

SLIM: Sparsified Late Interaction for Multi-Vector Retrieval with Inverted Indexes

Minghan Li
University of Waterloo
Waterloo, Canada
m692li@uwaterloo.ca

Sheng-Chieh Lin
University of Waterloo
Waterloo, Canada
s269lin@uwaterloo.ca

Xueguang Ma
University of Waterloo
Waterloo, Canada
x93ma@uwaterloo.ca

Jimmy Lin
University of Waterloo
Waterloo, Canada
jimmylin@uwaterloo.ca

ABSTRACT

This paper introduces Sparsified Late Interaction for Multi-vector (SLIM) retrieval with inverted indexes. Multi-vector retrieval methods have demonstrated their effectiveness on various retrieval datasets, and among them, ColBERT is the most established method based on the late interaction of contextualized token embeddings of pre-trained language models. However, efficient ColBERT implementations require complex engineering and cannot take advantage of off-the-shelf search libraries, impeding their practical use. To address this issue, SLIM first maps each contextualized token vector to a sparse, high-dimensional lexical space before performing late interaction between these sparse token embeddings. We then introduce an efficient two-stage retrieval architecture that includes inverted index retrieval followed by a score refinement module to approximate the sparsified late interaction, which is fully compatible with off-the-shelf lexical search libraries such as Lucene. SLIM achieves competitive accuracy on MS MARCO Passages and BEIR compared to ColBERT while being much smaller and faster on CPUs. To our knowledge, we are the first to explore using sparse token representations for multi-vector retrieval. Source code and data are integrated into the Pyserini IR toolkit.

CCS CONCEPTS

• Information systems → Retrieval models and ranking.

KEYWORDS

Neural IR, Late Interaction, Inverted Indexes, Sparse Retrieval

ACM Reference Format:

Minghan Li, Sheng-Chieh Lin, Xueguang Ma, and Jimmy Lin. 2023. SLIM: Sparsified Late Interaction for Multi-Vector Retrieval with Inverted Indexes. 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, 6 pages. <https://doi.org/10.1145/3539618.3591977>

1 INTRODUCTION

Pairwise token interaction [12, 19, 45] has been widely used in information retrieval tasks [31]. Interaction-based methods enable

deep coupling between queries and documents, which often outperform representation-based methods [2, 34] when strong text encoders are absent. However, with the rise of pre-trained transformer models [4, 27], representation-based methods such as DPR [17] and SPLADE [6] gain more popularity as the pre-trained representations capture rich semantics of the input texts. Moreover, these methods can leverage established search libraries such as FAISS [16] and Pyserini [22] for efficient retrieval. To combine the best of both worlds, models such as ColBERT [18] and COIL [9] that leverage late interaction are proposed, where the token interaction between queries and documents only happens at the last layer of contextualized embeddings. Their effectiveness and robustness are demonstrated on various retrieval and question-answering benchmarks. However, these models require token-level retrieval and aggregation, which results in large indexes and high retrieval latency. Therefore, different optimization schemes are proposed to improve efficiency in both time and space [13, 37], making multi-vector retrieval difficult to fit into off-the-shelf search libraries.

In this paper, we propose an efficient and compact approach called Sparsified Late Interaction for Multi-vector retrieval with inverted indexes, or SLIM. SLIM is as effective as the state-of-the-art multi-vector model ColBERT on MS MARCO Passages [32] and BEIR [40] without any distillation or hard negative mining while being more efficient. More importantly, SLIM is fully compatible with inverted indexes in existing search toolkits such as Pyserini, with only a few extra lines of code in the pre-processing and post-processing steps. In contrast, methods such as ColBERT and COIL require custom implementations and optimizations [13, 37] in order to be practical to use. This presents a disadvantage because some researchers and practitioners may prefer neural retrievers that are compatible with existing infrastructure based on inverted indexes so that the models can be easily deployed in production.

