

Exploring the Representation Power of SPLADE Models

Joel Mackenzie*
The University of Queensland
Brisbane, Australia
joel.mackenzie@uq.edu.au

Shengyao Zhuang*
The University of Queensland
Brisbane, Australia
s.zhuang@uq.edu.au

Guido Zuccon
The University of Queensland
Brisbane, Australia
g.zuccon@uq.edu.au

ABSTRACT

The SPLADE (SParse Lexical AnD Expansion) model is a highly effective approach to learned sparse retrieval, where documents are represented by term impact scores derived from large language models. During training, SPLADE applies regularization to ensure postings lists are kept sparse — with the aim of mimicking the properties of natural term distributions — allowing efficient and effective lexical matching and ranking. However, we hypothesize that SPLADE may encode additional signals into common postings lists to further improve effectiveness. To explore this idea, we perform a number of empirical analyses where we re-train SPLADE with different, controlled vocabularies and measure how effective it is at ranking passages. Our findings suggest that SPLADE can effectively encode useful ranking signals in documents even when the vocabulary is constrained to terms that are not traditionally useful for ranking, such as stopwords or even random words.

CCS CONCEPTS

• **Information systems** → **Query representation; Language models.**

KEYWORDS

Learned sparse retrieval, SPLADE, Term expansion

ACM Reference Format:

Joel Mackenzie, Shengyao Zhuang, and Guido Zuccon. 2023. Exploring the Representation Power of SPLADE Models. In *Proceedings of the 2023 ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR '23)*, July 23, 2023, Taipei, Taiwan. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3578337.3605129>

1 INTRODUCTION AND PRELIMINARIES

Learned sparse retrieval is a paradigm that exploits decades of work on efficient inverted index-based representations [21] and modern deep-learning based relevance modelling by replacing simple statistical payloads (usually term frequencies) with *learned impacts* [12]. A number of learned sparse retrieval models have been proposed [4, 6, 13, 15, 17, 22]; their key differences lie in terms of how or whether document expansion is applied [15, 20, 23], the way the model is trained to assign impacts, and the use of query term expansion and weighting. The SPLADE (SParse Lexical AnD Expansion)

*Equal contribution

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

ICTIR '23, July 23, 2023, Taipei, Taiwan

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 979-8-4007-0073-6/23/07...\$15.00
<https://doi.org/10.1145/3578337.3605129>

```
(##rogen', 251) ('receptor', 242) ('and', 225) ('receptors', 189)
('hormone', 179) ('definition', 162) ('meaning', 99) ('genus', 89)
('is', 70) (',', 68) ('define', 59) ('the', 56) ('drug', 53) ('for', 46)
('ring', 38) ('gene', 37) ('are', 32) ('god', 25) ('what', 18) ('##rus', 15)
('purpose', 12) ('defined', 10) ('doing', 8) ('a', 4) ('goal', 4)
```

Figure 1: An example of the SPLADE encoding for the query androgen receptor define. Each query term is represented as a ('token', weight) tuple, with blue terms representing the original input query terms, orange terms representing alternate inflections on those original terms, and pink terms representing entirely new terms. Figure reproduced with permission from Mackenzie et al. [16].

Table 1: Ten most commonly expanded tokens by SPLADEv2 on the 6,980 MS MARCO dev queries. These commonly expanded tokens have very long postings lists, and are typically not informative.

Term	Count	%	List length
him	2419	34.7	2,907,631
ring	1995	28.6	2,665,265
it	1941	27.8	2,737,546
for	1921	27.5	3,116,632
a	1838	26.3	2,873,128
are	1732	24.8	2,782,156
god	1704	24.4	1,293,359
.	1659	23.8	3,381,977
cause	1657	23.7	1,869,128
,	1606	23.0	2,553,418

family currently represents the most effective approach to learned sparse retrieval, with variations providing different efficiency and effectiveness trade-offs [4–6, 10]. Nguyen et al. [18] provide a unified overview and comparison of various learned sparse retrieval methods.

