

A Static Pruning Study on Sparse Neural Retrievers

Carlos Lassance
carlos.lassance@naverlabs.com
Naver Labs Europe
Meylan, France

Hervé Déjean
herve.dejean@naverlabs.com
Naver Labs Europe
Meylan, France

Simon Lupart
simon.lupart@naverlabs.com
Naver Labs Europe
Meylan, France

Stéphane Clinchant
stephane.clinchant@naverlabs.com
Naver Labs Europe
Meylan, France

Nicola Tonellotto
nicola.tonellotto@unipi.it
University of Pisa
Pisa, Italy

ABSTRACT

Sparse neural retrievers, such as DeepImpact, uniCOIL and SPLADE, have been introduced recently as an efficient and effective way to perform retrieval with inverted indexes. They aim to learn term importance and, in some cases, document expansions, to provide a more effective document ranking compared to traditional bag-of-words retrieval models such as BM25. However, these sparse neural retrievers have been shown to increase the computational costs and latency of query processing compared to their classical counterparts. To mitigate this, we apply a well-known family of techniques for boosting the efficiency of query processing over inverted indexes: static pruning. We experiment with three static pruning strategies, namely document-centric, term-centric and agnostic pruning, and we assess, over diverse datasets, that these techniques still work with sparse neural retrievers. In particular, static pruning achieves $2\times$ speedup with negligible effectiveness loss ($\leq 2\%$ drop) and, depending on the use case, even $4\times$ speedup with minimal impact on the effectiveness ($\leq 8\%$ drop). Moreover, we show that neural rerankers are robust to candidates from statically pruned indexes.

CCS CONCEPTS

• Information systems → Information retrieval.

KEYWORDS

neural sparse retrievers, static pruning, inverted indexes, SPLADE

ACM Reference Format:

Carlos Lassance, Hervé Déjean, Simon Lupart, Stéphane Clinchant, and Nicola Tonellotto. 2023. A Static Pruning Study on Sparse Neural Retrievers. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*, July 23–27, 2023, Taipei, Taiwan. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3539618.3591941>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '23, July 23–27, 2023, Taipei, Taiwan

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9408-6/23/07...\$15.00

<https://doi.org/10.1145/3539618.3591941>

1 INTRODUCTION

Pre-trained language models such as BERT [14] have been shown to significantly improve the effectiveness of Information Retrieval (IR) systems over traditional ranking models. Dense retrieval is one of the most promising applications of pre-trained language models in IR. In dense retrieval, a pre-trained language model computes dense representations of queries and documents, that must be stored and processed using approximate nearest neighbor algorithms [17], preferably on GPUs or TPUs [18]. However, the ranking of documents in dense retrieval incurs significant computational costs, even if this can be reduced through approaches such as ColBERT [19, 35], or by re-ranking only a small set of candidates [29]. Nevertheless, query processing in dense retrieval is much more expensive than processing documents stored in an inverted index with simpler ranking functions and optimized processing algorithms [36].

Sparse neural retrievers [15, 22, 27] aim to merge these two approaches, using a pre-trained language model to compute very compact document representations, up to a single value, to be stored in an inverted index. In doing so, these retrievers try to reach the effectiveness of the complex pre-trained language models, while retaining the efficiency of simple bag-of-words retrieval models and query processing strategies. Most sparse neural retrievers integrate two main learning goals: they learn which terms in a document should be indexed (document content learning), and how to represent these terms in a corresponding index (term impact learning). In document content learning, irrelevant terms are removed and relevant terms are added to the documents. During term impact learning, each term in each document is associated with an impact score that will be stored in a posting of the term's posting list.