In order to leverage standard inverted index search, SLIM projects each contextualized token embedding to a high-dimensional, sparse lexical space [6] before performing the late interaction operation. Ideally, the trained sparse representations of documents can be compressed and indexed into an inverted file system for search. Our method is the first to explore using sparse representations for multi-vector retrieval with inverted indexes to our knowledge. However, unlike previous supervised sparse retrieval methods such as uniCOIL [21] or SPLADE [6], which compute sequence-level representations for each query/document, deploying token-level sparse vectors into inverted indexes is problematic. As shown in Figure 1a, the high-level idea is to convert each sparse token embedding into a sub-query/document and perform token-level inverted list retrieval before aggregation (e.g., scatter sum and max). However, this is incredibly slow in practice as latency increases rapidly when the posting lists are too long.

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/10.1145/3539618.3591977>

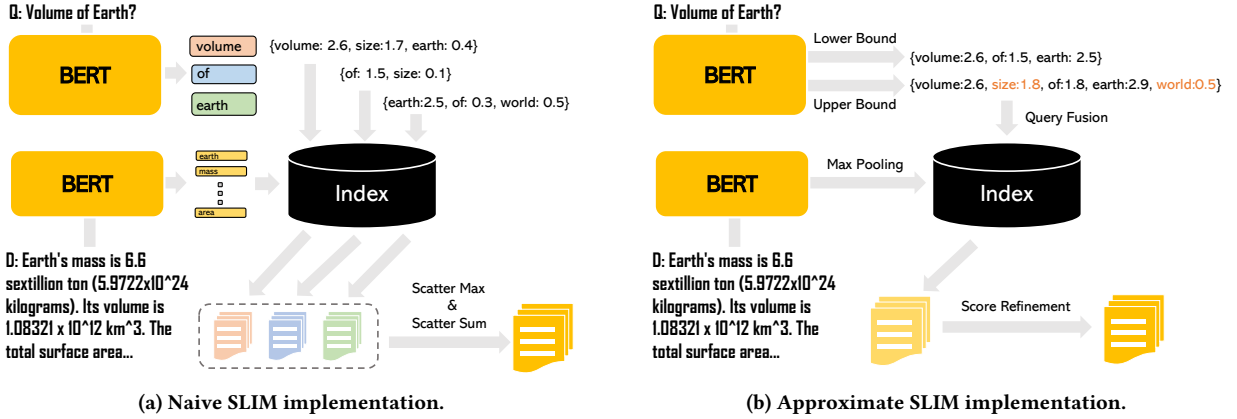


Figure 1: (a) The naive SLIM implementation takes the expansion of each token as sub-queries/documents. The token-level rankings are merged using scatter operations. (b) The approximate SLIM implementation first fuses the lower- and upper-bounds of scores for retrieval using Equations (8) and (9); the candidate list is then refined according to Equation (4).

Therefore, instead of using the naive approach, a more economical way is to approximate the sparsified late interaction with a two-stage system. We first calculate the upper and lower bounds of the sparsified late-interaction scores by swapping the order of the operators in late interaction (see Section 2). In this way, the token-level sparse vectors are pooled before indexing to obtain sequence-level representations for the inverted index, yielding fast retrieval speed. After retrieval, a lightweight score refinement step is applied, where we use the sparse token vectors stored using Scipy [41] to rerank the candidate lists. As shown in Figure 1b, this two-stage design allows us to optimize the latency of the retrieval stage without worrying about accuracy loss as the score refinement will compensate for errors from the first stage. Experiments on MS MARCO Passages show that SLIM is able to achieve a similar ranking accuracy compared to ColBERT-v2, while using 40% less storage and achieving an 83% decrease in latency. To sum up, our contributions in this paper are three-fold:

- We are the first to use sparse representations for multi-vector retrieval with standard inverted indexes.
- We provide a two-stage implementation of SLIM by approximating late interaction followed by a score refinement step.
- SLIM is fully compatible with off-the-shelf search toolkits such as Pyserini.