Recently, Mackenzie et al. [16] explored the efficiency properties of various learned sparse retrieval models, showing that they are often less efficient than traditional rankers due to so-called “wacky weight” distributions. However, it was not only the distribution of the impact weights that was abnormal; SPLADE was shown to expand terms that seem to have no relevance to the query, including the use of stopwords and punctuation. Figure 1 provides an example of this behavior, and Table 1 reports the top ten most commonly expanded tokens across the MS MARCO dev query set. While it is difficult to understand the reasoning behind these somewhat strange term expansions, we hypothesize that SPLADE may be using tokens that are typically not informative to encode additional signals that can be used during ranking. As such, the goal of this preliminary study is to better understand why SPLADE expands

such terms, and to determine whether SPLADE has the ability to encode signals into typically uninformative lists to improve ranking.

To gain insight into how “meaningless” tokens affect the effectiveness of the SPLADE model, we made changes to its MLM head. This allowed the model to expand and assign weights to specific sets of tokens during both training and inference. For instance, we experimented with restricting SPLADE’s expansion to a limited set of 150 stopwords (e.g., “is”, “the”, and “and”), which are considered to have little semantic meaning in the context of the content. The surprising outcome was that despite this restriction, SPLADE was still able to perform at the same level of effectiveness as BM25 on MS MARCO dev queries. This result suggests that SPLADE has the ability to encode semantic information effectively, regardless of the tokens used to represent documents. To further understand this phenomenon and explore which types of tokens have the greatest impact on SPLADE’s encoding, we tested other settings, such as limiting expansion to low-frequency tokens, and even restricting the vocabulary to “latent tokens” that do not represent *real* terms. We believe our findings have the potential to broaden the community’s understanding of the SPLADE model, and to motivate efforts to improve the explainability of learned sparse retrieval systems.

2 RELATED WORK

2.1 SPLADE models

The SPLADE family of models [4–6, 10] are BERT-based sparse retrieval models. They have been developed to project every term in queries and documents to a vocabulary-sized weight vector, where each dimension represents the weight of a token in the BERT’s vocabulary. This innovative approach estimates the weights of these vectors using the logits of masked language models, and then combines them (e.g., through sum or max pooling) to obtain a single representation of the query and document. This representation can be considered an expansion of the query or document, as it includes terms not originally present. To maintain efficiency, a sparsity regularization loss is applied during training to obtain a sparse representation that can be used more effectively with the inverted index. Compared to other learned sparse methods, such as docT5query [20], DeepImpact [17], uniCOIL [13] and TILDEv2 [22], SPLADE models perform the “expansion then assigning weights” process in an end-to-end manner. Hence, this model offers a more efficient and effective solution to the vocabulary mismatch problem for learned sparse retrieval.

2.2 SPLADEv2

This paper delves into the SPLADEv2 model [4]. This model can predict term importance over BERT WordPiece vocabulary ($|V_{BERT}| = 30522$) for any given query and passage. SPLADEv2 relies on BERT’s contextual embeddings (h_1, h_2, \dots, h_n) for a given sequence of text tokens (query tokens or passage tokens). Using the MLM head of its trained BERT, it maps each embedding into a weight vector comprising of the entire vocabulary size as follows:

$$\vec{w}_i = \text{MLM}(h_i), \quad (1)$$

In this equation, $\vec{w}_i \in \mathcal{R}^V$, where each element w_{ij} within this vector corresponds to the importance score of the j th token in the BERT vocabulary. To obtain the final importance score of the j th

token, a log-saturation effect is applied to the max pooling of all j th elements in the weight vectors of the sequence of n tokens. Negative scores are masked out by the ReLU function, resulting in the following equation:

$$w_j = \max_{i \in n} \log(1 + \text{ReLU}(w_{ij})). \quad (2)$$

With the importance score prediction scheme in place, SPLADEv2 is capable of creating a single vocabulary size vector \vec{q} or \vec{d} for the given query or document. Each element in these vectors represents the importance of the corresponding token to the original input text. Following this, the model is trained to learn sparse lexical matching between relevant query-document pairs, using contrastive loss and FLOPS losses. For a more detailed explanation of these losses, we refer the reader to the original paper [4]. However, it is crucial to note that, after training, any element in \vec{q} and \vec{d} has the potential to be assigned a high importance score, regardless of whether the token occurs in the original input text or not.

