCE.

3.4. Per-query Comparison

To address RQ3, in Figure 1 we present a per-query histogram comparing the $\Delta\text{NDCG@10}$ between ColBERT and ANCE on the TREC 2019 query set; positive deltas indicate a higher NDCG@10 for ColBERT than ANCE. In total, ColBERT outperform ANCE for 24 queries, while the opposite was true for 17 queries; Δ s less than 0.15 are omitted for clarity. On analysing Figure 1, it appears that many queries requesting a *definition* appear to perform well for ColBERT (e.g. 1124210, 490595). Indeed, on closer inspection of the TREC 2019 query set, out of 43 queries, we found 19 such definitional queries – of which 16 were more effective for ColBERT.

To illustrate other differences between the approaches, in Table 4 we select two non-definitional queries where one approach markedly outperformed the other (but not the queries with the

Table 4

Examples of passages retrieved by ANCE and ColBERT at top ranks. The Label column contains the assessment of that document for that query in the qrel file, with – denoting unjudged.

Document	System	Rank	Passage	Label
ColBERT > ANCE				
527433: types of dysarthria from cerebral palsy				
8617271	ColBERT	1	There are three major types of dysarthria in cerebral palsy: spastic, dyskinetic (athetosis) and ataxic. Speech impairments in spastic dysarthria involves four major abnormalities of voluntary movement: spasticity, weakness, limited range of motion and slowness of movement.	3
8306451	ANCE	2	The types of cerebral palsy are: 1 spastic: the most common type of cerebral palsy; reflexes are exaggerated and muscle movement is stiff. 2 dyskinetic: dyskinetic cerebral palsy is divided into two categories.	0
ANCE > ColBERT				
1063750: why did the us volunterilay enter ww1				
1300452	ColBERT	2	The main event that led the US to entering ww2 was Japan bombing Pearl Harbor. The day after the bombing u.s. joined the war On December 7, 1941, the Japanese Navy launched a surprise attack on the naval base at Pearl Harbor, Hawaii. Although the growing peril of Britain worried many, including Roosevelt, it was not until the US was directly attacked at Pearl Harbor that public and political opinion turned in favor of war with the Axis	–
7952971	ANCE	1	The U.S entered WW1 for several reasons. The U.S entered for two main reasons: one was that the Germans had declared unlimited German submarine warfare and the Zimmermann note. The German had totally disregarded the international laws protecting neutral nation's ships by sinking neutral ships. his note was the last straw, causing Wilson to join the war. The Zimmermann note and unlimited German submarine warfare were two of the biggest cause for the U.S to join the Allies and go to war with Germany. During the war Germany...	2

most extreme deltas, which may be outliers). Firstly, for query 527433 ('types of dysarthria from cerebral palsy'), ColBERT identifies a passage that clearly answers the query; in contrast, the non-relevant passage identified at rank 2 by ANCE appears to have focused solely on the 'cerebral palsy' aspect, omitting the dysarthria aspect of the query. Indeed, the Precision @10 of ANCE for this query was $\frac{3}{10}$, compared to $\frac{6}{10}$ for ColBERT. This suggests that ANCE's compression of a complex information need into one embedding has caused an information loss, with the model focusing on only a single aspect of the query, resulting in low effectiveness.

On the other hand, for query 1063750 ('why did the us volunterilay enter ww1'), ANCE identified a relevant passage, but ColBERT identified a passage (1300452) focusing entirely on the wrong World War ('ww2' rather than 'ww1'). At least some of the reason for the conflation of meanings is that neither 'ww1' nor 'ww2' do not appear in BERT's fixed vocabulary, e.g., the latter is tokenised into word pieces as 'w', '##w', '##2'. Hence distinguishing between 'ww1' and 'ww2' information needs require context to be distributed across the three embeddings. To analyse this passage further, Figure 2 shows the ColBERT interaction between the query and document embeddings for this passage and query⁶. In the figure, the darker shading in the matrix is indicative of higher similarity; the highest similarity that is selected for a given query embedding by the max-sim operator is indicated by a × symbol; the histogram at the top of the

⁶This figure can be reproduced using the `explain_text()` function within our `PyTerrier_ColBERT` library.

