End-to-End Retrieval with Learned Dense and Sparse Representations Using Lucene

Haonan Chen,[1] **Carlos Lassance,**[2] **and Jimmy Lin**[3]

[1] Department of Electrical and Computer Engineering, University of Waterloo
[2] Naver Labs Europe
[3] David R. Cheriton School of Computer Science, University of Waterloo

Abstract

The bi-encoder architecture provides a framework for understanding machine-learned retrieval models based on dense and sparse vector representations. Although these representations capture parametric realizations of the same underlying conceptual framework, their respective implementations of top-k similarity search require the coordination of different software components (e.g., inverted indexes, HNSW indexes, and toolkits for neural inference), often knitted together in complex architectures. In this work, we ask the following question: What's the simplest design, in terms of requiring the fewest changes to existing infrastructure, that can support end-to-end retrieval with modern dense and sparse representations? The answer appears to be that Lucene is sufficient, as we demonstrate in Anserini, a toolkit for reproducible information retrieval research. That is, effective retrieval with modern single-vector neural models can be efficiently performed directly in Java on the CPU. We examine the implications of this design for information retrieval researchers pushing the state of the art as well as for software engineers building production search systems.

1 Introduction

The popular bi-encoder architecture has emerged as the dominant conceptual framework around which retrieval models are organized today. This design formulates retrieval in terms of comparisons between vectors derived from queries and pieces of content (passages, documents, and even images and multimedia). More precisely, the core retrieval problem can be framed as nearest neighbour search: given a query vector, the system's task is to return the k "closest" content vectors, typically in terms of a simple operation such as the inner (dot) product.

Central to the bi-encoder design is representation learning, where deep learning techniques are typically applied to learn vector representations that maximize the inner product between queries and relevant content while minimizing the inner product between queries and non-relevant content. At a high-level, these vector representations can be categorized as *dense*, typically contextual representations derived from transformers [Vaswani et al., 2017] that capture latent semantics, or *sparse*, where the dimensions of the vectors are defined by the vocabulary space. Thus, we arrive today at retrieval models based on either learned dense or learned sparse representations.

Within this conceptual framework, we focus on the software infrastructure necessary to realize end-to-end retrieval with learned dense and sparse representations. In particular, we ask: What's the simplest design that could "work"? We operationalize this in terms of minimality, or requiring the smallest number of changes to existing, established infrastructure. Our answer might be somewhat surprising, at least to some: We find that Lucene is sufficient. Inverted indexes using the "fake words" trick enable retrieval with learned sparse representations. Lucene's HNSW indexes enable retrieval with learned dense representations. Finally, with the integration of the ONNX Runtime, query inference
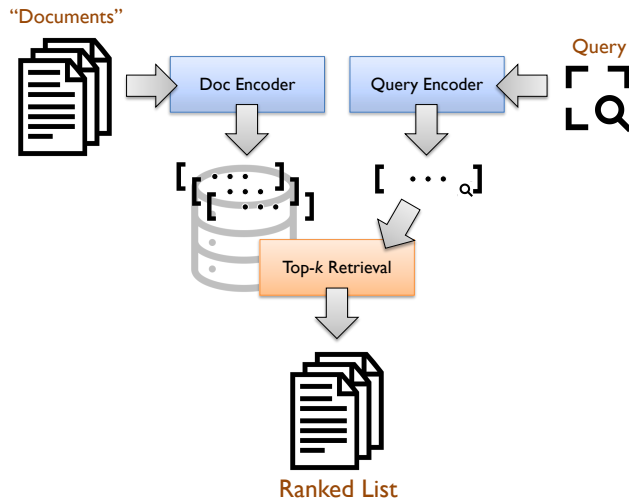
Figure 1: Illustration of the bi-encoder architecture, where encoders generate vector representations from queries and documents. In this design, retrieval is formulated as k-nearest neighbour search in vector space.

can be performed directly from Lucene. This means that effective end-to-end retrieval with modern neural models can be efficiently performed in Java, on the CPU. In particular, we were able to "pare away" two components that many have perhaps thought to be indispensable to neural retrieval today: a separate vector store to enable top-k retrieval on dense vectors and Python as a bridge to neural inference engines.