3 EXPERIMENTS

We now describe our experimental setup, baselines, and the controlled vocabulary mechanisms we explore. Our goal is to examine the effects of different expanded vocabularies on SPLADEv2. To demonstrate the impact, we compare the outcomes with BM25, a dense retriever trained under similar conditions, and the original SPLADEv2.

3.1 Experimental Setup

We experiment on the MS MARCO-v1 passage corpus consisting of around 8.8 million passages [19]. We evaluate our experiments with both the MS MARCO dev set (containing 6,980 queries and around 1.1 relevant documents per query) as well as the 2019 and 2020 TREC deep learning track passage ranking queries and judgments (43 and 54 queries for DL2019 and DL2020, respectively, and both with much deeper judgements) [1, 2].

3.2 Baseline Systems

To situate our findings, we report a few typical baselines as follows.

BM25. We use the Pyserini [14] toolkit to run BM25 with the tuned parameters $k_1 = 0.82$ and $b = 0.68$ as recommended by Dai and Callan [3].¹

Dense Retriever. We train a dense retriever with MS MARCO training dataset by following the example provided in Tevatron DR training toolkit [9].² Specifically, we employed CoCondenser [8] as our backbone encoder model. For each query in the MS MARCO training set, we took the top 200 passages retrieved by BM25 and randomly selected 7 hard negative passages, as well as one positive passage from the qrels. With a batch size of 8, we applied in-batch negatives to each training sample in the batch, which resulted in 63 negatives per training sample. To optimize our training, we set the learning rate to 5e-6 and the number of epochs to 3.

¹<https://github.com/castorini/pyserini/blob/master/docs/experiments-msmarco-passage.md>

²https://github.com/texttron/tevatron/blob/main/examples/example_msmarco.md

Table 2: Overall effectiveness of the models across all collections. The best results are highlighted in boldface. Superscripts denote significant differences in paired, two-sided Student’s t-test with $p \leq 0.01$.

#	Model	MS MARCO dev		DL2019		DL2020	
		RR@10	Recall@1000	NDCG@10	Recall@1000	NDCG@10	Recall@1000
A	BM25	18.7	85.7	49.7	74.5 ^{efk}	48.8	80.3 ^{efhk}
B	DR	35.9 ^{defghijklm}	97.8 ^{acdefghijklm}	65.8 ^{efghikm}	78.2 ^{efghikm}	66.1 ^{efghikm}	82.4 ^{efghikm}
C	SPLADEv2	35.5 ^{defghikm}	97.4 ^{defghikm}	66.9 ^{efghikm}	81.1 ^{efghikm}	62.3 ^{efhk}	81.9 ^{efghikm}
D	no-stop	35.7 ^{efghijklm}	97.5 ^{efghikm}	69.0 ^{efghikm}	82.8 ^{efghikm}	64.3 ^{efhk}	83.8 ^{efghikm}
E	stop-150	20.8 ^a	84.4	42.7	56.6	43.9	62.1
F	random-150	23.4 ^{ae}	87.0 ^e	51.6	60.0	50.3	68.8 ^e
G	random-768	28.4 ^{efhk}	92.9 ^{efhk}	56.0 ^e	65.1 ^{ek}	56.7 ^{eh}	72.5 ^{eh}
H	lowfreq-150	24.0 ^{ae}	87.8 ^{ae}	47.6	60.7	47.9	66.7
I	lowfreq-768	28.7 ^{efhk}	93.2 ^{efhk}	57.3 ^{eh}	69.3 ^{efhk}	60.2 ^{efhk}	73.5 ^{efhk}
J	added-latent-150	35.0 ^{efghikm}	97.2 ^{efghikm}	68.2 ^{efghikm}	79.7 ^{efghikm}	63.0 ^{efhk}	82.2 ^{efghikm}
K	latent-150	24.6 ^{af}	86.6 ^e	53.1 ^e	57.4	50.3	67.8 ^e
L	added-latent-768	35.3 ^{efghikm}	97.5 ^{efghijklm}	68.1 ^{efghikm}	81.8 ^{efghikm}	64.7 ^{efhk}	83.4 ^{efghikm}
M	latent-768	29.1 ^{efhk}	92.5 ^{efhk}	51.3	66.1 ^{efhk}	60.1 ^{efhk}	73.2 ^{efhk}