Despite the effectiveness and efficiency of sparse neural retrievers w.r.t. more complex neural ranking models, recent works have shown that, while effective retrieval is possible with learned sparse approaches, they are often still much slower than their traditional counterparts [27]. This efficiency loss in neural sparse retrievers has been addressed by introducing modifications to the layout of the posting lists in an inverted index, such as storing both the learned impact and the BM25 score in a single posting [28], splitting a posting list into two separate lists, depending on the impact values of their postings [24] or **accessing documents via clusters for anytime retrieval** [25, 26]. However, the actual content of the inverted index, in terms of terms and postings, does not change. A classical solution to increase the efficiency of query processing over inverted indexes is static pruning [7]. The static pruning of

inverted indexes deals with removing information stored in the inverted index to improve the efficiency of query processing with negligible or minimal negative impact on the effectiveness of the retrieval system. *Agnostic* pruning strategies [4, 7, 8, 31] focus on removing the low-scoring terms and/or documents from the inverted index. In *term-centric* [1, 5, 7, 13, 31] and *document-centric* [1, 6, 34] pruning, the removal of a posting depends on the distribution of the posting scores w.r.t. a posting list and a document, respectively, in which the posting appears in.

In this work, we investigate the impact of different static pruning strategies on inverted indexes built with sparse neural retrievers. Indeed, the new document representations learned by these systems alter the actual content of documents, and the statistical properties of the corresponding postings lists. **We differ from related previous work [23] as we focus on models that not only perform term weighting but allow for learning the entire retrieval process, also including document and query expansion and query weighting.** In particular, we investigate four major research questions:

- (1) Is static pruning effective on sparse neural retrievers on the dataset they were trained on, e.g., MSMARCO?
- (2) Can we extend static pruning of sparse neural retrievers to other datasets, e.g., TripClick, even in zero-shot scenarios?
- (3) Does static pruning affect the quality of reranking?
- (4) Can pruning be combined with other approaches to improve the trade-off between effectiveness and efficiency?

Our experiments demonstrate that static pruning can improve the efficiency of query processing of sparse neural retrievers by a factor of 2× with a negligible loss in effectiveness ($\leq 2\%$) and up to 4× with an acceptable loss of effectiveness ($\leq 8\%$). Moreover, our experiments show that neural rerankers are robust to candidates from statically pruned indexes and that we are able to further push the efficiency/effectiveness first-stage trade-off by combining static pruning with neural network quantization, better pre-trained language models and better distillation data. **Due to space constraints we make Supplementary Material/Appendix available at , in order to provide more information and clarify questions from reviewers.**

2 METHODOLOGY

In the following paragraphs, we define the static pruning methods that we use throughout our experiments, **and we refer the reader to the supplementary material for a summary of more complex methods that did not improve over the reported ones.** These methods can be divided into term-centric (T), document-centric (D), and agnostic (A) pruning. Each method is tested separately to verify its general behavior and assess the potential advantages of a given approach over the other.

Remove low scores from posting lists based on score quantile (T). In T we choose a given quantile for each posting list and remove all values that are smaller than that quantile. For example, if we pick 50% quantile we remove all values equal or smaller than the median value stored on the posting list. In more detail, we test the following quantiles: 50%, 75%, 80%, 85%. The argument for choosing such a method is: if a document-term pair has a score that is on the low end of its respective posting list, there is a small chance for it to be important in score computation, especially if we look into measures with small cutoff values, e.g., MRR@10.

Remove low scores from documents based on a max size (D):. This strategy retains, for each document, the 4, 8, 16, 32, 64 terms with the highest impact values. While in T, pruning is made on the terms posting lists, in D it is made on documents representations before creating posting lists. The argument for the strategy is: the less important terms of the document should be the less useful ones.

Prune low scores (A). Finally, we also verify an approach that is agnostic to both terms and documents, where we simply remove the lowest scores of the inverted index. The idea of this approach is almost the same as for the quantiles-based pruning T, but simplifies the application as it uses the same threshold for all terms. It can also be seen as an extension of L1-regularization-based approaches [32], **as A prunes values that L1 would push close, but not equal, to 0.**

3 EXPERIMENTS

We perform different experiments to assess if static pruning still works for the readily available sparse neural retrievers, such as DeepImpact [27], uniCOIL with TILDev2 expansion [38], and the eff-V-{S/M/L} SPLADE family [20], together with a BM25 baseline. **Experimental details are made available in the Appendix and at the codebase¹.** Statistical significance is measured using paired *t*-tests with $p \leq 0.05$ and Bonferroni's correction.