We demonstrate end-to-end retrieval with learned dense and sparse representations in Anserini [Yang et al., 2018], an existing Lucene toolkit for reproducible information retrieval research. This design holds implications for information retrieval researchers as well as for engineers who build production search systems. For researchers, simplicity accelerates progress by lowering hurdles and streamlining research iteration cycles. For system builders, Anserini primarily provides a prototype of what's possible: Since Lucene resides at its core, anything that can be implemented in Anserini can just as easily be integrated into Elasticsearch, OpenSearch, Solr, or any other software component in the Lucene ecosystem. This suggests that enterprises can take advantage of the latest advances in neural retrieval models without abandoning existing infrastructure investments (e.g., the ELK stack). Capabilities that are already present in Lucene provide a compelling adoption path.

2   Background and Related Work

An illustration of the bi-encoder architecture is shown in Figure 1, where retrieval is framed as nearest neighbour search in vector space. This design assumes the existence of encoders that convert documents and queries into vector representations. Note that throughout this work we use "documents" in a generic sense to refer to content that we wish to search, even though in actuality they may represent passages, web pages, or even images and other multimedia.

Given the vector representation of a query, the system's task is to efficiently retrieve the top-k most similar document vectors with respect to a simple operation such as the inner (dot) product. Lin [2021] observed that in this framework, there are two important characteristics of these representations:

- the basis of the vector space—for example, dense vectors where the dimensions capture latent semantics, or sparse vectors where the dimensions are defined by the vocabulary;
- whether the vectors are generated by a machine-learned model (typically in a supervised manner) or via some heuristic weighting function (typically in an unsupervised manner).

Cast in this framework, "traditional" bag-of-words BM25 retrieval [Robertson and Zaragoza, 2009] can be characterized as sparse (lexical) representations generated by a heuristic weighting func-

tion (that is, the BM25 scoring function). On the other hand, dense retrieval models such as DPR [Karpukhin et al., 2020] can be characterized as learned dense representations. These models, along with vector stores for manipulating dense vectors, dominate today's discourse.

2.1 The Vector Store Narrative

Let us begin by recounting the "standard narrative" behind "semantic search" today. The story, parroted endlessly on social media, blog posts, and in other venues, typically begins with the awesome power of transformers to generate dense (semantic) representations that are able to capture the "meaning" of text. Such learned dense representations, or embeddings, are supposedly light-years ahead of bag-of-words lexical representations in terms of retrieval effectiveness. The truth is far less hyperbolic, as dense retrieval models still underperform traditional lexical models in at least some scenarios [Thakur et al., 2021, Sciavolino et al., 2021, Kamalloo et al., 2023]. Embeddings are also touted as the key ingredient to retrieval augmentation, a popular method of enhancing the capabilities of large language models (LLMs) [Mialon et al., 2023, Asai et al., 2023].

But there's one catch: top-k retrieval with dense vector representations requires a completely new software stack built around "vector stores". Listening to the breathless hype surrounding these new-fangled contraptions, one could easily imagine the obsolescence of the venerable inverted index, the workhorse on which bag-of-words retrieval models rely. One goal of this paper, to butcher a quote often (mis-)attributed to Mark Twain, is to push back against this narrative: "reports of the death of inverted indexes are greatly exaggerated".

While learned dense representations have dominated social media and blog posts—with countless startups peddling vector stores—many researchers have been developing learned *sparse* representations in (relative) obscurity, e.g., Dai and Callan [2020], Bai et al. [2020], Gao et al. [2021], Mallia et al. [2021], Lin and Ma [2021], Nguyen et al. [2023], just to cite a few. This class of models takes advantage of the same innovations that have fed learned dense retrieval models, but instead generate sparse vectors that are grounded in the vocabulary space. With recent advances, models such as SPLADE++ [Formal et al., 2022] have achieved retrieval effectiveness at least on par with dense retrieval on a variety of standard benchmarks [Kamalloo et al., 2023].

Retrieval with sparse representations can be implemented with inverted indexes using the widely-known "fake words" trick [Mackenzie et al., 2022], which means that they are compatible with existing libraries such as Lucene. This has already been demonstrated in the Anserini IR toolkit [Yang et al., 2018] built around Lucene, which has incorporated sparse retrieval models since 2021 [Lin and Ma, 2021]. Furthermore, Devins et al. [2022] showed that what's possible in Anserini is equally possible in Elasticsearch and any other component in the broader Lucene ecosystem.