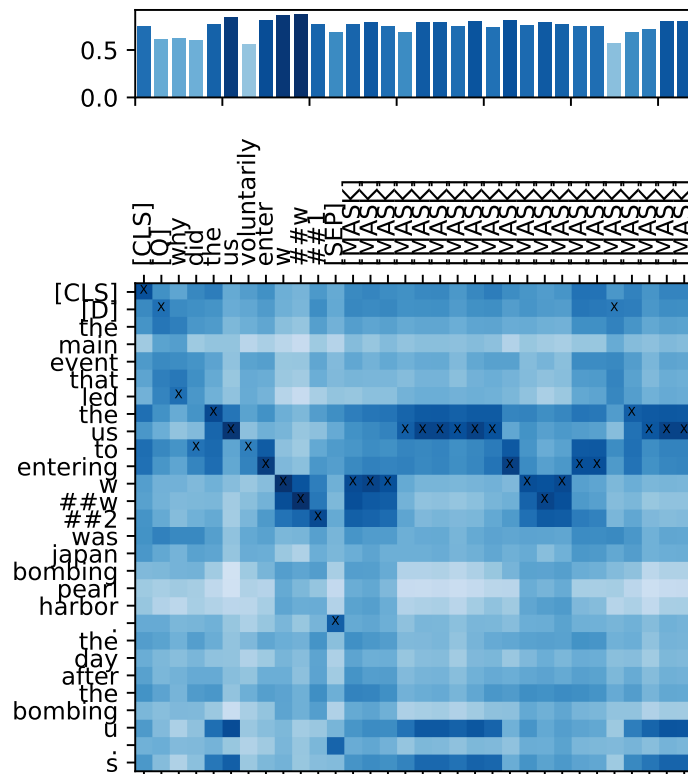


Figure 2: ColBERT interaction between query and document embeddings for query 106375 and passage 1300452 (see Table 4). In the interaction matrix, darker shading is indicative of higher similarity; the document embedding (row) with highest similarity for each query embedding (column) is indicated with a \times symbol. The histogram at the top portrays the contribution of each query embedding to the final score of the passage, with shading also indicative of the magnitude of contribution.

figure indicates the contribution of each query embedding to the final passage score. Indeed, on inspection of the max similarities for this passage shows that the highest contributions to the passage’s score comes from the ‘##w’ token, with ‘##1’ query embedding being highly similar to the ‘##2’ document embedding. This suggests that the embeddings for ‘##1’ and ‘##2’ are not sufficiently contextualised when following ‘##w’, or that ColBERT’s max similarity computation could be adapted to better address proximity. In contrast, ANCE retrieved passage 1300452 at rank 155, showing that the single representations for the passages sufficiently distinguish between World War 1 vs. World War 2.

In summary, in addressing RQ3, we observed that there exists some large differences between ANCE and ColBERT for some queries. Our analysis found that ColBERT performs better than ANCE for definitional type queries. Moreover, our analysis suggests that in ANCE, the use of a single embedding representation risks misinterpreting complex queries with multiple aspects as shown by results in the previous subsection; For ColBERT, the max similarity operator can overly focus on highly similar embeddings at the risk of mis-interpreting a query.

4. Conclusions

Despite their recency, dense passage retrieval approaches have the effectiveness potential to supplant the traditional inverted index data structure. Yet, different families of dense retrieval are emerging, for which the comparative advantages and disadvantages are not yet clear. In this work, we made a systematic study of single vs. multiple representation dense retrieval approaches, namely ANCE and ColBERT. We found that while both significantly outperformed BM25 baselines across various metrics, ColBERT significantly outperformed ANCE for MAP on TREC2019 and MRR@10 on the MSMARCO Dev query set, was more effective for queries that BM25 found hard, and was better at definitional queries as well as queries that had complex information needs. On the other hand, ANCE has desirable qualities in terms of mean response time and memory occupancy (see Table 1). We postulate that research should be directed toward hybrid solutions, either reducing the size of the ColBERT embedding index, e.g., through adaptations to static pruning, or through using multiple embeddings within ANCE for complex queries/passages.