The sparse neural retrievers we evaluate vary in some design choices, i.e., i) vocabulary, and ii) expansion. DeepImpact uses a white-space tokenized collection vocabulary with documents being expanded via DocT5 [30], but it does not modify the queries. UniCoil uses the sub-word tokenized BERT vocabulary, with documents expanded using TILDev2 [38] and reweighted queries without expansion. Finally, SPLADE also uses the BERT vocabulary but learns both query and document expansions during training using a uniformity regularizer to minimize a *FLOPS* metric [20, 32].

3.1 Static Pruning on MSMARCO

To verify that the pruning methods proposed in Section 2 can be extended to sparse neural retrievers, we apply them on the MSMARCO passage dataset [3]. In this set of experiments, models are both trained and evaluated on MSMARCO, on train and “dev.small” query sets, respectively. We also extend our evaluation to the query sets from TREC-DL [10, 12], namely the 2019 and 2020 editions.

Figure 1 depicts a comparison of speed-up and relative effectiveness (Recall@1k and MRR@10), i.e., the ratio over the unpruned performance, for all models, under all pruning methods. We note that for BM25 (black) all pruning strategies are too aggressive, with only a few points appearing before the effectiveness decreased too much. About the different pruning methods, the term-centric pruning (T) works for the models with BERT vocabulary, i.e., Unicoil (grey) and SPLADE (blue), but it is too aggressive for the DeepImpact model (light red), exploiting a larger vocabulary. For the other two pruning methods, i.e., document-centric and agnostic pruning, the results are very similar, with agnostic pruning (cross marker) having a slight advantage, but it needs finer calibration depending on the model as the term impact statistics are different between models. Finally, we note that overall the models that do not learn to perform document expansion, but use a predefined expansion, i.e. DeepImpact and Unicoil, generate inverted indexes that look easier

¹ URL to github

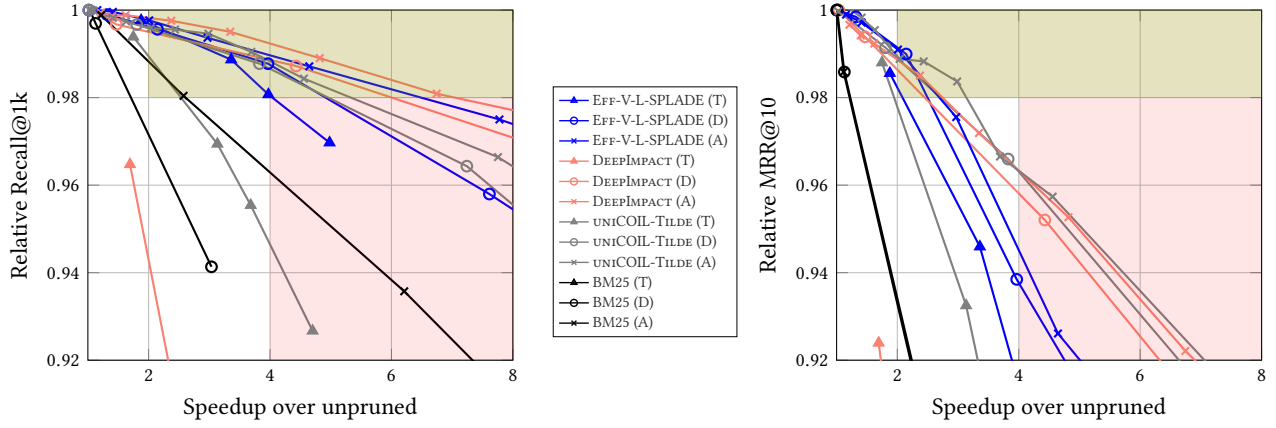


Figure 1: Speedup vs relative effectiveness comparison. Rectangles show the two desired conditions (Green $\geq 2\times$ speedup with $\leq 2\%$ effectiveness drop, Red $\geq 4\times$ speedup with $\leq 8\%$ effectiveness drop).

to prune. In particular, for SPLADE, stronger pruning entails larger MRR@10 losses. In summary, all neural sparse retrievers we test can get a $2\times$ speedup due to pruning without meaningful loss of effectiveness ($\leq 2\%$ drop in effectiveness) and a $4\times$ speedup is possible with negligible impact ($\leq 8\%$ drop), represented by the green and red rectangles. As we verified that all three methods reported similar trade-offs and that it worked for all models, from now on we shift the focus to document-centric pruning over SPLADE, in order to make the presentation more concise.²