In truth, *both* dense and sparse representations are indispensable, as the combination of evidence from both approaches leads to higher-quality output than either alone [Ma et al., 2021, Kamalloo et al., 2023]. Yet, the fact remains that top-k retrieval with dense and sparse representations require different machinery: while sparse representations are amenable to inverted indexes, current best practices recommend the use of hierarchical navigable small-world network (HNSW) indexes for top-k retrieval on dense representations [Malkov and Yashunin, 2020]. However, Lin et al. [2023] argued that the need for the capabilities offered by HNSW indexes doesn't necessarily require full-blown vector stores as independent software components. In fact, the authors pointed out that Lucene provides implementations of HNSW indexes[1] that are adequate for servicing dense retrieval models. They argued that, from a simple cost–benefit analysis, there does not appear to be a compelling need for dedicated vector stores, since vector search capabilities already exist in Lucene. In this sense, both techniques are converging, with Lucene adding vector search and most vector stores developers starting to realize the need for basic search techniques, such as filters and inverted indexes.[2]

2.2 The ONNX Runtime

There is one final piece to the puzzle: retrieval with learned dense and sparse representations requires encoders for generating the representation vectors, typically using transformers. This breaks down into two distinct processing phases: encoding the corpus and encoding the queries. Encoding a (static)

---

corpus is a one-time operation that demands high throughput and is usually impractical without GPUs. We treat this as a form of preprocessing and assume that the entire corpus has already been converted into vector representations prior to ingestion into Lucene. Query encoding, on the other hand, must be performed at search time, and unlike corpus encoding, is latency-sensitive. Furthermore, it would be desirable if query encoding could be performed on the CPU, since such an implementation would result in significantly lower costs, better scaling characteristics (since it is easier to add CPU cores than GPUs), and expand the range of hardware that search applications can be deployed to. This is exactly the scenario we explore: query encoding on the CPU.

When discussing neural inference, Python immediately comes to mind, as it has emerged as the dominant front-end language to popular neural toolkits such as PyTorch [Paszke et al., 2019]. But Lucene is implemented in Java, which presents a language mismatch. Previously, researchers have addressed this issue in different ways. Pyserini [Lin et al., 2021] provides Anserini bindings in Python, and thus for sparse retrieval models, query encoding can be performed using PyTorch before the query is passed to Anserini (via a Python-to-Java bridge) for top-k retrieval. Previously, without the benefit of PyTorch bindings via Python, Anserini by itself relied on pre-encoded queries for both dense and sparse retrieval with neural models. That is, inference is applied to the queries "outside the system" and cached. This reflects the scenario where a Lucene-based search engine calls, for example, an external inference API to perform query encoding. While this is certainly an acceptable design pattern, it adds latency (from inter-process communication) and lacks the self-contained elegance that would come with an in-process query encoding solution (i.e., Java-based neural inference).

This is where ONNX (Open Neural Network Exchange), an open standard for computational graphs, comes in. The ONNX Runtime provides cross-platform, cross-language, and cross-device support for neural inference—including using Java, the language of Lucene's implementation. This means that with ONNX integration in Anserini, it is possible to perform retrieval with learned dense and sparse models *directly* in Java, on the CPU—that is, without needing PyTorch, a vector store, or even a separate HNSW library such as Faiss [Johnson et al., 2019].

2.3   Putting Everything Together

In other words, we find that Lucene is all you need for end-to-end retrieval with learned dense and sparse representations. To be more precise, this work focuses on single-vector representations, setting aside multi-vector approaches such as ColBERT for dense representations [Khattab and Zaharia, 2020] and SLIM for sparse representations [Li et al., 2023]. Lucene is the answer to the motivating question we posed: What's the simplest design, in terms of requiring the fewest changes to existing established infrastructure, that suffices to provide all our desired capabilities?

Although our experiments are conducted with Anserini, the capabilities we describe are easily exposed in Elasticsearch, OpenSearch, Solr, or any other component in the Lucene ecosystem. Pyserini still provides Anserini bindings in Python, thereby offering seamless connections to other components of the modern Python-centric deep learning ecosystem.

To be clear, this is not the first work to examine or integrate ONNX into Lucene or any other search application. In fact, our group has been experimenting with ONNX since 2017 [Tu et al., 2018]. More recently, ONNX has been integrated into Elasticsearch, currently as an experimental feature.[3] Similarly, Vespa has integrated ONNX for some time now.[4] However, the contribution of this paper is to "put all the pieces together". To our knowledge, our implementation is the first that (1) is fully open-source, (2) covers both learned dense and sparse models, (3) provides end-to-end retrieval, including query encoding, (4) integrates with infrastructure that already has a massive install base— for the benefit of practitioners who focus on building production systems, and (5) provides support for formal evaluations on common IR benchmarks today—for the benefit of researchers who focus on model development.