SPLADEv2. Our final baseline is the “default” SPLADEv2 system. To train SPLADEv2, we follow the example in the Tevatron toolkit from the original SPLADE authors.³ To ensure a fair comparison with our dense retriever baseline, we utilized the same backbone model and training hyperparameters, and note that all systems were trained on the *title augmented* passage collection [11].

3.3 SPLADEv2 with Controlled Vocabularies

To determine how important the original representation of each document is for the effectiveness of SPLADEv2, we re-train SPLADEv2 using different vocabularies. During training, we limit the attention to only tokens appearing in the provided vocabulary; as such, every document and query is represented by just a subset of these vocabulary tokens. We tested the following vocabularies.

No Stopword Tokens. Our first model is a simple modification to the SPLADEv2 baseline discussed above. We apply the NLTK⁴ English stopword list to the BERT vocabulary to ensure stopwords are not seen by the model during training. We refer to this approach as no-stop.

Stopword Tokens. Our second approach applies *only* the NLTK English stopword list as the full vocabulary, resulting in a total of $|V| = 150$ unique tokens that can be used to represent each document and/or query. These tokens are typically considered to be uninformative for ranking and retrieval on English corpora due to their commonality. We refer to this setting as stop-150.

Random Tokens. We also randomly sample tokens from the original BERT vocabulary ($|V_{BERT}| = 30,522$ tokens). We try both $|V| = 150$ (random-150) and $|V| = 768$ (random-768) terms corresponding to the size of the *stopwords-tokens-only* and typical dense retrieval dimensionality, respectively.

Low Frequency Tokens. Another approach is to simply take the *lowest frequency* tokens (the rarest tokens). In particular, we tokenize the original corpus into BERT tokens and then use the least frequent tokens as the vocabulary, again experimenting with $|V| = 150$ and $|V| = 768$ (called lowfreq-150 and lowfreq-768, respectively).

Latent Tokens. In this context, we introduce 150 and 768 new tokens that we call latent tokens. These tokens do not exist in the original BERT vocabulary and have randomly initialized embeddings. We conduct experiments in four different settings to explore the impact of these latent tokens. The first setting is called *added latent tokens* (added-latent-*), where we add our 150 or 768 latent tokens to the original BERT vocabulary, resulting in a total of $|V| = |V_{BERT}| + 150$ and $|V| = |V_{BERT}| + 768$ tokens. The second setting is called *latent tokens only* (latent-*), where we use our sets of latent tokens as the full vocabulary, resulting in $|V| = 150$ and $|V| = 768$.

4 RESULTS

Table 2 presents the effectiveness of the baselines and alternative vocabulary approaches across the MS MARCO dev, DL2019, and DL2020 topics. We break these results down in the following paragraphs, all in reference to Table 2.