In Table 1, we look into the TREC-DL (2019 and 2020) query sets [9, 11]. The TREC-DL query sets have fewer queries than the MSMARCO dev set, but have a richer annotation set. Furthermore, the same results still holds: static pruning allows for around $2\times$ speedup (max size of 64) without meaningful loss of effectiveness ($\leq 2\%$ drop) and around $4\times$ speedup (max size between 16 and 32) is possible with negligible impact ($\leq 8\%$ drop).

Table 1: Experiments on document-centric pruning of eff-V-L SPLADE on the TREC-DL queries (2019 and 2020). The \times value represents the speedup compared to the baseline. \dagger indicates statistical significance compared to the baseline.

Max size	TREC-DL 2019			TREC-DL 2020		
	Latency (\times)	nDCG@10	R@1k	Latency (\times)	nDCG@10	R@1k
30 522	28.2 (1.0 \times)	71.4	84.8	30.2 (1.0 \times)	71.5	86.2
Prune by document size (D)						
64	13.2 (2.1 \times)	69.3 \dagger	83.5 \dagger	13.7 (2.2 \times)	70.4	84.8 \dagger
32	8.8 (3.2 \times)	67.9 \dagger	79.0 \dagger	8.9 (3.4 \times)	67.2 \dagger	81.8 \dagger
16	6.0 (4.7 \times)	64.4 \dagger	74.6 \dagger	6.5 (4.6 \times)	65.0 \dagger	79.4 \dagger
8	3.4 (16.6 \times)	58.82 \dagger	67.0 \dagger	3.5 (8.6 \times)	59.5 \dagger	70.2 \dagger
4	1.7 (16.6 \times)	52.4 \dagger	54.5 \dagger	1.7 (17.8 \times)	48.6 \dagger	60.7 \dagger

Finally, we take advantage of the fact that the SPLADE models come in three “sizes”, namely (S)mall, (M)edium and (L)arge, and look if it is better to generate sparser retrievers by regularization or if pruning by itself is sufficient for achieving a good trade-off

² Full-scale experiments can be found in the Appendix.

between efficiency/effectiveness. The results in Figure 2 for Recall@1k show that all sizes are *Pareto optimal* at some point, in particular for high recall values, which means that L1/FLOPS regularization at training time still plays a role for SPLADE models. The same experiment with the MRR@10 (not shown here) leads to a slightly different conclusion: *pruning the medium-sized SPLADE is overall better*, but it could be explained by the fact that both the medium and large versions have equivalent MRR@10 values, i.e. the MRR@10 is already saturated for the (M)edium version.




Figure 2: SPLADE models under document-centric pruning. First point on the right of each curve is the unpruned index.

3.2 Static Pruning beyond MSMARCO

In the previous experiments we validated that static pruning is effective for sparse neural retrievers on MSMARCO. However, there are still two questions about its effectiveness: i) does this technique generalize to other datasets? ii) does static pruning of sparse neural retrieval depend on the fact that the training and test sets use the same document collection? In this section, we aim to answer these questions using the TripClick dataset [33].

Static pruning generalizes to TripClick: The experiments below use the best model described in [21] for Tripclick, for which we obtained the models from the authors. As shown in the 3rd column in Table 2, the pruning effect has even a positive effect on the MRR@10: by pruning documents down to 16 tokens, the effectiveness of the pruned index is similar to the baseline, but with a 4.6× speedup. We observe the same phenomena whatever the query category (high, medium, low frequency), but due to a lack of space, we do not provide these figures. Nevertheless, we note that recall at 1000 is more impacted than MRR at 10: all pruned versions show a statistical significant reduction compared to the baseline.

Table 2: Experiments on TripClick HEAD queries. [†] indicates statistical significance difference compared to the baseline.