3   Design and Experimental Setup

The core technical contribution of this paper is the integration of the ONNX Runtime into the Anserini IR toolkit to provide Java-based in-process query encoding. To accomplish this, we implemented an

---

[3] https://www.elastic.co/guide/en/machine-learning/master/ml-nlp-deploy-models.html

[4] https://blog.vespa.ai/stateful-model-serving-how-we-accelerate-inference-using-onnx-runtime/

	dev		DL19			DL20		
	RR@10	R@1k	AP	nDCG@10	R@1k	AP	nDCG@10	R@1k
BM25	0.184	0.853	0.301	0.506	0.750	0.286	0.480	0.786
cosDPR-distil	0.389	0.975	0.466	0.725	0.822	0.487	0.703	0.852
SPLADE++ ED	0.383	0.983	0.505	0.731	0.873	0.500	0.720	0.900

Table 1: Effectiveness of the learned dense model (cosDPR-distil) and the learned sparse model (SPLADE++ ED) on the MS MARCO passage corpus.

abstract class called `OnnxEncoder`, which defines the (abstract) `encode` method for converting a string (the query) into a dense or sparse vector representation. This abstract class has two sub-classes (also abstract), called `DenseEncoder` and `SparseEncoder`, for generating dense and sparse representations, respectively. Concrete retrieval models instantiate one or the other class, as appropriate. We experimented with an illustrative model from each class:

- For the dense model, we selected the cosDPR-distil model [Ma et al., 2023a], primarily because Lucene's implementation of HNSW is restricted to top-k retrieval using cosine similarity (as opposed to inner products in the general case). The key difference lies in vector normalization, which most existing models do not perform. Nevertheless, cosDPR-distil is trained according to best practices: an initial version of the model is used to mine hard negatives, which are then used to train a cross-encoder reranker and then distilled into the final retrieval model.

- For the sparse model, we used SPLADE++ Cocondenser Ensemble-Distil [Formal et al., 2022] (SPLADE++ ED for short). As this model is well known, we simply refer interested readers to the cited reference for details.

To prepare the models for the ONNX Runtime, we first began by converting the PyTorch implementation of the encoders from HuggingFace [Wolf et al., 2020]. We loaded the models using the HuggingFace API and then exported model weights using the `torch.onnx.export()` function available in PyTorch. To further optimize the ONNX compute graphs extracted from the models, we used the `onnxruntime.transformers.optimizer` tool from the ONNX Runtime library, specifically developed for transformer models. Note that further optimization is possible by quantizing weights to lower precision, but we decided to stay aligned with existing model weights and results.

As the focus of this work is query encoding at search time, we assumed that corpus encoding is handled separately and that we are provided with the raw representation vectors for ingestion into Anserini. Our indexing implementation is based on the work described in Ma et al. [2023a], except based on Lucene 9.8.0. More details are provided in the results section.

We performed evaluations on the MS MARCO passage ranking test collection [Craswell et al., 2021], a standard benchmark dataset used by researchers today. Results are reported on the 6980 queries from the development (dev) set, as well as on queries from the TREC 2019 and 2020 Deep Learning Tracks [Craswell et al., 2019, 2020]. All our experiments followed standard evaluation methodology and were performed on a desktop system with an AMD Ryzen 9 5900X 12-Core processor.[5] Our test system ran Ubuntu 22.04.3 LTS, OpenJDK 11.0.20.1, and Python 3.10.13.

4 Results

Effectiveness results are shown in Table 1, comparing cosDPR-distil and SPLADE++ ED to baseline BM25 in terms of standard metrics across three sets of queries on the MS MARCO passage corpus. At a high level, these scores are competitive in terms of retrieval quality to models that are available today, although they remain below the state of the art [Ma et al., 2023b]. One easy way to increase effectiveness is to perform ensembling of the dense and sparse models (in this case, average of normalized scores), which improves RR@10 on the dev queries to 0.408. This fusion run is significantly better than both models individually and closer to the state of the art. However, our main goal here is to provide a point of reference and also to show that dense and sparse retrieval models are roughly on par in terms of effectiveness.

