

BiCA: Effective Biomedical Dense Retrieval with Citation-Aware Hard Negatives

Aarush Sinha^{1*†}, Pavan Kumar S^{2,3,4†}, Roshan Balaji^{2,3,4}, Nirav Pravinbhai Bhatt^{2,3,4‡}

¹Vellore Institute of Technology (VIT) Chennai, Tamil Nadu, India

²BioSystems Engineering and Control (BiSECT) Lab, Department of Biotechnology

³Wadhvani School of Data Science and AI, Indian Institute of Technology (IIT) Madras, Tamil Nadu, India

⁴The Centre for Integrative Biology and Systems medicinE (IBSE), IIT Madras, Chennai, Tamil Nadu, India
aarush.sinha@gmail.com, {spavaniitm, roshan, niravbhatt}@smail.iitm.ac.in

Abstract

Hard negatives are essential for training effective retrieval models. Hard-negative mining typically relies on ranking documents using cross-encoders or static embedding models based on similarity metrics such as cosine distance. Hard negative mining becomes challenging for biomedical and scientific domains due to the difficulty in distinguishing between source and hard negative documents. However, referenced documents naturally share contextual relevance with the source document but are not duplicates, making them well-suited as hard negatives. In this work, we propose BiCA: Biomedical Dense Retrieval with Citation-Aware Hard Negatives, an approach for hard-negative mining by utilizing citation links in 20,000 PubMed articles for improving a domain-specific small dense retriever. We fine-tune the GTE_{small} and GTE_{Base} models using these citation-informed negatives and observe consistent improvements in zero-shot dense retrieval using nDCG@10 for both in-domain and out-of-domain tasks on BEIR and outperform baselines on long-tailed topics in LoTTE using Success@5. Our findings highlight the potential of leveraging document link structure to generate highly informative negatives, enabling state-of-the-art performance with minimal fine-tuning and demonstrating a path towards highly data-efficient domain adaptation.

Code — <https://github.com/bisect-group/BiCA>

Datasets —

<https://huggingface.co/collections/bisectgroup/bica>

Extended Version — <https://arxiv.org/abs/2511.08029>

Introduction

Information Retrieval (IR) is a fundamental discipline focused on extracting relevant information from vast collections of unstructured data, primarily text. IR systems employ various algorithms to match user queries with pertinent documents, integrating both exact lexical matching and semantic understanding techniques. These systems are essential for

*Worked done as a UG student and currently at the University of Copenhagen, Denmark

†These authors contributed equally.

‡Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

search engines, digital libraries, and question-answering applications, enabling users to efficiently navigate large volumes of information (Manning 2009).

Despite these advances, retrieving precise information from the rapidly expanding biomedical literature indexed in PubMed (Sayers et al. 2011) remains a significant challenge. This difficulty is often compounded by the prevalence of low-quality, keyword-based queries, which may lack the specificity required to pinpoint relevant documents within such a specialized and nuanced domain. To address this issue, we propose an effective alternative by taking advantage of advanced training strategies and model architectures tailored for this complex environment.

One such strategy is Hard Negative mining, which involves selecting challenging examples that closely resemble positive samples yet are ultimately irrelevant (Allan et al. 2003; Yang et al. 2024). By compelling models to learn finer-grained distinctions between these difficult-to-distinguish negatives and true positives, the resulting systems exhibit more accurate rankings and improved retrieval effectiveness. Specifically for biomedical IR, the challenge lies not only in the sheer volume of literature but also in the intricate terminology and the subtle semantic relationships between concepts.

In this work, we introduce BiCA (Biomedical Citation Aware) retrievers, a family of models designed to enhance biomedical information retrieval and out of domain retrieval. We propose a novel hard negative mining technique based on multi-hop citation chains within the PubMed database. This approach, combined with efficient model architectures, allows us to develop systems that are not only highly effective but also computationally efficient. We demonstrate that our models, BiCA_{Base} and the significantly smaller BiCA_{small}, achieve state of the art or highly competitive results on several biomedical and general domain IR benchmarks, often outperforming models that are substantially larger.

Our Contributions

The main contributions of this work are as follows:

- We introduce a novel hard negative mining strategy that constructs multi hop citation chains from PubMed, using the pubmed parser, to generate high quality, challenging negative examples for training retrieval models for