Are stopwords useful? Given that SPLADEv2 was shown to expand seemingly “unhelpful” tokens in many queries, we hypothesized that SPLADE may be using these tokens to encode additional ranking signals. We now try to determine whether this is indeed the case, using stopwords as our proxy for such tokens. First, let us examine the stop-only SPLADEv2 (row E) to the baseline systems. Surprisingly, even if we only allow SPLADEv2 to represent the entire index with 150 stopwords, it manages to outperform the BM25 baseline on the dev set with RR@10. Although this finding does

³<https://github.com/texttron/tevatron/tree/main/examples/splade>

⁴<https://www.nltk.org/>

not hold across the other metrics or collections, it is still surprisingly effective. This motivated us to further explore how SPLADEv2 leverages stopwords for ranking.

Comparing our SPLADEv2 baseline (B) to the same system *without* the inclusion of stopwords during training (D), we observe some (but not significant) improvements when the stopwords were *removed*, which contradicts our hypothesis that these tokens are being used by the original model in a meaningful way.

To try to better understand this difference, we also experimented with removing the top 100 most commonly expanded terms on the original SPLADEv2 model during *query time* (that is, without re-training the baseline SPLADEv2 model). Similarly, we observed only a very minor drop in effectiveness (from an RR@10 of 35.5 to 35.2 on the dev queries). Overall, this suggests that our original hypothesis does not hold. So how do we explain the surprising effectiveness of the stop-150 system, which represents all documents and queries through just 150 stopwords?

Does the vocabulary actually matter? Next, we try to understand why training SPLADEv2 on just 150 stopword tokens yields such a surprisingly effective model. Our first idea was to consider alternative vocabularies (with 150 tokens) – systems F (random tokens), H (low frequency tokens), and K (latent tokens) – to see if the findings hold. Here, we observe even better effectiveness than stop-150, with all of the alternatives outperforming stop-150 (and rivalling BM25) across most metrics and collections. This suggests that SPLADEv2 is capable of modelling useful ranking information in *arbitrary vocabularies*, irrespective of the terms within them.

Increasing the vocabulary size. Next, we re-examine the controlled vocabulary SPLADEv2 systems by increasing the vocabulary size from 150 to 768 tokens. Interestingly, we see large (and often significant) improvements (from F to G, H to I, and K to M) just by increasing the vocabulary size. This is intuitive, because larger vocabularies can provide SPLADEv2 with more “space” to accurately model and represent term interactions. Based on the notion that the vocabulary terms themselves do not seem to make a huge difference to the quality of retrieval, but that the size of the vocabulary seems important, we began to hypothesize whether SPLADEv2 can be thought of as a type of *dense retriever*, in the sense that the lexical matching only serves to activate some learned representation for relevance. However, comparing our results to the baseline dense retriever (also using 768 dimensions) demonstrates that these systems (G, I, and M) are significantly less effective than the dense retriever, suggesting that training SPLADEv2 with a constrained representation may have weaker semantic encoding abilities. Despite this finding, it would be premature to conclude that SPLADEv2 cannot function as a dense retriever, especially since additional experimentation is necessary to verify or refute this hypothesis.

Can latent tokens help SPLADEv2? Since even arbitrary lexicons were able to represent documents for ranking tasks, we now test whether *adding* tokens to the original SPLADEv2 model can improve performance further. Our intuition is that SPLADEv2 may be able to take advantage of the *latent tokens* to form a type of hybrid retriever. During training, we allow SPLADEv2 to use the latent tokens thus increasing the original vocabulary by either 150 or 768. Rows J and L report the outcome. Interestingly, we find

that adding latent tokens can increase the effectiveness of the original SPLADEv2 baseline, though not significantly, suggesting that SPLADEv2 may be able to employ these tokens to its advantage. However, further experimentation is required to understand this phenomenon.