---

[5] `https://www.amd.com/en/products/cpu/amd-ryzen-9-5900x`

	1 thread throughput (qps)		12 threads throughput (qps)	
	pre-encoded	ONNX	pre-encoded	ONNX
BM25	29.8 ± 1.2	-	278.9 ± 9.2	-
cosDPR-distil	27.5 ± 0.9	25.5 ± 0.3	158.0 ± 0.3	72.1 ± 0.5
SPLADE++ ED	5.1 ± 0.2	5.0 ± 0.3	48.8 ± 2.0	35.3 ± 0.2

Table 2: Performance of Lucene on the MS MARCO dev queries, comparing single-thread query throughput using pre-encoded queries vs. ONNX and query throughput on 12 threads using pre-encoded queries vs. ONNX.

The HNSW index for cosDPR-distil was constructed using the indexer implemented in Anserini, with the parameter M set to 16 and efC set to 1000, using 16 threads, but optimized down to a single index segment (which is a very time-consuming operation). The cosDPR-distil index is around 27G. To better understand the impact of different parameter settings, we refer the reader to Ma et al. [2023a].[6] Due to the non-determinism associated with HNSW indexing, we report scores to three digits on our specific index instance, which are slightly different from the "reference values" presented in the Anserini reproducibility documentation. Effectiveness may vary from run to run, with further dependence on index construction settings (M, efC, and the number of indexing threads).

In contrast, SPLADE++ ED provides deterministic results using an inverted index of only 2.3G, which is about the same size as the inverted index used for BM25, which is 2.6G. Here, our experiments used an index configuration that only stores impact scores and is also optimized down into a single index segment. Overall, the sparse index has a smaller memory footprint and can be built faster than the HNSW index. However, as we will see below, search is slower with inverted indexes than with HNSW indexes.

Having established the effectiveness of our representative learned dense and sparse retrieval models, we turn our attention to efficiency. Results are shown in Table 2, where we examine two conditions: In the first set of experiments, we performed retrieval on a single thread with both models to accurately measure query latency with pre-encoded queries (i.e., cached query representations) and with "on-the-fly" query encoding using the ONNX Runtime. This comparison quantifies the overhead of neural inference for query encoding. In the second set of experiments, we performed retrieval using 12 threads, fully saturating the CPU of our test machine, to measure throughput. Once again, we compared pre-encoded queries with on-the-fly query encoding using the ONNX Runtime.

For both sets of experiments, we used the 6980 dev set queries from MS MARCO, retrieving 1000 hits; there are too few TREC queries to obtain reliable measurements. We quantified performance in terms of throughput, measuring queries per second. For all conditions, we conducted three warm-up runs to mitigate noise during benchmarking. We computed the mean over three subsequent runs, reporting 95% confidence intervals. The throughput of BM25 is provided as a point of reference.

These results show that on a single thread, the overhead of query encoding with the ONNX Runtime is relatively modest: we see this by comparing "pre-encoded" vs. "ONNX" under the "1 thread" condition for both models. With a single thread, retrieval with HNSW indexes is just slightly slower than retrieval with inverted indexes computing BM25 rankings. However, retrieval with SPLADE++ ED is much slower than with cosDPR-distil. We note (again) here that the HNSW indexes have been optimized down to a single segment, which may be an unrealistic assumption for real-world document collections that are not static.

As expected, none of the models exhibit linear scaling (i.e., $12\times$ more threads does not yield a $12\times$ increase in throughput), indicating contention in accessing the index structures (under the pre-encoded condition). The performance gap between BM25 and cosDPR-distil (pre-encoded queries) widens as we scale from 1 to 12 threads, but the gap between cosDPR-distil (pre-encoded) and SPLADE++ ED (pre-encoded) narrows. These results suggest that, overall, inverted indexes scale better than HNSW indexes under multiple threads. Interestingly, the gap between "pre-encoded" and "ONNX" is much larger for 12 threads than for 1 thread, indicating that contention is more severe for ONNX query inference than for accessing index structures.

---

[6] Although note that those experiments were conducted with Lucene 9.4.2; the version we used, Lucene 9.8.0, contains additional improvements.

5   Discussion

This paper operationalizes simplicity of design in terms of the minimal number of changes to a baseline architecture that is necessary to gain a certain set of capabilities—in this case, end-to-end retrieval with learned dense and sparse representations. As previously articulated in Lin et al. [2023], we begin with the simple observation that search is a brownfield application, which Wikipedia defines as "a term commonly used in the information technology industry to describe problem spaces needing the development and deployment of new software systems in the immediate presence of existing (legacy) software applications/systems." Additionally, "this implies that any new software architecture must take into account and coexist with live software already in situ."