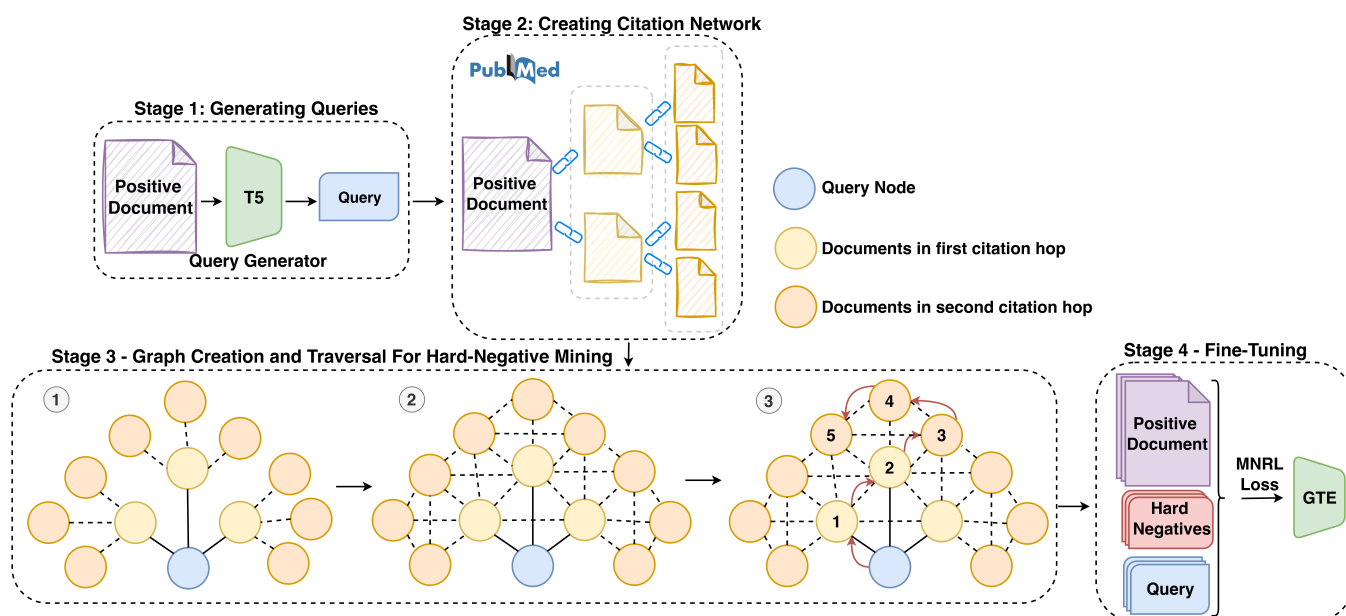


Figure 1: Our four-stage data generation and training pipeline. **Stage 1:** A query is synthetically generated from a positive document’s abstract using a T5 model. **Stage 2:** A 2-hop citation neighborhood is constructed by retrieving papers cited by the positive document (1-hop) and papers cited by them (2-hop) via the PubMed API. **Stage 3:** Hard negatives are mined via semantic graph traversal. First, similarities are computed between the query and 1-hop documents. Second, a dense, pairwise similarity graph is built for all 1-hop and 2-hop documents. Third, a 5-step greedy traversal is initiated from the 1-hop document most similar to the query, creating a path of five hard negatives. **Stage 4:** The resulting (Query, Positive Document, Hard Negatives) triplet is used to fine-tune the GTE model using the multiple negative ranking loss.

biomedical domains.

- We introduce BiCA, two dense retrieval models (base and small) specifically tailored for the biomedical domain using the proposed citation aware hard negative mining, which also demonstrate strong performance on general domain retrieval tasks.
- Extensive zero shot evaluations of our BiCA models on 14 BEIR tasks and 4 LoTTE tasks, outperforming all baselines on several tasks in BEIR and all sub topics on LoTTE.
- We provide a detailed latency analysis, demonstrating the practical efficiency of our models, particularly BiCA_{small}, on a single V100 GPU, highlighting their suitability for real world deployment.

Related Work

Biomedical Information Retrieval

Recent advances in biomedical IR have focused on the integration of novel methods and the use of large-scale data to improve retrieval performance. One such approach is Bibliometric Data Fusion (Breuer et al. 2023), which incorporates bibliometric metadata such as citation counts and altmetrics into retrieval systems. Using these implicit relevance signals, this method aims to improve retrieval performance, particularly for patient users, without relying on explicit relevance labeling.

A more recent development, Self-Learning Hypothetical Document Embeddings (SL-HyDE) (Li et al. 2024), introduces a zero-shot approach to medical IR by utilizing large language models (LLMs) to generate hypothetical documents based on a given query. This framework, which self-learns both pseudo-document generation and retrieval processes, improves retrieval accuracy without needing labeled data. The approach has shown notable performance across various LLM and retriever configurations, indicating its potential for enhancing zero-shot retrieval tasks.

Another important contribution to biomedical IR is the development of Neural Retrievers (NRs) (Luo et al. 2022), which address data scarcity in the biomedical domain. By proposing a template-based question generation method and introducing pre-training tasks aligned with the downstream retrieval task, NRs have made substantial strides. The “Poly-DPR” model, which encodes each context into multiple vectors, has been particularly effective, outperforming traditional methods like BM25 in certain retrieval settings.

Finally, MedCPT (Jin et al. 2023) employs contrastive pre-training to enhance zero-shot retrieval for biomedical information. Leveraging a large collection of user click logs from PubMed, MedCPT utilizes contrastive learning to train an integrated retriever and re-ranker model. This methodology has set new state-of-the-art benchmarks, outperforming several Baselines, including larger models like GPT-3-sized cpt-text-XL.

Biomedical Language Models

The development of domain-specific language models has addressed the unique challenges posed by biomedical text. Models like SciFive (Phan et al. 2021), BioMegatron (Shin et al. 2020), and PubMedBERT (Gu et al. 2021) have been trained on extensive biomedical corpora, enabling them to better understand specialized language and concepts. Additionally, other models such as BioBERT (Lee et al. 2019), PMC-LLaMA (Wu et al. 2024), ELECTRAMed (Miolo, Mantoan, and Orsenigo 2021), BioBART (Yuan et al. 2022), and BioMedLM (Bolton et al. 2024) have significantly advanced biomedical text mining and natural language processing (NLP).

Recent advancements in biomedical language modeling have explored graph-based approaches to represent biomedical literature as knowledge graphs, effectively capturing complex relationships among entities and concepts. These knowledge graphs enhance accuracy by providing a structured framework that reflects the intricate interconnections inherent in biomedical data. Works of (Saxena, Tripathi, and Talukdar 2020), (Yasunaga, Leskovec, and Liang 2022), and (Yasunaga et al. 2022) show significant improvements in question-answering systems and biomedical text understanding using knowledge graphs and multi-hop frameworks.