Max Size	Finetuning			Zero-Shot		
	Latency (×)	MRR@10	R@1k	Latency (×)	MRR@10	R@1k
30 522	5.8 (1.0×)	54.5	88.8	4.2 (1.0×)	32.0	83.6
Prune by document size (D)						
64	3.1 (1.9×)	56.7 [†]	88.3 [†]	2.5 (1.7×)	31.6	82.8 [†]
32	1.9 (3.0×)	55.8	86.8 [†]	1.7 (2.5×)	31.5	80.1 [†]
16	1.2 (4.6×)	54.7	83.4 [†]	1.1 (3.8×)	29.9 [†]	73.4 [†]
8	1.0 (5.8×)	51.3 [†]	76.8 [†]	1.0 (4.1×)	27.7 [†]	58.9 [†]
4	1.0 (5.6×)	42.9 [†]	66.0 [†]	0.8 (5.2×)	22.5 [†]	41.3 [†]

Static pruning works even on zero-shot evaluation. We also conduct a zero-shot evaluation using the SPLADE model previously finetuned with MSMARCO and evaluated on the TripClick dataset. Results reported in the last three columns of Table 2 show again similar results to our previous experiments: we obtain 4× speedup with a small impact on effectiveness ($\leq 8\%$ drop on MRR@10).

3.3 Impact of First Stage Pruning for Reranking

We now look into what happens if we use the pruned first stage as the input for candidate reranking. The goal is not to reduce the time of the overall system (first stage + reranker), but to look into what would happen if you add a reranker to an already existing pruned index. We use a reimplementation of RankT5 [37] to perform this step, over different pruned versions of eff-V-SPLADE-L and present the results for reranking 1000 candidates in Table 3. We note that the results are pretty much similar to what we observed before, 2× *first-stage* speedup (size=32) has almost no loss of effectiveness, while for 3× to 4× *first-stage* speedup (size=8), drops in effectiveness are still small enough ($\leq 8\%$), thus we show that the effectiveness of the reranker does not drop when it is added over a pruned index.

3.4 A More Efficient and Effective SPLADE

Having shown that static pruning can improve the first stage retrievers, we look into how pruning could be combined with other common tools in order to reduce total latency and improve effectiveness. In [summary](#)³, we start from the eff-V-SPLADE L [20], which has 0.388 MRR@10 and a total latency of 73 ms (28 retrieval + 45 inference) and aim at greatly reducing latency while keeping a similar MRR@10: i) we use an ensemble of rerankers to generate

³ Due to space constraints, we refer the reader to the Appendix for a more details.

Table 3: Reranking eff-V-SPLADE-L with static pruning. Latency considers only the first stage step. [†] indicates statistical significant drop compared to the reranked baseline.

Max size	MSMARCO dev		TREC-DL 2019	
	Latency (×)	MRR@10	Latency (×)	nDCG@10
Baseline without reranking				
30 522	28.0 (1.0×)	38.9	28.2 (1.0×)	71.4
Baseline with reranking				
30 522	28.0 (1.0×)	42.8	28.2 (1.0×)	77.1
Prune by document size (D)				
64	13.0 (2.2×)	42.9	14.1 (2.1×)	77.0
32	7.1 (4.0×)	42.7	10.3 (3.2×)	75.8
16	3.7 (7.6×)	42.4 [†]	8.4 (4.7×)	76.4
8	2.0 (13.9×)	41.5 [†]	5.9 (6.4×)	75.6
4	1.0 (27.0×)	38.8 [†]	3.4 (16.6×)	74.1

the distillation scores and change the document encoder to CoCondenser [16], improving MRR@10 from 0.388 to 0.406, with a slight increase in retrieval latency from 28 to 32 ms; ii) recent improvements from model compilation and quantization [2] allow us to reduce average inference latency from 45ms to 7ms; iii) document-centric pruning reduces the retrieval latency from 32 to 12 ms (2.7×) for an effectiveness of 0.396 MRR@10 ($\downarrow 2.5\%$).

In summary, the new model has better effectiveness than the best model presented in [20], while having a total latency of 19ms (12ms retrieval + 7ms inference) which is comparable to one of the most efficient models from [20], which has an effectiveness of 0.376 MRR@10 and a total latency of 13ms (12 ms retrieval + 1ms inference). In other words, we show that pruning can be integrated alongside other improvements to yield state-of-the-art results.