That is, most enterprises today already have substantial existing investments in search infrastructure, specifically in Lucene-powered platforms such as Elasticsearch, OpenSearch, and Solr. The Lucene ecosystem is stable, mature, well supported, and has a massive install base. In this context, we have shown with Anserini that Lucene coupled with the ONNX Runtime provides all the components necessary to support end-to-end retrieval with the latest dense and sparse models on the CPU, which makes "bets" on less mature alternatives (such as dedicated vector stores) less compelling. Continued investments to grow capabilities in existing Lucene-based stacks are the most likely adoption path we see moving forward.

As advocated in Devins et al. [2022], alignment between research in neural retrieval and the practice of building search applications brings benefits to both industry and academia (and the research community in general). For researchers, more rapid technology transfer is an attractive upside, and from this perspective, toolkits such as Anserini and Pyserini help accelerate progress. Based on our experience, the most important desideratum for researchers is rapid experimental cycles, and thus a pure Java solution (or a Java solution with a thin Python wrapper) can be compelling. One piece of feedback we have received about the current implementation of Pyserini is its complex dependency chain, which makes installation and configuration a non-trivial barrier for many.[7] With the capabilities presented here, we see the potential of simplifying first-stage retrieval infrastructure, thereby enabling faster iteration for researchers.

6   Conclusion

The dominant narrative today promotes dense retrieval models as a revolutionary advance in search technology that requires entirely new data management infrastructure built around vector stores. The hype suggests that these models render lexical-based techniques obsolete, and by extension, the venerable inverted index.

Our work pushes back against this narrative. While transformers have indeed led to significant advances in search effectiveness, the improvements are better characterized as evolutionary. Furthermore, transformer-based retrieval models are not limited to generating dense representations. Experimental evidence by many researchers and reproduced here has shown that an entirely different class of learned *sparse* models are on par with dense retrieval models in terms of effectiveness.

The truth is that we need *both* dense and sparse representations, and that their fusion yields the best results. However, implementation differences remain, which present practical obstacles to building real-world search systems. Our position, previously argued in Lin et al. [2023], is based on the observation that search is a brownfield application, and as a result, many organizations have already made substantial investments that they are unlikely to abandon. In this context, we believe that the most compelling path forward is the design that requires the fewest changes to existing infrastructure—for most, already built on Lucene—that enables the largest gain in capabilities. In that sense, Lucene is all you need.

---

[7]Python "dependency hell" is no doubt a common experience.

References

A. Asai, S. Min, Z. Zhong, and D. Chen. Retrieval-based language models and applications. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 6: Tutorial Abstracts)*, pages 41–46, Toronto, Canada, July 2023.

Y. Bai, X. Li, G. Wang, C. Zhang, L. Shang, J. Xu, Z. Wang, F. Wang, and Q. Liu. SparTerm: Learning term-based sparse representation for fast text retrieval. *arXiv:2010.00768*, 2020.

N. Craswell, B. Mitra, E. Yilmaz, D. Campos, and E. M. Voorhees. Overview of the TREC 2019 Deep Learning Track. In *Proceedings of the Twenty-Eighth Text REtrieval Conference Proceedings (TREC 2019)*, Gaithersburg, Maryland, 2019.

N. Craswell, B. Mitra, E. Yilmaz, and D. Campos. Overview of the TREC 2020 Deep Learning Track. In *Proceedings of the Twenty-Ninth Text REtrieval Conference Proceedings (TREC 2020)*, Gaithersburg, Maryland, 2020.

N. Craswell, B. Mitra, D. Campos, E. Yilmaz, and J. Lin. MS MARCO: Benchmarking ranking models in the large-data regime. In *Proceedings of the 44th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2021)*, pages 1566–1576, 2021.

Z. Dai and J. Callan. 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 (SIGIR 2020)*, pages 1533–1536, 2020.

J. Devins, J. Tibshirani, and J. Lin. Aligning the research and practice of building search applications: Elasticsearch and Pyserini. In *Proceedings of the 15th ACM International Conference on Web Search and Data Mining (WSDM 2022)*, pages 1573–1576, 2022.

T. Formal, C. Lassance, B. Piwowarski, and S. Clinchant. 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 (SIGIR 2022)*, pages 2353–2359, Madrid, Spain, 2022.