Methodology

First, we construct a rich, 2-hop citation neighborhood around a set of seed documents. Second, we perform a novel hard-negative mining technique by converting these citation graphs into dense semantic graphs and performing a series of diverse, stochastic traversals to find documents that are semantically close but not directly relevant. We provide an overview of our entire pipeline in Figure 1.

Data Curation: 2-Hop Citation Neighborhood Construction

The foundation of our dataset is a large-scale, localized citation graph. The process begins with a seed collection of PubMed abstracts from the uiyunkim-hub/pubmed-abstract dataset on Hugging Face¹. Our goal was to generate a final dataset of approximately 20,000 query-positive pairs, each with a corresponding set of high-quality hard negatives. We have ensured that the selected corpus of 20,000 documents is a fair representation of the much larger PubMed database (Refer extended version for details).

To create a candidate pool for these negatives, we performed the following steps for each seed article, which we designate as the “positive” document (P_0):

- **1-Hop Citation Retrieval:** Using the PubMed Identifier (PMID) of P_0 , we employed the pubmed-parser library to query the NCBI E-utilities API and retrieve a list of all PMIDs cited by P_0 . These form the 1-hop neighborhood (C_1). We then fetched the abstract for each paper in C_1 .
- **2-Hop Citation Retrieval:** For each paper $P_1 \in C_1$, we repeated the process, fetching the PMIDs of all papers it

cites. This collection of PMIDs forms the 2-hop neighborhood (C_2). We then fetched the abstract for each paper in C_2 .

- **Data Aggregation:** The final curated data structure for each positive document P_0 consists of its own abstract, a list of all 1-hop abstracts, and a list of all 2-hop abstracts. To ensure a sufficiently rich neighborhood for mining, we only retained records where abstracts could be successfully retrieved.

This data collection was heavily parallelized across 80 worker processes to manage the high volume of API calls to NCBI. The result is a JSONL file containing 20,000 complex objects, each representing a positive document and its extensive 2-hop citation context.

Hard-Negative Mining via Diverse Semantic Traversal

With the 2-hop citation neighborhoods established, we proceeded to the core of our hard-negative mining strategy. To enhance diversity and prevent the model from overfitting to a single type of negative, our approach, detailed in Algorithm 1, transforms the structural citation graph into a semantic space and explores it using multiple, stochastic paths.

The mining process unfolds in three steps for each of the 20,000 curated data points:

- **Query Generation:** We first generate a synthetic query from the positive abstract (A_{positive}) using the Doc2Query (doc2query/all-t5-base-v1) model (Nogueira et al. 2019). This creates a realistic search query that the positive document is expected to be relevant for.
- **Dense Graph Construction:** We then construct a dense, semantically weighted graph. All abstracts from the 1-hop and 2-hop neighborhoods are encoded into high-dimensional vectors using the Pubmedbert-base embeddings (NeuML 2025). We compute a complete pairwise cosine similarity matrix between all documents in the 1-hop and 2-hop pools.
- **Diverse Semantic Traversal:** With the dense graph constructed, we identify a varied set of hard negatives. The process is designed to be robust and avoid overfitting:
 - **Multiple Start Points:** Instead of one, we initiate three separate traversal paths, starting from the three 1-hop documents most semantically similar to the generated query.
 - **Stochastic Path Selection:** At each step of a traversal, rather than taking a purely greedy step to the single most similar node, we perform weighted random sampling from the top five most similar, unvisited nodes. This stochasticity ensures a wider exploration of the semantic space.
 - **Global Visited Set:** A single global set of visited nodes is maintained across all traversals for a given query, guaranteeing that each path explores unique documents and maximizing the diversity of the final negative set.
 - **Random Negative Augmentation:** Finally, to further improve training stability, one additional negative is