2 METHODOLOGY

ColBERT [18] proposes late interaction between the tokens in a query $q = \{q_1, q_2, \dots, q_N\}$ and a document $d = \{d_1, d_2, \dots, d_M\}$:

$$s(q, d) = \sum_{i=1}^N \max_{j=1}^M v_{q_i}^T v_{d_j}, \quad (1)$$

where v_{q_i} and v_{d_j} denote the last-layer contextualized token embeddings of BERT. This operation exhaustively compares each query token to all document tokens. The latency and storage of ColBERT are bloated as many tokens do not contribute to either query or document semantics, and thus complex engineering optimizations are needed to make the model practical [37, 38].

Unlike ColBERT, which only uses contextualized token embeddings for computing similarity, SPLADE [5, 6] further utilizes the pre-trained Mask Language Modeling (MLM) layer to project each c -dimensional token embedding to a high-dimensional, lexical space V . Each dimension corresponds to a token and has a non-negative value due to the following activation:

$$\phi_{d_j} = \log(1 + \text{ReLU}(W^T v_{d_j} + b)), \quad (2)$$

where $\phi_{d_j} \in \mathbb{R}^{|V|}$, v_{d_j} is the token embedding of the j th token of document d ; W and b are the weights and biases of the MLM layer. To compute the final similarity score, SPLADE pools all the token embeddings into a sequence-level representation and uses the dot product between a query and a document as the similarity:

$$s(q, d) = (\max_{i=1}^N \phi_{q_i})^T (\max_{j=1}^M \phi_{d_j}). \quad (3)$$

Here, max pooling is an element-wise operation and is feasible for sparse representations as the dimensions of all the token vectors are aligned (lexical space) and non-negative. These two properties play a vital role in the success of SPLADE and are also important for making SLIM efficient, as we shall see later.

Similar to ColBERT's late interaction, sparsified late interaction also takes advantages of the contextualized embeddings of BERT's last layer. But different from ColBERT, we first apply the sparse activation in Equation (2) before calculating the similarity in Equation (1):

$$s(q, d) = \sum_{i=1}^N \max_{j=1}^M \phi_{q_i}^T \phi_{d_j}, \quad (4)$$

where ϕ_{d_j} and ϕ_{q_i} are the representations in Equation (2). However, as shown in Figure 1a, to keep the token-level interaction, we must convert the sparse representation of each token in a query to a sub-query and retrieve the sub-documents from the index. Moreover, the token-level rankings need to be merged to yield the final ranked list using scatter operations, which further increases latency.

Due to the impracticality of the naive implementation, a natural solution would be to approximate Equation (4) during retrieval. We

Models	MARCO Dev		TREC DL19		TREC DL20		BEIR (13 tasks)	Space&Time Efficiency	
	MRR@10	R@1k	nDCG@10	R@1k	nDCG@10	R@1k	nDCG@10	Disk (GB)	Latency (ms/query)
<i>Models trained with only BM25 hard negatives from MS MARCO Passages</i>									
BM25	0.188	0.858	0.506	0.739	0.488	0.733	0.440	0.7	40
DPR	0.319	0.941	0.611	0.742	0.591	0.796	0.375	26.0	2015
SPLADE	0.340	0.965	0.683	0.813	0.671	0.823	0.453	2.6	475
COIL	0.353	0.967	0.704	0.835	0.688	0.841	0.483	78.5	3258
ColBERT	0.360	0.968	0.694	0.830	0.676	0.837	0.453	154.3	-
SLIM	0.358	0.962	0.701	0.824	0.640	0.854	0.451	18.2	580
<i>Models trained with further pre-training/hard-negative mining/distillation</i>									
coCondenser	0.382	0.984	0.674	0.820	0.684	0.839	0.420	26.0	2015
SPLADE-v2	0.368	0.979	0.729	0.865	0.718	0.890	0.499	4.1	2710
ColBERT-v2	0.397	0.985	0.744	0.882	0.750	0.894	0.500	29.0	3275
SLIM++	0.404	0.968	0.714	0.842	0.702	0.855	0.490	17.3	550