5 CONCLUSION AND FUTURE WORK

In conclusion, our study demonstrated that SPLADEv2 is a robust model that can effectively represent documents by a small set of meaningless stopword tokens and still deliver impressive effectiveness. Moreover, our experiments with randomly sampled and low frequency sets of tokens revealed that SPLADEv2 can outperform BM25, even when the expanded tokens are completely irrelevant to the original text. This surprising finding suggests that SPLADEv2 can effectively encode semantic information into any tokens it is given, effectively acting like a dense retriever. Our attempt to make SPLADEv2 act like a hybrid model by adding latent tokens to the original vocabulary showed some improvement in effectiveness, albeit not statistically significant. Our preliminary study sheds light on the capabilities of the SPLADEv2 model and can help the community gain a deeper understanding of its potential applications. For future work, we suggest to investigate the relationship between vocabulary size (and subset) and SPLADE’s representation capabilities. Additionally, exploring the potential use of latent tokens could potentially enhance the semantic encoding power of the SPLADE model. It would also be interesting to examine other models to understand if these findings generalize [7]. These avenues of inquiry will help the community to fully comprehend the capabilities of such models and maximize its potential for future applications.

We have made our code publicly available at <https://github.com/ielab/understanding-splade>, allowing others to easily reproduce the results presented in this paper.

REFERENCES

- [1] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2020. Overview of the TREC 2020 Deep Learning Track. In *Proceedings of the Twenty-Ninth Text REtrieval Conference, TREC 2020, Virtual Event [Gaithersburg, Maryland, USA], November 16-20, 2020 (NIST Special Publication)*, Ellen M. Voorhees and Angela Ellis (Eds.), Vol. 1266. National Institute of Standards and Technology (NIST). <https://trec.nist.gov/pubs/trec29/papers/OVERVIEW.DL.pdf>
- [2] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M. Voorhees. 2020. Overview of the TREC 2019 deep learning track. *CoRR abs/2003.07820* (2020). arXiv:2003.07820 <https://arxiv.org/abs/2003.07820>
- [3] Zhuyun Dai and Jamie Callan. 2020. Context-Aware Term Weighting For First Stage Passage Retrieval. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, China) (SIGIR '20). Association for Computing Machinery, New York, NY, USA, 1533–1536. <https://doi.org/10.1145/3397271.3401204>
- [4] Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. 2021. SPLADE v2: Sparse Lexical and Expansion Model for Information Retrieval. *CoRR abs/2109.10086* (2021). arXiv:2109.10086 <https://arxiv.org/abs/2109.10086>
- [5] Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. 2022. From Distillation to Hard Negative Sampling: Making Sparse Neural IR Models More Effective. 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, 2353–2359. <https://doi.org/10.1145/3477495.3531857>
- [6] Thibault Formal, Benjamin Piwowarski, and Stéphane Clinchant. 2021. SPLADE: Sparse Lexical and Expansion Model for First Stage Ranking. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, Canada) (SIGIR '21). Association for Computing Machinery, New York, NY, USA, 2288–2292. <https://doi.org/10.1145/3404835.3463098>
- [7] Thibault Formal, Benjamin Piwowarski, and Stéphane Clinchant. 2021. A White Box Analysis of ColBERT. In *Advances in Information Retrieval*, Djoerd Hiemstra,