¹huggingface.co/datasets/uiyunkim-hub/pubmed-abstract

Model	Size	COVID	NFC	SCIFACT	SCIDOCs	QUORA	ArguAna	Climate-Fever	NQ	CQADup	DBPedia	Touche-2020	HotpotQA	FEVER	FiQA	Avg.	Macro Avg.
TAS-B	66M	0.481	0.319	0.643	0.149	0.835	0.434	0.221	0.463	0.315	0.384	0.162	0.584	0.700	0.300	0.399	0.399
R-GPL	66M	0.760	0.342	0.678	0.162	0.808	0.464	0.231	0.504	0.348	0.381	0.264	0.567	0.791	0.336	0.474	0.474
GPL	66*5M	0.700	0.345	0.674	0.169	0.832	0.483	0.227	0.467	0.345	0.360	0.266	0.636	0.758	0.344	0.472	0.472
DPR	110M	0.332	0.189	0.318	0.077	0.248	0.175	0.148	0.474	0.153	0.263	0.131	0.456	0.562	0.112	0.274	0.274
ANCE	110M	0.650	0.230	0.507	0.122	0.852	0.415	0.198	0.446	0.296	0.281	0.240	0.584	0.669	0.295	0.414	0.414
Contriever	110M	0.596	0.328	0.677	0.165	0.865	0.446	0.237	0.495	0.284	0.413	0.230	0.638	0.758	0.329	0.463	0.463
ColBERT	110M	0.677	0.305	0.671	0.145	0.854	0.233	0.184	0.524	0.350	0.392	0.202	0.593	0.771	0.317	0.445	0.445
ColBERTv2	110M	0.738	0.338	0.693	0.154	0.852	0.463	0.176	0.562	0.359	<u>0.446</u>	0.278	0.667	0.785	0.356	0.490	0.490
LexMAE	110M	<u>0.763</u>	0.347	0.710	0.159	-	0.500	0.219	0.562	-	0.424	<u>0.290</u>	0.716	<u>0.800</u>	0.352	-	0.487
DRAGON+	110M	0.759	0.339	0.679	0.159	0.875	0.469	0.227	0.537	0.354	0.414	0.263	0.662	0.781	0.359	0.491	0.491
SpladeV3	110M	0.748	<u>0.357</u>	0.710	0.158	0.814	0.509	0.233	0.586	-	0.450	0.293	<u>0.692</u>	0.796	0.374	-	<u>0.517</u>
SpladeV2	110M	0.710	0.334	0.693	0.158	0.838	0.479	0.235	0.521	0.341	0.435	0.272	0.684	0.786	0.336	0.487	0.487
RetroMae	110M	0.772	0.308	0.653	0.133	0.847	0.433	0.232	0.518	0.297	0.356	0.219	0.635	0.774	0.325	0.464	0.464
GenQ	220M	0.610	0.310	0.644	0.143	0.830	0.493	0.175	0.358	0.347	0.328	0.182	0.534	0.669	0.308	0.424	0.424
GTR _{Base}	110M	0.539	0.308	0.600	0.149	0.881	0.511	0.241	0.495	0.357	0.347	0.205	0.535	0.660	0.349	0.441	0.441
GTR-Large	335M	0.557	0.329	0.639	0.158	<u>0.890</u>	0.525	0.262	0.547	0.384	0.391	0.219	0.579	0.712	0.424	0.473	0.473
GTRxl	1.2B	0.580	0.343	0.635	0.159	<u>0.890</u>	0.531	<u>0.270</u>	0.559	0.388	0.396	0.230	0.591	0.717	<u>0.444</u>	0.481	0.481
GTRxxl	4.8B	0.500	0.342	0.662	0.161	0.892	0.540	0.267	<u>0.568</u>	<u>0.399</u>	0.408	0.256	0.599	0.740	0.467	0.486	0.486
BiCA_{small}	33M	0.661	0.347	0.727	0.214	0.880	0.555	0.264	0.502	0.399	0.391	0.222	0.637	0.815	0.393	<u>0.501</u>	0.501
BiCA_{Base}	110M	0.684	0.378	0.762	0.231	0.882	0.571	0.279	0.529	0.428	0.411	0.220	0.657	0.815	0.407	0.518	0.518

Table 1: Evaluation on all 14 BEIR tasks in a zero-shot setting using nDCG@10. **Bold** and underline denote the best and second-best scores, respectively.

selected uniformly at random from the remaining pool of unvisited documents.

The final output is a dataset of approximately 20,000 entries, each containing a query, a single positive abstract, and a diverse list of hard negatives (averaging 6.5 per query). This results in a total corpus of approximately 150,000 documents, specifically curated to train and evaluate retrieval models on their ability to make fine-grained relevance distinctions.

Experiments

Fine-Tuning

We fine tune two models the GTE_{small} and the GTE_{Base} (Li et al. 2023). GTE_{base} (110M parameters, 768-dim) and GTE_{small} (33M parameters, 384-dim) are BERT-based embedding models trained with multi-stage contrastive learning, balancing accuracy with efficiency. We describe our choice of fine-tuning for only 20 steps in the extended version.

The fine-tuning was conducted on a single NVIDIA V100 GPU (32GB), enabling efficient handling of large batch sizes and complex models without memory constraints. The Multiple Negative Ranking Loss (MNR) function (Henderson et al. 2017) is used and defined as:

$$\mathcal{L}_{MNR} = -\log \left(\frac{\exp(\mathbf{q} \cdot \mathbf{d}_+)}{\exp(\mathbf{q} \cdot \mathbf{d}_+) + \sum_{i=1}^K \exp(\mathbf{q} \cdot \mathbf{d}_i^-)} \right)$$

where \mathbf{q} denotes the query embedding, \mathbf{d}_+ the positive doc-

ument embedding, \mathbf{d}_i^- the i -th negative document embedding, and K the number of negatives.

Evaluation

BEIR We evaluate our models on fourteen BEIR (Thakur et al. 2021) datasets in a zero-shot setting. Our primary evaluation metric is Normalized Discounted Cumulative Gain at 10 (nDCG@10), which assesses the ranking quality of the top 10 retrieved documents. The comprehensive results, comparing our models against a wide range of existing methods, are presented in Table 1. We also provide the improvements over the base GTE models in the extended version. As shown in Table 1, our $BiCA_{Base}$ model (110M parameters) achieves the highest average nDCG@10 score of 0.518 across all fourteen tasks, setting a new state-of-the-art on BEIR and surpassing significantly larger models such as GTR_{xxl} (4.8B parameters, 0.486). $BiCA_{Base}$ excels in both biomedical and general domains, leading on NFCORPUS (0.378), SCIFACT (0.762), SCIDOCs (0.231), ARGUANA (0.571), CLIMATE-FEVER (0.279), and CQADUP (0.428), while tying for the highest on FEVER (0.815) and performing strongly on HOTPOTQA (0.657). Our smaller $BiCA_{small}$ model (33M parameters) also demonstrates remarkable performance, achieving an average nDCG@10 of 0.501, ranking second overall and outperforming many larger baselines, including GTR_{xxl} . Notably, it secures the top score on FEVER (0.815) and second-highest on SCIDOCs (0.214), ARGUANA (0.555), and CQADUP (0.399). Its ability to rival or surpass models up to 145 times larger highlights the