Table 1: In-domain and out-of-domain evaluation on MS MARCO Passages, TREC DL 2019/2020, and BEIR. “-” means not practical to evaluate on a single CPU. Latency is benchmarked on a single CPU and query encoding time is excluded.

first unfold the dot product in Equation (4):

$$s(q, d) = \sum_{i=1}^N \max_{j=1}^M \sum_{k=1}^{|V|} \phi_{q_i}^{(k)} \phi_{d_j}^{(k)}, \quad (5)$$

where $\phi_{q_i}^{(k)}$ is the k th elements of ϕ_{q_i} . As each ϕ_{q_i} across different q_i all shares the same lexical space and the values are non-negative, we can easily derive its upper-bound and lower-bound:

$$\sum_{i=1}^N \max_{j=1}^M \sum_{k=1}^{|V|} e_{q_i}^{(k)} \phi_{d_j}^{(k)} \leq s(q, d) \leq \sum_{i=1}^N \sum_{k=1}^{|V|} \max_{j=1}^M \phi_{q_i}^{(k)} \phi_{d_j}^{(k)}, \quad (6)$$

$$e_{q_i}^{(k)} = \begin{cases} \phi_{q_i}^{(k)}, & \text{if } k = \operatorname{argmax}_k \phi_{q_i}^{(k)} \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

Then, the lower-bound score $s_l(q, d)$ and upper-bound score $s_h(q, d)$ can be further factorized as:

$$\begin{aligned} s_l(q, d) &= \sum_{i=1}^N \max_{j=1}^M \sum_{k=1}^{|V|} e_{q_i}^{(k)} \phi_{d_j}^{(k)} = \left(\sum_{i=1}^N e_{q_i} \right)^T \left(\max_{j=1}^M \phi_{d_j} \right); \\ s_h(q, d) &= \sum_{i=1}^N \sum_{k=1}^{|V|} \max_{j=1}^M \phi_{q_i}^{(k)} \phi_{d_j}^{(k)} = \left(\sum_{i=1}^N \phi_{q_i} \right)^T \left(\max_{j=1}^M \phi_{d_j} \right) \end{aligned} \quad (8)$$

where both the sum and max operations will be element-wise if the targets are vectors. We see that these two equations resemble the form of SPLADE in Equation (3), where queries and documents are encoded into sequence-level representations independently. In this way, $\max_{j=1}^M \phi_{d_j}$ in Equation (8) can be pre-computed and indexed offline. To approximate $s(q, d)$ in Equation (5), we use the linear interpolation between the lower- and the upper-bound:

$$\begin{aligned} s_a(q, d) &= \beta \cdot s_l(q, d) + (1 - \beta) \cdot s_h(q, d) \\ &= \left(\sum_{i=1}^N (\beta \cdot e_{q_i} + (1 - \beta) \cdot \phi_{q_i}) \right)^T \left(\max_{j=1}^M \phi_{d_j} \right) \end{aligned} \quad (9)$$

where $s_a(q, d)$ is the approximate score of SLIM and $\beta \in [0, 1]$ is the interpolation coefficient. As shown in Equation (9), we can first

fuse the query representation before retrieval and fine-tune the coefficient β on a validation set. It is worth mentioning that the approximation is applied after the SLIM model is trained using Equation (4). During indexing, we use Pyserini [22] to index the sequence level-representation $\max_{j=1}^M \phi_{d_j}$ and use Scipy [41] to store the sparse token vectors. During retrieval, we use Equation (9) to retrieve the top- k candidates with the LuceneImpactSearcher in Pyserini. For score refinement, we extract the stored sparse token vectors for the top- k documents and use Equation (4) to refine the candidate list. This two-stage breakdown yields a good effectiveness-efficiency trade-off, as we can aggressively reduce the first-stage retrieval latency without too much accuracy loss due to the score refinement.