- Marie-Francine Moens, Josiane Mothe, Raffaele Perego, Martin Potthast, and Fabrizio Sebastiani (Eds.). Springer International Publishing, Cham, 257–263.
- [8] Luyu Gao and Jamie Callan. 2022. Unsupervised Corpus Aware Language Model Pre-training for Dense Passage Retrieval. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Dublin, Ireland, 2843–2853. <https://doi.org/10.18653/v1/2022.acl-long.203>
- [9] Luyu Gao, Xueguang Ma, Jimmy Lin, and Jamie Callan. 2022. Tevatron: An Efficient and Flexible Toolkit for Dense Retrieval. *CoRR* abs/2203.05765 (2022). <https://doi.org/10.48550/arXiv.2203.05765> arXiv:2203.05765
- [10] 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* (Madrid, Spain) (SIGIR '22). Association for Computing Machinery, New York, NY, USA, 2220–2226. <https://doi.org/10.1145/3477495.3531833>
- [11] Carlos Lassance and Stéphane Clinchant. 2023. The tale of two MS MARCO - and their unfair comparisons. *CoRR* abs/2304.12904 (2023). <https://doi.org/10.48550/arXiv.2304.12904> arXiv:2304.12904
- [12] Jimmy Lin. 2022. A Proposed Conceptual Framework for a Representational Approach to Information Retrieval. *SIGIR Forum* 55, 2, Article 4 (mar 2022), 29 pages. <https://doi.org/10.1145/3527546.3527552>
- [13] Jimmy Lin and Xueguang Ma. 2021. A Few Brief Notes on DeepImpact, COIL, and a Conceptual Framework for Information Retrieval Techniques. *CoRR* abs/2106.14807 (2021). arXiv:2106.14807 <https://arxiv.org/abs/2106.14807>
- [14] Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. 2021. Pyserini: A Python Toolkit for Reproducible Information Retrieval Research with Sparse and Dense Representations. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, Canada) (SIGIR '21). Association for Computing Machinery, New York, NY, USA, 2356–2362. <https://doi.org/10.1145/3404835.3463238>
- [15] Xueguang Ma, Ronak Pradeep, Rodrigo Nogueira, and Jimmy Lin. 2022. Document Expansion Baselines and Learned Sparse Lexical Representations for MS MARCO V1 and V2. 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, 3187–3197. <https://doi.org/10.1145/3477495.3531749>
- [16] 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 Trans. Inf. Syst.* 41, 4, Article 96 (mar 2023), 28 pages. <https://doi.org/10.1145/3576922>
- [17] Antonio Mallia, Omar Khattab, Torsten Suel, and Nicola Tonello. 2021. Learning Passage Impacts for Inverted Indexes. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, Canada) (SIGIR '21). Association for Computing Machinery, New York, NY, USA, 1723–1727. <https://doi.org/10.1145/3404835.3463030>
- [18] Thong Nguyen, Sean MacAvaney, and Andrew Yates. 2023. A Unified Framework For Learned Sparse Retrieval. In *Advances in Information Retrieval: 45th European Conference on Information Retrieval, ECIR 2023, Dublin, Ireland, April 2–6, 2023, Proceedings, Part III* (Dublin, Ireland). Springer-Verlag, Berlin, Heidelberg, 101–116. https://doi.org/10.1007/978-3-031-28241-6_7
- [19] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A Human Generated MACHINE Reading COMprehension Dataset. In *Proceedings of the Workshop on Cognitive Computation: Integrating neural and symbolic approaches 2016 co-located with the 30th Annual Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain, December 9, 2016 (CEUR Workshop Proceedings)*, Tarek Richard Besold, Antoine Bordes, Artur S. d'Avila Garcez, and Greg Wayne (Eds.), Vol. 1773. CEUR-WS.org. https://ceur-ws.org/Vol-1773/CoCoNIPS_2016_paper9.pdf
- [20] Rodrigo Nogueira and Jimmy Lin. 2019. From doc2query to docTTTTTquery. https://cs.uwaterloo.ca/~jimmylin/publications/Nogueira_Lin_2019_docTTTTTquery-latest.pdf Unpublished report, David R. Cheriton School of Computer Science, University of Waterloo, Canada.
- [21] Nicola Tonello, Craig Macdonald, and Iadh Ounis. 2018. Efficient Query Processing for Scalable Web Search. *Found. Trends Inf. Retr.* 12, 4–5 (dec 2018), 319–500. <https://doi.org/10.1561/15000000057>
- [22] Shengyao Zhuang and Guido Zuccon. 2021. Fast Passage Re-ranking with Contextualized Exact Term Matching and Efficient Passage Expansion. *CoRR* abs/2108.08513 (2021). arXiv:2108.08513 <https://arxiv.org/abs/2108.08513>
- [23] Shengyao Zhuang and Guido Zuccon. 2021. TILDE: Term Independent Likelihood MoDEL for Passage Re-Ranking. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, Canada) (SIGIR '21). Association for Computing Machinery, New York, NY, USA, 1483–1492. <https://doi.org/10.1145/3404835.3462922>