Corpus	ColBERT	BM25	ANCE	RocketQAv2	SPLADEv2	ColBERTv2	BiCA _{small}	BiCA _{Base}
LoTTE Search Test Queries (Success@5)								
Writing	74.7	60.3	74.4	78.0	77.1	<u>80.1</u>	79.8	81.6
Recreation	68.5	56.5	64.7	72.1	69.0	72.3	<u>76.1</u>	79.7
Science	53.6	32.7	53.6	55.3	55.4	56.7	<u>58.5</u>	60.6
Lifestyle	80.2	63.8	82.3	82.1	82.3	84.7	<u>86.8</u>	87.7
LoTTE Forum Test Queries (Success@5)								
Writing	71.0	64.0	68.8	71.5	73.0	76.3	<u>78.1</u>	80.8
Recreation	65.6	55.4	63.8	65.7	67.1	70.8	<u>75.6</u>	77.5
Science	41.8	37.1	36.5	38.0	43.7	<u>46.1</u>	44.6	47.1
Lifestyle	73.0	60.6	73.1	73.7	74.0	76.9	<u>82.2</u>	84.0

Table 2: Retrieval performance (Success@5) of different models on LoTTE search and forum queries. **Bold** is best, underline is second best.

parameter efficiency of our approach.

LOTTE We evaluate on long tailed topics, which refer to specific and less frequently searched queries, using four subtopics from the LoTTE benchmark (Santhanam et al. 2022): Science, Writing, Recreation, and Lifestyle. As detailed in Table 2, we report zero shot Success@5 on its test set. The benchmark includes two query formats: concise Search queries from GooAQ logs and more descriptive Forum queries from StackExchange user questions. Our BiCA_{Base} model sets a new state-of-the-art, achieving the highest Success@5 across all four categories for both LoTTE query types. On Search queries, it scores 87.7 on Lifestyle, 81.6 on Writing, 79.7 on Recreation, and 60.6 on Science. On the more challenging Forum queries, it attains 84.0 on Lifestyle, 80.8 on Writing, 77.5 on Recreation, and 47.1 on Science. The smaller BiCA_{small} model consistently ranks second, with Search scores of 86.8 on Lifestyle, 76.1 on Recreation, and 58.5 on Science, and Forum scores of 82.2 on Lifestyle, 78.1 on Writing, and 75.6 on Recreation, demonstrating strong performance and parameter efficiency on long-tailed topics.

Latency To assess model efficiency, we measured latency using the TAS-B setup on a single NVIDIA V100 with 32GB memory. We encoded 10,000 MS MARCO passages and indexed them with FAISS (IndexFlatIP). We then timed two steps: query encoding and retrieval of the top 1000 results. Tests were run on query batches of size 1, 10, and 2000. We report average and 99th percentile latencies in milliseconds over 100 iterations (1 and 10) or 10 iterations (2000).

Table 3 compares BiCA_{Base} (110M), BiCA_{small} (33M), ColBERTv2, RetroMAE, and SpladeV3. For batch size 1, BiCA_{small} is fastest overall with 13 ms total and 4 ms retrieval. ColBERTv2 has the quickest encoding at 8 ms and a total of 15 ms. The others average 16 ms, with BiCA_{Base} showing slightly higher tail times.

At batch size 10, BiCA_{small} again leads in total time (19 ms), driven by retrieval at 5 ms. ColBERTv2, RetroMAE, and SpladeV3 encode slightly faster (11 ms vs 14 ms for BiCA_{small}). ColBERTv2 has the best tail latency at 23 ms, while SpladeV3 peaks at 32 ms.

At batch size 2000, BiCA_{small} outperforms all others with

994 ms total (554 ms encoding, 441 ms retrieval). RetroMAE follows at 1837 ms, then ColBERTv2 (1844 ms) and SpladeV3 (1847 ms). BiCA_{Base} is slowest at 1904 ms.

Effect of Traversal Parameters

To determine appropriate values for the traversal parameters, we conduct an ablation study varying the *Number of Traversal Paths* (N_{paths}) and the *Length of the Path* (L_{path}) in the range of 1–5. For each study, we fix one parameter at 3 while varying the other, using a BERT-base fine-tuned for 1 epoch on the entire corpus with a batch size of 16 and an MNR loss. As shown in Table 4, the choice of $N_{paths} = 3$ and $L_{path} = 3$ consistently provides a strong balance across datasets, achieving the highest overall average performance (0.2739). While other configurations occasionally yield the best score on a single dataset (e.g., $N_{paths} = 5$ for SCIFACT or $L_{path} = 1$ for ArguAna), they underperform on others, leading to a lower overall average. We therefore select $N_{paths} = 3$ and $L_{path} = 3$ as the default configuration for our final results, as it offers the most stable and robust performance across benchmarks.

Robustness and Scalability

To examine the effect of training data size, we fine-tuned bert-base-uncased (Devlin et al. 2019) on randomly sampled subsets of our 20,000-record dataset (1k, 5k, 10k, 15k, and full). Each subset reserved 10% for validation. Models were trained for up to 1 epoch using MNR Loss with a batch size of 16, applying early stopping based on the highest triplet accuracy on validation. The best checkpoints were evaluated zero-shot on three biomedical tasks and one BEIR task. Results in Table 5 show a clear positive correlation between data size and retrieval performance.