To further reduce the memory footprint and improve latency for the first-stage retrieval, we apply two post-hoc pruning strategies on the inverted index to remove: (1) tokens that have term importance below a certain threshold; (2) postings that exceed a certain length (i.e., the inverse document frequency is below a threshold). Since all the term importance values are non-negative, tokens with smaller weights contribute less to the final similarity. Moreover, tokens with long postings mean that they frequently occur in different documents, which yield low inverse document frequencies (IDFs) and therefore can be safely discarded.

3 EXPERIMENTS

We evaluate our model and baselines on MS MARCO Passages [32] and its shared tasks, TREC DL 2019/2020 passage ranking tasks [3]. We train the baseline models on MS MARCO Passages and report results on its dev-small set and TREC DL 2019/2020 test queries following the same setup in CITADEL [20]. We further evaluate the models on the BEIR benchmark [40], which consists of a diverse set of 18 retrieval tasks across 9 domains. Following previous work [5, 38], we only evaluate 13 datasets due to license restrictions.¹ The evaluation metrics are MRR@10, nDCG@10, and

¹CQADupstack, Signal-1M(RT), BioASQ, Robust04, TREC-NEWS are excluded.




Figure 2: Effectiveness-Efficiency trade-offs of SLIM (w/ and w/o score refinement) on MS MARCO Passages.

Recall@1000 (i.e., R@1K). Latency is benchmarked on a single Intel(R) Xeon(R) Platinum 8275CL CPU @ 3.00GHz and both the batch size and the number of threads are set to 1.

We follow the training procedure of CITADEL [20] to train SLIM and apply ℓ_1 regularization on the sparse token representations for sparsity control. SLIM is trained only with BM25 hard negatives in MS MARCO Passages, while SLIM⁺⁺ is trained with cross-encoder distillation and hard negative mining. For indexing and pruning, the default token weight pruning threshold is 0.5 and the IDF threshold is 3. We use Pyserini to index and retrieve the documents and use Scipy to store the sparse token vectors (CSR matrix). For retrieval, the fusion weight β in Equation (9) is set to 0.01. We first retrieve the top-4000 candidates using Pyserini’s LuceneImpactSearcher and then use the original sparse token vectors stored by Scipy to refine the ranked list and output the top-1000 documents.

Table 1 shows the in-domain evaluation results on MS MARCO Passages and TREC DL 2019/2020. SLIM and SLIM⁺⁺ manage to achieve accuracy comparable to ColBERT and ColBERT-v2 on MS MARCO Passages. However, the results seem unstable on TREC DL, which is expected as TREC DL has far fewer queries. For out-of-domain evaluation, the results are a little bit worse than the current state of the art but still close to ColBERT. For latency and storage, SLIM has a much smaller disk requirement and lower latency compared to ColBERT. For example, SLIM⁺⁺ achieves effectiveness comparable to ColBERT-v2 with an 83% decrease in latency on a single CPU and using 40% less storage.

Next, we show that a two-stage implementation of SLIM provides a good trade-off between effectiveness and efficiency. Figure 2 plots CPU latency vs. MRR@10/Recall@1000 on MS MARCO Passages using IDF thresholding for SLIM (w/o hard-negative mining and distillation). We set the minimum IDF threshold to $\{0, 0.5, 1, 1.5, 2, 2.5, 3\}$ and first-stage retrieval candidates top- k to $\{1000, 1500, 2000, 2500, 3000, 3500, 4000\}$ before score refinement (if any). Without the score refinement, the MRR@10 and Recall@1000 scores are high when the IDF threshold is small, but effectiveness drops rapidly when we increase the IDF threshold to trade accuracy for latency. In contrast, with the score refinement step, we see that with about 0.003 loss in MRR@10 and 0.01 loss in Recall@1000, latency improves drastically from about 1800 ms/query to 500 ms/query. The reason is that the score refinement step, which takes a small amount of time compared to retrieval, compensates for the errors from the aggressive pruning at the first stage, which shows that

our novel lower- and upper-bound fusion provides a good approximation of SLIM when the IDF threshold is low, and that the score refinement step is important as it allows more aggressive pruning for the first-stage retrieval.