L. Gao, Z. Dai, and J. Callan. 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*, pages 3030–3042, 2021.

J. Johnson, M. Douze, and H. Jégou. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547, 2019.

E. Kamalloo, N. Thakur, C. Lassance, X. Ma, J.-H. Yang, and J. Lin. Resources for brewing BEIR: Reproducible reference models and an official leaderboard. *arXiv:2306.07471*, 2023.

V. Karpukhin, B. Oğuz, S. Min, P. Lewis, L. Wu, S. Edunov, D. Chen, and W.-t. Yih. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, 2020.

O. Khattab and M. Zaharia. ColBERT: Efficient and effective passage search via contextualized late interaction over BERT. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2020)*, pages 39–48, 2020.

M. Li, S.-C. Lin, X. Ma, and J. Lin. SLIM: Sparsified late interaction for multi-vector retrieval with inverted indexes. In *Proceedings of the 46th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2023)*, pages 1954–1959, Taipei, Taiwan, 2023.

J. Lin. A proposed conceptual framework for a representational approach to information retrieval. *arXiv:2110.01529*, 2021.

J. Lin and X. Ma. A few brief notes on DeepImpact, COIL, and a conceptual framework for information retrieval techniques. *arXiv:2106.14807*, 2021.

J. Lin, X. Ma, S.-C. Lin, J.-H. Yang, R. Pradeep, and R. Nogueira. Pyserini: A Python toolkit for reproducible information retrieval research with sparse and dense representations. In *Proceedings of the 44th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2021)*, pages 2356–2362, 2021.

J. Lin, R. Pradeep, T. Teofili, and J. Xian. Vector search with OpenAI embeddings: Lucene is all you need. *arXiv:2308.14963*, 2023.

X. Ma, K. Sun, R. Pradeep, and J. Lin. A replication study of dense passage retriever. *arXiv:2104.05740*, 2021.

X. Ma, T. Teofili, and J. Lin. Anserini gets dense retrieval: Integration of Lucene's HNSW indexes. In *Proceedings of the 32nd International Conference on Information and Knowledge Management (CIKM 2023)*, pages 5366–5370, Birmingham, the United Kingdom, 2023a.

X. Ma, L. Wang, N. Yang, F. Wei, and J. Lin. Fine-tuning LLaMA for multi-stage text retrieval. *arXiv:2310.08319*, 2023b.

J. Mackenzie, A. Trotman, and J. Lin. Efficient document-at-a-time and score-at-a-time query evaluation for learned sparse representations. *ACM Transactions on Information Systems*, 41: Article No. 96, 2022.

Y. A. Malkov and D. A. Yashunin. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *Transactions on Pattern Analysis and Machine Intelligence*, 42(4):824–836, 2020.

A. Mallia, O. Khattab, T. Suel, and N. Tonellotto. Learning passage impacts for inverted indexes. In *Proceedings of the 44th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2021)*, pages 1723–1727, 2021.

G. Mialon, R. Dessì, M. Lomeli, C. Nalmpantis, R. Pasunuru, R. Raileanu, B. Rozière, T. Schick, J. Dwivedi-Yu, A. Celikyilmaz, E. Grave, Y. LeCun, and T. Scialom. Augmented language models: a survey. *arXiv:2302.07842*, 2023.

T. Nguyen, S. MacAvaney, and A. Yates. A unified framework for learned sparse retrieval. In *European Conference on Information Retrieval*, pages 101–116. Springer, 2023.

A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32, 2019.

S. Robertson and H. Zaragoza. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends in Information Retrieval*, 3(4):333–389, 2009.

C. Sciavolino, Z. Zhong, J. Lee, and D. Chen. Simple entity-centric questions challenge dense retrievers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6138–6148, Online and Punta Cana, Dominican Republic, 2021.

N. Thakur, N. Reimers, A. Rücklé, A. Srivastava, and I. Gurevych. 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)*, 2021.

Z. Tu, M. Li, and J. Lin. Pay-per-request deployment of neural network models using serverless architectures. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 6–10, New Orleans, Louisiana, 2018.

A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30*, pages 5998–6008, Long Beach, California, 2017.

T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest, and A. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, 2020.

P. Yang, H. Fang, and J. Lin. Anserini: Reproducible ranking baselines using Lucene. *Journal of Data and Information Quality*, 10(4):Article 16, 2018.