4 CONCLUSION

In this work, we have studied if basic static pruning is effective given the shift to neural sparse retrievers. We verify experimentally that even the simpler techniques still work and that for most methods we are able to achieve a 2x speedup with practically no loss (less than 2% loss of MRR@10) and 4x speedup with negligible loss (less than 8% loss of MRR@10). Furthermore, we have shown that static pruning still works beyond MSMARCO in a zero-shot manner and that efficient pruning does not hurt the reranking performance. Finally, by combining static pruning with other techniques (quantization, training improvements), a more effective and efficient SPLADE model can be obtained, thus advancing the state of the art on the efficiency-effectiveness trade-off.

ACKNOWLEDGMENTS

This work is supported, in part, by the spoke “FutureHPC & BigData” of the ICSC – Centro Nazionale di Ricerca in High-Performance Computing, Big Data and Quantum Computing funded by European Union – NextGenerationEU, and the FoReLab project (Departments of Excellence)

REFERENCES

- [1] Ismail S. Altıngövdü, Rifat Özcan, and Özgür Ulusoy. 2012. Static Index Pruning in Web Search Engines: Combining Term and Document Popularities with Query Views. *ACM TOIS* 30, 1 (2012), 1–28.
- [2] Junjie Bai, Fang Lu, Ke Zhang, et al. 2019. ONNX: Open Neural Network Exchange. <https://github.com/onnx/onnx>.
- [3] Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. 2016. Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268* (2016).
- [4] Roi Blanco and Alvaro Barreiro. 2007. Static Pruning of Terms in Inverted Files. In *Proc. ECIR*. 64–75.
- [5] Roi Blanco and Alvaro Barreiro. 2010. Probabilistic Static Pruning of Inverted Files. *ACM TOIS* 28, 1 (2010), 1–33.
- [6] Stefan Büttcher and Charles L. A. Clarke. 2006. A Document-Centric Approach to Static Index Pruning in Text Retrieval Systems. In *Proc. CIKM*. 182–189.
- [7] David Carmel, Doron Cohen, Ronald Fagin, Eitan Farchi, Michael Herscovici, Yoëlle S. Maarek, and Aya Soffer. 2001. Static Index Pruning for Information Retrieval Systems. In *SIGIR '01: Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 43–50.
- [8] Ruey-Cheng Chen and Chia-Jung Lee. 2013. An Information-Theoretic Account of Static Index Pruning. In *Proc. SIGIR*. 163–172.
- [9] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2021. Overview of the TREC 2020 deep learning track.
- [10] 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* (2020).
- [11] 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* (2020).
- [12] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, Ellen M Voorhees, and Ian Soboroff. 2021. TREC deep learning track: Reusable test collections in the large data regime. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*. 2369–2375.
- [13] Edleno S. de Moura, Célia F. dos Santos, Daniel R. Fernandes, Altigran S. Silva, Pavel Calado, and Mario A. Nascimento. 2005. Improving Web Search Efficiency via a Locality Based Static Pruning Method. In *Proc. WWW*. 235–244.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. NAACL*. 4171–4186.
- [15] Thibault Formal, Benjamin Piwowarski, and Stéphane Clinchant. 2021. SPLADE: Sparse Lexical and Expansion Model for First Stage Ranking. In *Proc. SIGIR*. 2288–2292.
- [16] Luyu Gao and Jamie Callan. 2021. Unsupervised corpus aware language model pre-training for dense passage retrieval. *arXiv preprint arXiv:2108.05540* (2021).
- [17] J. Johnson, M. Douze, and H. Jegou. 2021. Billion-Scale Similarity Search with GPUs. *IEEE Trans. Big Data* 7, 03 (2021), 535–547.
- [18] Norman P. Jouppi et al. 2017. In-Datcenter Performance Analysis of a Tensor Processing Unit. In *Proc. ISCA*. 1–12.