4 RELATED WORK

Dense retrieval [17] gains much popularity as it is supported by multiple approximate nearest neighbor search libraries [11, 16]. To improve effectiveness, hard negative mining [42, 43] and knowledge distillation [14, 26] are often deployed. Recently, further pre-training for retrieval [7, 8, 10, 15, 28] is proposed to improve the fine-tuned effectiveness of downstream tasks.

Sparse retrieval systems such as BM25 [35] and tf-idf [36] encode documents into bags of words. Recently, pre-trained language models are used to learn contextualized term importance [1, 6, 21, 23, 30]. These models leverage current search toolkits such as Pyserini [22] to perform sparse retrieval, or contribute to hybrid approaches with dense retrieval [13, 24, 25, 39].

Besides ColBERT [18], COIL [9] accelerates retrieval by combining exact match and inverted index search. CITADEL [20] further introduces lexical routing to avoid the lexical mismatch issue in COIL. ME-BERT [29] and MVR [44] propose to use a portion of token embeddings for late interaction. ALIGNER [33] frames multi-vector retrieval as an alignment problem and uses entropy-regularized linear programming to solve it.

5 CONCLUSION

In this paper, we propose an efficient yet effective implementation for multi-vector retrieval using sparsified late interaction, which yields fast retrieval speed, small index size, and full compatibility with off-the-shelf search toolkits such as Pyserini. The key ingredients of SLIM include sparse lexical projection using the MLM layer, computing the lower- and upper-bounds of the late interaction, as well as token weight and postings pruning. Experiments on both in-domain and out-of-domain information retrieval datasets show that SLIM achieves comparable accuracy to ColBERT-v2 while being more efficient in terms of both space and time.

ACKNOWLEDGEMENTS

This research was supported in part by the Natural Sciences and Engineering Research Council (NSERC) of Canada; computational resources were provided by Compute Canada.