Performance of Different Architectures

To assess generalizability, we fine-tune models for a maximum of one epoch with early stopping (patience=3), where evaluation is performed every 10 steps. We experiment with two pretrained checkpoints: e5-base-V2² (Wang et al. 2022)

²<https://huggingface.co/intfloat/e5-base-v2>

Model	Batch Size	Enc. Avg. (ms)↓	Enc. 99th (ms)↓	Ret. Avg. (ms)↓	Ret. 99th (ms)↓	Total Avg. (ms)↓	Total 99th (ms)↓
BiCA _{Base}	1	9	14	7	9	16	21
	10	11	16	9	10	20	25
	2000	1292	1475	612	622	1904	2082
BiCA _{small}	1	9	11	4	4	13	14
	10	14	19	5	5	19	24
	2000	554	850	441	504	994	1341
ColBERTv2	1	8	9	7	7	15	16
	10	11	13	9	10	20	23
	2000	1249	1423	594	612	1844	2004
RetroMAE	1	9	11	7	8	16	20
	10	11	13	9	12	20	25
	2000	1246	1403	591	607	1837	1985
SpladeV3	1	9	11	7	9	16	19
	10	11	15	9	13	21	32
	2000	1250	1437	598	609	1847	2045

Table 3: Latency analysis for BiCA_{Base}, BiCA_{small}, and other baselines on a V100 (32GB) GPU. Cell colors highlight timings from lowest (lightest orange) to highest (darkest orange) for each metric across models within the same batch size. All times in milliseconds (ms).

N_{paths}	L_{path}	NFC	SCIDOCs	SCIFACT	ArguAna	FIQA	Average
<i>Ablation on Number of Traversals (fixed $L_{path} = 3$)</i>							
1	3	0.1803	0.1201	0.5114	0.3974	0.1301	0.2679
2	3	0.1390	0.0984	0.3934	0.3174	0.0860	0.2068
4	3	0.1400	0.1073	0.4392	0.3024	0.1030	0.2184
5	3	0.1891	0.1230	0.5180	0.4190	0.1178	0.2734
<i>Ablation on Path Length (fixed $N_{paths} = 3$)</i>							
3	1	0.1875	0.1245	0.5053	0.4211	0.1240	0.2725
3	2	0.1299	0.0965	0.3960	0.2920	0.1062	0.2041
3	3	0.1987	0.1234	0.5156	0.4094	0.1225	0.2739
3	4	0.1861	0.1202	0.5102	0.3854	0.1324	0.2669
3	5	0.1820	0.1183	0.5058	0.3730	0.1110	0.2580

Table 4: Ablation study on the number of traversals (N_{paths}) and path length (L_{path}). All models are based on BERT-base fine-tuned for one epoch. We report NDCG@10 scores and highlight the best result in each column in **bold**. The yellow row indicates the true parameters used in BiCA.

Dataset	Baseline	1k	5k	10k	15k	Full (20k)
NFCorpus	0.043	0.082	0.171	0.164	0.171	0.185
SciDocs	0.028	0.061	0.117	0.116	0.114	0.121
SciFact	0.130	0.262	0.469	0.468	0.492	0.493
ArguAna	0.283	0.384	0.364	0.385	0.405	0.444

Table 5: Scaling ablation results for fine tuning bert base uncased on our citation aware negatives. Scores are nDCG@10 on biomedical BEIR tasks. The baseline represents zero shot performance without any fine tuning. The results show consistent performance improvement as the amount of training data increases.

and a DistilBERT model³ (Sanh et al. 2020) fine-tuned on MS MARCO. For evaluation, we select five tasks from the BEIR benchmark: three from the biomedical domain (NF-

³<https://huggingface.co/GPL/msmarco-distilbert-margin-mse>

Dataset	DB _{base}	DB _{fine tune}	E5 _{base}	E5 _{fine tune}
NFCorpus	24.8	25.2 (+0.4)	35.3	34.8 (-0.5)
SciFact	51.6	55.8 (+4.2)	71.0	71.9 (+0.9)
SCIDOCs	13.4	14.9 (+1.5)	18.3	20.4 (+2.1)
ArguAna	39.7	39.9 (+0.2)	51.6	52.7 (+1.1)
FiQA	18.2	19.7 (+1.5)	37.3	37.9 (+0.6)
Average Δ	-	+1.56	-	+0.84

Table 6: NDCG@10 (%) comparison between DistilBERT (DB) and E5 models across BEIR datasets. Parentheses show the absolute improvement of fine-tuned models over base versions.

Corpus, SciDocs, SciFact) and two from non-biomedical domains (ArguAna, FiQA). Table 6 shows the performance gains of fine-tuning the models on our corpus, and Table 7 shows the number of fine-tuning steps selected for the

Algorithm 1: Hard Negative Mining via Diverse Semantic Traversal

Require:

A_{pos} : Abstract of the positive document.
 A_{cands} : Set of candidate abstracts from citation hops.
 $N_{\text{paths}}, L_{\text{path}}, K_{\text{sample}}$: Traversal control parameters.

Ensure:

L_{negs} : A diverse list of hard negative abstracts.

```
1: procedure MINEHARDNEGATIVES( $A_{\text{pos}}, A_{\text{cands}}$ )
2:    $\triangleright$  Step 1: Construct a semantic graph of documents.
3:    $Q \leftarrow \text{GENERATEQUERY}(A_{\text{pos}})$ 
4:    $S_{\text{graph}} \leftarrow \text{BUILDSIMILARITYGRAPH}(A_{\text{cands}})$ 
5:    $\triangleright$  Step 2: Initiate N traversals from query-relevant
   starts.
6:    $I_{\text{start}} \leftarrow \text{FINDTOPNSTARTS}(Q, A_{\text{cands}}, N_{\text{paths}})$ 
7:    $L_{\text{negs}} \leftarrow \emptyset, V_{\text{visited}} \leftarrow \emptyset$ 
8:    $\triangleright$  Step 3: Perform stochastic walks to find diverse
   negatives.
9:   for each  $i_{\text{start}}$  in  $I_{\text{start}}$  do
10:     $i_{\text{curr}} \leftarrow i_{\text{start}}$ 
11:    for  $l \leftarrow 1$  to  $L_{\text{path}}$  do
12:      if  $i_{\text{curr}} \in V_{\text{visited}}$  then break
13:      Add  $A_{\text{cands}}[i_{\text{curr}}]$  to  $L_{\text{negs}}$  and  $V_{\text{visited}}$ 
14:       $\triangleright$  Select next node: top-K unvisited
   neighbors,
    $\triangleright$  sampled probabilistically by similarity.
15:       $I_{\text{topK}} \leftarrow \text{GETTOPKUNVISITEDNEIGHBORS}$ 
    $\{i_{\text{curr}}, S_{\text{graph}}, K_{\text{sample}}, V_{\text{visited}}\}$ 
16:       $i_{\text{curr}} \leftarrow \text{SAMPLEPROBABILISTICALLY}$ 
    $\{I_{\text{topK}}, S_{\text{graph}}[i_{\text{curr}}, I_{\text{topK}}]\}$ 
17:       $\triangleright$  Step 4: Add a random negative for robustness.
18:      Add one random, unvisited abstract to  $L_{\text{negs}}$ .
19:   return Unique( $L_{\text{negs}}$ )
```