- [19] Omar Khattab and Matei Zaharia. 2020. ColBERT: Efficient and effective passage search via contextualized late interaction over BERT. In *Proc. SIGIR*. 39–48.
- [20] Carlos Lassance and Stéphane Clinchant. 2022. An efficiency study for SPLADE models. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2220–2226.
- [21] Carlos Lassance, Hervé Déjean, and Stéphane Clinchant. 2023. An Experimental Study on Pretraining Transformers from Scratch for IR. <https://doi.org/10.48550/ARXIV.2301.10444>
- [22] Jimmy Lin and Xueguang Ma. 2021. A Few Brief Notes on DeepImpact, COIL, and a Conceptual Framework for Information Retrieval Techniques. *Preprint: arXiv:2106.14807* (2021).
- [23] J. Mackenzie, Z. Dai, L. Gallagher, and J. Callan. 2020. Efficiency Implications of Term Weighting for Passage Retrieval. In *Proc. SIGIR*. 1821–1824.
- [24] Joel Mackenzie, Antonio Mallia, Alistair Moffat, and Matthias Petri. 2022. Accelerating Learned Sparse Indexes Via Term Impact Decomposition. In *Findings of the Association for Computational Linguistics: EMNLP 2022*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 2830–2842. <https://aclanthology.org/2022.findings-emnlp.205>
- [25] Joel Mackenzie, Matthias Petri, and Alistair Moffat. 2021. Anytime ranking on document-ordered indexes. *ACM Transactions on Information Systems (TOIS)* 40, 1 (2021), 1–32.
- [26] Joel Mackenzie, Andrew Trotman, and Jimmy Lin. 2023. Efficient document-at-a-time and score-at-a-time query evaluation for learned sparse representations. *ACM Transactions on Information Systems* 41, 4 (2023), 1–28.
- [27] Antonio Mallia, Omar Khattab, Torsten Suel, and Nicola Tonello. 2021. Learning Passage Impacts for Inverted Indexes. In *Proc. SIGIR*. 1723–1727.
- [28] Antonio Mallia, Joel Mackenzie, Torsten Suel, and Nicola Tonello. 2022. Faster Learned Sparse Retrieval with Guided Traversal. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (Madrid, Spain) (SIGIR '22)*. Association for Computing Machinery, New York, NY, USA, 1901–1905. <https://doi.org/10.1145/3477495.3531774>
- [29] Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage Re-ranking with BERT. *arXiv:Preprint: arXiv:1901.04085*
- [30] Rodrigo Nogueira and Jimmy Lin. 2019. From doc2query to docTTTTTquery. (2019).
- [31] Alexandros Ntoulas and Junghoo Cho. 2007. Pruning Policies for Two-Tiered Inverted Index with Correctness Guarantee. In *Proc. SIGIR*. 191–198.
- [32] Biswajit Paria, Chih-Kuan Yeh, Ian E.H. Yen, Ning Xu, Pradeep Ravikumar, and Barnabás Póczos. 2020. Minimizing FLOPs to Learn Efficient Sparse Representations. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=SygpC6Ntvr>
- [33] Navid Rekabsaz, Oleg Lesota, Markus Schedl, Jon Brassey, and Carsten Eickhoff. 2021. TripClick: The Log Files of a Large Health Web Search Engine. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2507–2513. <https://doi.org/10.1145/3404835.3463242>
- [34] Sree Lekha Thota and Ben Carterette. 2011. Within-Document Term-Based Index Pruning with Statistical Hypothesis Testing. In *Proc. ECIR*. 543–554.
- [35] Nicola Tonello and Craig Macdonald. 2021. Query Embedding Pruning for Dense Retrieval. In *Proc. CIKM*. 3453–3457.
- [36] Nicola Tonello, Craig Macdonald, and Iadh Ounis. 2018. Efficient Query Processing for Scalable Web Search. *Found. Trends in Inf. Retr.* 12, 4–5 (2018), 319–492.
- [37] Honglei Zhuang, Zhen Qin, Rolf Jagerman, Kai Hui, Ji Ma, Jing Lu, Jianmo Ni, Xuanhui Wang, and Michael Bendersky. 2022. RankT5: Fine-tuning T5 for text ranking with ranking losses. *arXiv preprint arXiv:2210.10634* (2022).
- [38] Shengyao Zhuang and Guido Zuccon. 2021. Fast passage re-ranking with contextualized exact term matching and efficient passage expansion. *Preprint: arXiv:2108.08513* (2021).