REFERENCES

- [1] Yang Bai, Xiaoguang Li, Gang Wang, Chaoliang Zhang, Lifeng Shang, Jun Xu, Zhaowei Wang, Fangshan Wang, and Qun Liu. 2020. SparTerm: Learning Term-based Sparse Representation for Fast Text Retrieval. <https://doi.org/10.48550/ARXIV.2010.00768>
- [2] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. Universal Sentence Encoder. <https://doi.org/10.48550/ARXIV.1803.11175>
- [3] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2021. Overview of the TREC 2020 deep learning track. <https://doi.org/10.48550/ARXIV.2102.07662>
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [5] Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. 2021. SPLADE v2: Sparse Lexical and Expansion Model for Information Retrieval. <https://doi.org/10.48550/ARXIV.2109.10086>
- [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*. Association for Computing Machinery. <https://doi.org/10.1145/3404835.3463098>
- [7] Luyu Gao and Jamie Callan. 2021. Condenser: a Pre-training Architecture for Dense Retrieval. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 981–993. <https://doi.org/10.18653/v1/2021.emnlp-main.75>
- [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, ZhuYun Dai, and Jamie Callan. 2021. COIL: Revisit Exact Lexical Match in Information Retrieval with Contextualized Inverted List. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Online, 3030–3042. <https://doi.org/10.18653/v1/2021.naacl-main.241>
- [10] Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 6894–6910. <https://doi.org/10.18653/v1/2021.emnlp-main.552>
- [11] Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. 2020. Accelerating Large-Scale Inference with Anisotropic Vector Quantization. In *Proceedings of the 37th International Conference on Machine Learning, ICMML 2020, 13-18 July 2020, Virtual Event (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 3887–3896. <http://proceedings.mlr.press/v119/guo20h.html>
- [12] Hua He and Jimmy Lin. 2016. Pairwise Word Interaction Modeling with Deep Neural Networks for Semantic Similarity Measurement. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, San Diego, California, 937–948. <https://doi.org/10.18653/v1/N16-1108>
- [13] Sebastian Hofstätter, Omar Khattab, Sophia Althammer, Mete Sertkan, and Allan Hanbury. 2022. Introducing Neural Bag of Whole-Words with ColBERT: Contextualized Late Interactions Using Enhanced Reduction. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. Association for Computing Machinery, 737–747. <https://doi.org/10.1145/3511808.3557367>
- [14] Sebastian Hofstätter, Sheng-Chieh Lin, Jheng-Hong Yang, Jimmy Lin, and Allan Hanbury. 2021. Efficiently Teaching an Effective Dense Retriever with Balanced Topic Aware Sampling. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. Association for Computing Machinery, 113–122. <https://doi.org/10.1145/3404835.3462891>
- [15] Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. Unsupervised Dense Information Retrieval with Contrastive Learning. <https://doi.org/10.48550/ARXIV.2112.09118>
- [16] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. Billion-Scale Similarity Search with GPUs. *IEEE Transactions on Big Data* 7, 03 (jul 2021), 535–547. <https://doi.org/10.1109/TBDATA.2019.2921572>
- [17] 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)*. Association for Computational Linguistics, Online, 6769–6781. <https://doi.org/10.18653/v1/2020.emnlp-main.550>
- [18] Omar Khattab and Matei Zaharia. 2020. ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. Association for Computing Machinery, 39–48. <https://doi.org/10.1145/3397271.3401075>
- [19] Wuwei Lan and Wei Xu. 2018. Neural Network Models for Paraphrase Identification, Semantic Textual Similarity, Natural Language Inference, and Question Answering. In *Proceedings of the 27th International Conference on Computational Linguistics*. Association for Computational Linguistics, Santa Fe, New Mexico, USA, 3890–3902. <https://aclanthology.org/C18-1328>
- [20] Minghan Li, Sheng-Chieh Lin, Barlas Oguz, Asish Ghoshal, Jimmy Lin, Yashar Mehdad, Wen-tau Yih, and Xilun Chen. 2022. CITADEL: Conditional Token Interaction via Dynamic Lexical Routing for Efficient and Effective Multi-Vector Retrieval. <https://doi.org/10.48550/ARXIV.2211.10411>
- [21] Jimmy Lin and Xueguang Ma. 2021. A Few Brief Notes on DeepImpact, COIL, and a Conceptual Framework for Information Retrieval Techniques. <https://doi.org/10.48550/ARXIV.2106.14807>
- [22] 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*. Association for Computing Machinery, 2356–2362. <https://doi.org/10.1145/3404835.3463238>
- [23] Sheng-Chieh Lin, Akari Asai, Minghan Li, Barlas Oguz, Jimmy Lin, Yashar Mehdad, Wen-tau Yih, and Xilun Chen. 2023. How to Train Your DRAGON: Diverse Augmentation Towards Generalizable Dense Retrieval. <https://doi.org/10.48550/ARXIV.2302.07452>
- [24] Sheng-Chieh Lin, Minghan Li, and Jimmy Lin. 2022. Aggretriever: A Simple Approach to Aggregate Textual Representation for Robust Dense Passage Retrieval. <https://doi.org/10.48550/ARXIV.2208.00511>
- [25] Sheng-Chieh Lin and Jimmy Lin. 2022. A Dense Representation Framework for Lexical and Semantic Matching. <https://doi.org/10.48550/ARXIV.2206.09912>
- [26] Sheng-Chieh Lin, Jheng-Hong Yang, and Jimmy Lin. 2021. In-Batch Negatives for Knowledge Distillation with Tightly-Coupled Teachers for Dense Retrieval. In *Proceedings of the 6th Workshop on Representation Learning for NLP (ReplANLP-2021)*. Association for Computational Linguistics, Online, 163–173. <https://doi.org/10.18653/v1/2021.repl4nlp-1.17>
- [27] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. <https://doi.org/10.48550/ARXIV.1907.11692>
- [28] Shuqi Lu, Di He, Chenyan Xiong, Guolin Ke, Waleed Malik, Zhicheng Dou, Paul Bennett, Tie-Yan Liu, and Arnold Overwijk. 2021. Less is More: Pretrain a Strong Siamese Encoder for Dense Text Retrieval Using a Weak Decoder. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 2780–2791. <https://doi.org/10.18653/v1/2021.emnlp-main.220>
- [29] Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2021. Sparse, Dense, and Attentional Representations for Text Retrieval. *Transactions of the Association for Computational Linguistics* 9 (2021), 329–345. https://doi.org/10.1162/tacl_a_00369
- [30] 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*. Association for Computing Machinery, 1723–1727. <https://doi.org/10.1145/3404835.3463030>
- [31] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to information retrieval*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511809071>
- [32] 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 *CoCo@ NIPS*. <https://www.microsoft.com/en-us/research/publication/ms-marco-human-generated-machine-reading-comprehension-dataset/>
- [33] Yujie Qian, Jinhyuk Lee, Sai Meher Karthik Duddu, Zhuyun Dai, Siddhartha Brahma, Iftekhar Naim, Tao Lei, and Vincent Y. Zhao. 2022. Multi-Vector Retrieval as Sparse Alignment. <https://doi.org/10.48550/ARXIV.2211.01267>
- [34] 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 and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, Hong Kong, China, 3982–3992. <https://doi.org/10.18653/v1/D19-1410>
- [35] Stephen E. Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Foundations and Trends in Information Retrieval* 3, 4 (2009), 333–389. <https://doi.org/10.1561/1500000019>