Model	No. Fine-Tuning Steps
DistilBERT	1150
e5-base-v2	290

Table 7: Number of fine-tuning steps on our constructed corpus before doing zero-shot evaluation on BEIR

chosen models, after which we do zero-shot evaluation on BEIR. We see consistent improvements in using our corpus for fine-tuning over different architectures. DistilBERT sees an average improvement of 1.56 points, and e5-base-v2 sees an improvement of 0.84 points.

Conclusions

In this work, we present BiCA_{Base} and BiCA_{small}, two dense retrieval models designed to address the unique challenges of biomedical and general-domain information retrieval. At the core of our approach is a novel hard negative mining strategy that exploits multi-hop citation chains extracted from PubMed. This citation-aware technique provides semantically challenging yet relevant negative examples, encouraging the models to learn fine-grained distinctions es-

sential for high-precision retrieval. Through extensive experiments on the BEIR benchmark, BiCA (110M) demonstrated strong performance across both biomedical and non-biomedical tasks, consistently outperforming several larger state of the art models. Notably, it achieved the highest average nDCG@10 scores in both domains, indicating its effectiveness and generalizability. Despite its smaller size, BiCA (33M) also delivered competitive results, often closely trailing BiCA (110M) while offering substantially lower inference latency, making it well suited for real time and resource constrained applications. Evaluations on the LoTTE dataset further highlighted the robustness of our models in handling retrieval over long-tailed, diverse topics. BiCA_{Base} led across all sub-domains, while BiCA_{small} ranked consistently among the top performers, demonstrating the broad applicability and efficiency of our approach.

Limitations

The citation-aware hard negative mining strategy improves retrieval performance, but faces challenges in scalability and efficiency. Constructing multi-hop citation chains requires iterative PubMed API requests for abstracts and cited PMIDs, a process hindered by rate limits, network latency, and the parsing of large text data. As a result, generating large training sets can take week(s), depending on the number of seed documents and citation depth. Furthermore, our current work is restricted to PubMed; extending this approach to other sources such as Wikipedia, where scientific and technical articles contain rich citation trails, may enable the construction of semantically meaningful hard negatives for general-domain retrieval while preserving citation-aware principles. We acknowledge that our latency evaluation setup may not fully reflect the efficiency advantages of the ColBERTv2 model. However, we adopted this configuration to ensure a uniform and fair comparison across all systems.

References

- Allan, J.; Aslam, J.; Belkin, N.; Buckley, C.; Callan, J.; Croft, B.; Dumais, S.; Fuhr, N.; Harman, D.; Harper, D. J.; Hiemstra, D.; Hofmann, T.; Hovy, E.; Kraaij, W.; Lafferty, J.; Lavrenko, V.; Lewis, D.; Liddy, L.; Manmatha, R.; McCallum, A.; Ponte, J.; Prager, J.; Radev, D.; Resnik, P.; Robertson, S.; Rosenfeld, R.; Roukos, S.; Sanderson, M.; Schwartz, R.; Singhal, A.; Smeaton, A.; Turtle, H.; Voorhees, E.; Weischedel, R.; Xu, J.; and Zhai, C. 2003. Challenges in information retrieval and language modeling: report of a workshop held at the center for intelligent information retrieval, University of Massachusetts Amherst, September 2002. *SIGIR Forum*, 37(1): 31–47.
- Bolton, E.; Venigalla, A.; Yasunaga, M.; Hall, D.; Xiong, B.; Lee, T.; Daneshjou, R.; Frankle, J.; Liang, P.; Carbin, M.; and Manning, C. D. 2024. BioMedLM: A 2.7B Parameter Language Model Trained On Biomedical Text. arXiv:2403.18421.
- Breuer, T.; Kreutz, C. K.; Schaer, P.; and Tunger, D. 2023. Bibliometric Data Fusion for Biomedical Information Re-