Acknowledgements

Nicola Tonello was partially supported by the Italian Ministry of Education and Research (MIUR) in the framework of the CrossLab project (Departments of Excellence). Craig Macdonald and Iadh Ounis acknowledge EPSRC grant EP/R018634/1: Closed-Loop Data Science for Complex, Computationally- & Data-Intensive Analytics.

References

- [1] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, in: Proc. NAACL, 2019.
- [2] J. Lin, R. Nogueira, A. Yates, Pretrained transformers for text ranking: BERT and beyond, 2020. [arXiv:2010.06467](https://arxiv.org/abs/2010.06467).
- [3] O. Khattab, M. Zaharia, ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT, in: Proc. SIGIR, 2020, p. 39–48.
- [4] S. Hofstätter, A. Hanbury, Let’s measure run time! Extending the IR replicability infrastructure to include performance aspects, in: OSIRRC@SIGIR, 2019.
- [5] H. Zamani, M. Dehghani, W. B. Croft, E. Learned-Miller, J. Kamps, From neural re-ranking to neural ranking: Learning a sparse representation for inverted indexing, in: Proc. CIKM, 2018, pp. 497–506.
- [6] S. MacAvaney, A. Yates, A. Cohan, N. Goharian, CEDR: Contextualized embeddings for document ranking, in: Proc. SIGIR, 2019, pp. 1101–1104.
- [7] S. MacAvaney, F. M. Nardini, R. Perego, N. Tonello, N. Goharian, O. Frieder, Efficient document re-ranking for transformers by precomputing term representations, in: Proc. SIGIR, 2020, pp. 49–58.
- [8] S. MacAvaney, F. M. Nardini, R. Perego, N. Tonello, N. Goharian, O. Frieder, Expansion via prediction of importance with contextualization, in: Proc. SIGIR, 2020, p. 1573–1576.

- [9] Z. Dai, J. Callan, Deeper text understanding for IR with contextual neural language modeling, in: Proc. SIGIR, 2019, pp. 985–988.
- [10] J. Johnson, M. Douze, H. Jégou, Billion-scale similarity search with GPUs, 2017. [arXiv:1702.08734](https://arxiv.org/abs/1702.08734).
- [11] V. Karpukhin, B. Oguz, S. Min, P. Lewis, L. Wu, S. Edunov, D. Chen, W.-t. Yih, Dense passage retrieval for open-domain question answering, in: Proc. EMNLP, 2020, pp. 6769–6781.
- [12] L. Xiong, C. Xiong, Y. Li, K.-F. Tang, J. Liu, P. Bennett, J. Ahmed, A. Overwijk, Approximate nearest neighbor negative contrastive learning for dense text retrieval, in: Proc. ICLR, 2021.
- [13] C. Macdonald, N. Tonello, On approximate nearest neighbour selection for multi-stage dense retrieval, in: Proc. CIKM, 2021.
- [14] C. Macdonald, N. Tonello, Declarative experimentation in information retrieval using PyTerrier, in: Proc. ICTIR, 2020, pp. 161–168.
- [15] C. Macdonald, N. Tonello, S. MacAvaney, I. Ounis, PyTerrier: Declarative experimentation in Python from BM25 to dense retrieval, in: Proc. CIKM, 2021.
- [16] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, RoBERTa: A Robustly Optimized BERT Pretraining Approach, 2019. [arXiv:1907.11692](https://arxiv.org/abs/1907.11692).
- [17] L. Wang, P. N. Bennett, K. Collins-Thompson, Robust ranking models via risk-sensitive optimization, in: Proc. SIGIR, 2012, p. 761–770.
- [18] J. Mothe, L. Laporte, A.-G. Chifu, Predicting Query Difficulty in IR: Impact of Difficulty Definition, in: Proc. KSE, 2019, pp. 1–6.
- [19] N. Craswell, B. Mitra, D. Campos, E. Yilmaz, Overview of the TREC 2019 Deep Learning Track, in: Proc. TREC 2019, 2020.