- [36] Gerard Salton and Christopher Buckley. 1988. Term-weighting approaches in automatic text retrieval. *Information Processing & Management* 24, 5 (1988), 513–523. [https://doi.org/10.1016/0306-4573\(88\)90021-0](https://doi.org/10.1016/0306-4573(88)90021-0)
- [37] Keshav Santhanam, Omar Khattab, Christopher Potts, and Matei Zaharia. 2022. PLAD: An Efficient Engine for Late Interaction Retrieval. <https://doi.org/10.48550/ARXIV.2205.09707>
- [38] Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2022. ColBERTv2: Effective and Efficient Retrieval via Lightweight Late Interaction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Seattle, United States, 3715–3734. <https://doi.org/10.18653/v1/2022.naacl-main.272>
- [39] Tao Shen, Xiubo Geng, Chongyang Tao, Can Xu, Kai Zhang, and Daxin Jiang. 2022. Unifire: A Unified Retriever for Large-Scale Retrieval. <https://doi.org/10.48550/ARXIV.2205.11194>
- [40] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogeneous Benchmark for Zero-shot Evaluation of Information Retrieval Models. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*. <https://openreview.net/forum?id=wCu6T5xFje>
- [41] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, António H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- [42] Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N. Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. Approximate Nearest Neighbor Negative Contrastive Learning for Dense Text Retrieval. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. <https://openreview.net/forum?id=zeFrFgyZln>
- [43] Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Jiafeng Guo, Min Zhang, and Shaoping Ma. 2021. Optimizing Dense Retrieval Model Training with Hard Negatives. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. Association for Computing Machinery, 1503–1512. <https://doi.org/10.1145/3404835.3462880>
- [44] Shunyu Zhang, Yaobo Liang, Ming Gong, Daxin Jiang, and Nan Duan. 2022. Multi-View Document Representation Learning for Open-Domain Dense Retrieval. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Dublin, Ireland, 5990–6000. <https://doi.org/10.18653/v1/2022.acl-long.414>
- [45] Yinan Zhang, Raphael Tang, and Jimmy Lin. 2019. Explicit Pairwise Word Interaction Modeling Improves Pretrained Transformers for English Semantic Similarity Tasks. <https://doi.org/10.48550/ARXIV.1911.02847>