- trieval. In *2023 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, 107–118.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J.; Doran, C.; and Solorio, T., eds., *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)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Gu, Y.; Tinn, R.; Cheng, H.; Lucas, M.; Usuyama, N.; Liu, X.; Naumann, T.; Gao, J.; and Poon, H. 2021. Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing. ArXiv:2007.15779.
- Henderson, M.; Al-Rfou, R.; Strope, B.; hsuan Sung, Y.; Lukacs, L.; Guo, R.; Kumar, S.; Miklos, B.; and Kurzweil, R. 2017. Efficient Natural Language Response Suggestion for Smart Reply. arXiv:1705.00652.
- Jin, Q.; Kim, W.; Chen, Q.; Comeau, D. C.; Yeganova, L.; Wilbur, W. J.; and Lu, Z. 2023. MedCPT: Contrastive Pre-trained Transformers with large-scale PubMed search logs for zero-shot biomedical information retrieval. *Bioinformatics*, 39(11): btad651.
- Lee, J.; Yoon, W.; Kim, S.; Kim, D.; Kim, S.; So, C. H.; and Kang, J. 2019. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4): 1234–1240.
- Li, L.; Zhang, X.; Zhou, X.; and Liu, Z. 2024. AutoMIR: Effective Zero-Shot Medical Information Retrieval without Relevance Labels. arXiv:2410.20050.
- Li, Z.; Zhang, X.; Zhang, Y.; Long, D.; Xie, P.; and Zhang, M. 2023. Towards General Text Embeddings with Multi-stage Contrastive Learning. arXiv:2308.03281.
- Luo, M.; Mitra, A.; Gokhale, T.; and Baral, C. 2022. Improving Biomedical Information Retrieval with Neural Retrievers. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10): 11038–11046.
- Manning, C. D. 2009. *An introduction to information retrieval*. Cambridge: Cambridge University Press.
- Miolo, G.; Mantoan, G.; and Orsenigo, C. 2021. ELEC-TRAMed: a new pre-trained language representation model for biomedical NLP. arXiv:2104.09585.
- NeuML. 2025. NeuML/pubmedbert-base-embeddings · Hugging Face. [Online; accessed 2025-07-31].
- Nogueira, R.; Yang, W.; Lin, J.; and Cho, K. 2019. Document Expansion by Query Prediction. ArXiv:1904.08375.
- Phan, L. N.; Anibal, J. T.; Tran, H.; Chanana, S.; Bahadroglu, E.; Peltekian, A.; and Altan-Bonnet, G. 2021. SciFive: a text-to-text transformer model for biomedical literature. ArXiv:2106.03598.
- Sanh, V.; Debut, L.; Chaumond, J.; and Wolf, T. 2020. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. ArXiv:1910.01108 version: 4.
- Santhanam, K.; Khattab, O.; Saad-Falcon, J.; Potts, C.; and Zaharia, M. 2022. ColBERTv2: Effective and Efficient Retrieval via Lightweight Late Interaction. In Carpuat, M.; de Marneffe, M.-C.; and Meza Ruiz, I. V., eds., *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3715–3734. Seattle, United States: Association for Computational Linguistics.
- Saxena, A.; Tripathi, A.; and Talukdar, P. 2020. Improving Multi-hop Question Answering over Knowledge Graphs using Knowledge Base Embeddings. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4498–4507. Online: Association for Computational Linguistics.
- Sayers, E. W.; Barrett, T.; Benson, D. A.; Bolton, E.; Bryant, S. H.; Canese, K.; Chetvermin, V.; Church, D. M.; DiCuccio, M.; Federhen, S.; Feolo, M.; Fingerman, I. M.; Geer, L. Y.; Helmberg, W.; Kapustin, Y.; Landsman, D.; Lipman, D. J.; Lu, Z.; Madden, T. L.; Madej, T.; Maglott, D. R.; Marchler-Bauer, A.; Miller, V.; Mizrahi, I.; Ostell, J.; Panchenko, A.; Phan, L.; Pruitt, K. D.; Schuler, G. D.; Sequeira, E.; Sherry, S. T.; Shumway, M.; Sirotkin, K.; Slotta, D.; Souvorov, A.; Starchenko, G.; Tatusova, T. A.; Wagner, L.; Wang, Y.; Wilbur, W. J.; Yaschenko, E.; and Ye, J. 2011. Database resources of the National Center for Biotechnology Information. *Nucleic Acids Research*, 39(suppl_1): D38–D51.
- Shin, H.-C.; Zhang, Y.; Bakhturina, E.; Puri, R.; Patwary, M.; Shoeybi, M.; and Mani, R. 2020. BioMegatron: Larger Biomedical Domain Language Model. ArXiv:2010.06060.
- Thakur, N.; Reimers, N.; Rücklé, A.; Srivastava, A.; and Gurevych, I. 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*.
- Wang, L.; Yang, N.; Huang, X.; Jiao, B.; Yang, L.; Jiang, D.; Majumder, R.; and Wei, F. 2022. Text Embeddings by Weakly-Supervised Contrastive Pre-training. *arXiv preprint arXiv:2212.03533*.
- Wu, C.; Lin, W.; Zhang, X.; Zhang, Y.; Xie, W.; and Wang, Y. 2024. PMC-LLaMA: toward building open-source language models for medicine. *Journal of the American Medical Informatics Association*, 31(9): 1833–1843.
- Yang, Z.; Shao, Z.; Dong, Y.; and Tang, J. 2024. TriSampler: A Better Negative Sampling Principle for Dense Retrieval. arXiv:2402.11855.
- Yasunaga, M.; Bosselut, A.; Ren, H.; Zhang, X.; Manning, C. D.; Liang, P.; and Leskovec, J. 2022. Deep Bidirectional Language-Knowledge Graph Pretraining. In *Neural Information Processing Systems (NeurIPS)*.
- Yasunaga, M.; Leskovec, J.; and Liang, P. 2022. LinkBERT: Pretraining Language Models with Document Links. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 8003–8016. Dublin, Ireland: Association for Computational Linguistics.
- Yuan, H.; Yuan, Z.; Gan, R.; Zhang, J.; Xie, Y.; and Yu, S. 2022. BioBART: Pretraining and Evaluation of A Biomedical Generative Language Model. In *Proceedings of the*

21st Workshop on Biomedical Language Processing, 97–109. Dublin, Ireland: Association for Computational Linguistics.