

QABISAR: Query-Article Bipartite Interactions for Statutory Article Retrieval

Santosh T.Y.S.S, Hassan Sarwat, Matthias Grabmair

School of Computation, Information, and Technology;

Technical University of Munich, Germany

{santosh.tokala, hassan.sarwat, matthias.grabmair}@tum.de

Abstract

In this paper, we introduce QABISAR, a novel framework for statutory article retrieval, to overcome the semantic mismatch problem when modeling each query-article pair in isolation, making it hard to learn representation that can effectively capture multi-faceted information. QABISAR leverages bipartite interactions between queries and articles to capture diverse aspects inherent in them. Further, we employ knowledge distillation to transfer enriched query representations from the graph network into the query bi-encoder, to capture the rich semantics present in the graph representations, despite absence of graph-based supervision for unseen queries during inference. Our experiments on a real-world expert-annotated dataset demonstrate its effectiveness.

1 Introduction

In an age where legal complexities are challenging for many individuals, there is a pressing need to bridge the gap between legal expertise and public understanding (Ponce et al., 2019). One critical step in this process is Statutory Article Retrieval (SAR), which involves identifying relevant statutes for a legal question. SAR plays a vital role in providing initial legal assistance by offering foundational insights into the law. Beyond statutes, previous works have explored retrieval of similar prior cases (Goebel et al., 2023; Santosh et al., 2024a) or pertinent information from case documents in response to legal queries (Santosh et al., 2024b). Leveraging advanced technologies for SAR can enhance the accuracy of legal insights, ultimately making legal knowledge more accessible and understandable to a broader audience.

Traditionally, SAR methods have been explored using the COLIEE Statute Law Corpus (Rabelo et al., 2021), containing questions linked to relevant articles from the Japanese Civil Code. However, these questions which are obtained from legal

bar exam yes or no questions, are quite different from those posed by ordinary citizens, often being vague and underspecified. To address this, Louis and Spanakis (2022) developed the Belgian Statutory Article Retrieval Dataset (BSARD), featuring french legal questions from Belgian citizens labeled by legal experts with references to relevant articles from Belgian legislation, which we use in our study.

Traditional SAR techniques included BM25, TF-IDF (Yoshioka et al., 2018), Indri (Strohman et al., 2005) and Word Movers’ Distance (Kusner et al., 2015). With the rise of pre-trained models, BERT and their ensembles have become popular (Kim et al., 2019; Rabelo et al., 2021, 2022). Recently, dense retrieval methods have gained attention (Louis and Spanakis, 2022) and were enhanced further through synthetic query generation and legal domain-oriented pre-training (Louis et al., 2023). Additionally, Louis et al. (2023) has demonstrated that articles are not completely independent of each other but an ensemble of interdependent rules organized into different codes, books, titles, chapters, and sections. They leveraged the hierarchical organization of statute law and utilized graph neural networks to enrich article representations by exploiting the interdependencies among articles within the topological structure. Orthogonal to these improvements, Santosh et al. (2024c) introduced a curriculum-based negative sampling strategy to make the model distinguish easier negatives in the initial stages of learning and progressively tackle more difficult ones.

Existing SAR works primarily focus on capturing the semantic relevance between individual query and article pairs in isolation. However, we argue that this approach may lead to sub-optimal representations, particularly in scenarios where both queries and articles contain multifaceted information. Articles, for instance, can cover a variety of semantics relevant to different queries, while

each query may require multiple relevant articles to comprehensively address its various aspects. Recognizing this inherent many-to-many relationship, our work takes a different approach. We leverage a query-article bipartite graph where nodes represent either a query or an article, and edges between them signify their relevance. We further augment this bipartite graph with hierarchical organization of statutes, with additional structural links such as sections, chapters, titles etc, to facilitate cross-article dependencies through their neighbourhood hops (Louis et al., 2023). We employ a Graph Attention Network on this augmented graph to aggregate information across the graph, allowing to capture the multiple interactions between queries and articles simultaneously, leading to effective capture of the diverse aspects inherent in each of them.

During inference, we face the challenge of utilizing only article representations from the graph network and fall back on query encoder to obtain query embeddings, as unseen queries at test time are absent in the constructed graph. To ensure that the query encoder representations are as expressive as the graph representations, we adopt knowledge distillation (KD) (Hinton et al., 2015) which aims to improve student models with the aid of teacher models which usually have same architecture with greater number of layers and dimensions (Wang et al., 2020). KD has been explored in IR tasks earlier, different from conventional setting, where they employ more expressive cross-encoder models as teachers and bi-encoders as students (Qu et al., 2021; Lu et al., 2022; Hofstätter et al., 2020; Choi et al., 2021). In this work, we use KD to facilitate representation transfer of queries from the graph network to the query encoder. By doing so, we aim to equip the query encoder with the ability to capture the rich semantics present in the graph representations, despite the absence of explicit graph-based supervision during inference.

We apply our approach, QABISAR, on the publicly available BSARD dataset (Louis and Spanakis, 2022), demonstrating the effectiveness of leveraging bipartite interactions between queries and articles, as well as knowledge distillation for representational transfer.

2 Our Method: QABISAR

Statutory Article Retrieval: Given a question q and corpus of statutes $P = \{p_1, p_2, \dots, p_m\}$, the task of SAR is to retrieve a smaller set of statutes

P_q ($|P_q| \ll |P|$) ranked in terms of their relevancy to answer the query. We mainly deal with optimizing the recall of the SAR system acting as pre-fetcher component, leaving the re-ranker component optimized for precision, for future.

QABISAR involves two stages of training. The first stage employs a dense bi-encoder which maps query and article into representations independently, while the second stage utilizes a graph encoder, designed to capture the bipartite interactions between queries and articles, enhancing retrieval by learning multi-faceted representations.

2.1 Dense bi-encoder

We use a dual-encoder architecture (Karpukhin et al., 2020) with query and statute encoder mapping each of them into a k -dimensional vector and the relevance score is computed using dot product between the encodings of query q and statute p_i as $f(q, p_i) = E_q(q) \cdot E_p(p_i)$ where E_q, E_p denote query and statute encoder. We use BERT-based model (Devlin et al., 2018) as query encoder to obtain query embedding from [CLS] representation. To account for longer length of articles, we use hierarchical article encoder (Pappagari et al., 2019) where the article is split into different chunks and each of them is independently encoded using a BERT-based model and then [CLS] representations from each chunk along with learnable position encodings are passed into a transformer which are then max-pooled to obtain the article embedding.

Dense bi-encoder module is trained with contrastive loss whose objective is to pull the representations of the query q and relevant articles S_q together (as positives), while pushing apart irrelevant ones $P'_q = P - P_q$ (as negatives). However, training with all the negatives is inscalable given larger corpus. To alleviate this issue, negative sampling has been employed where some irrelevant documents are sampled for each query during training making the final objective function as follows:

$$L(q, P_q, P'_q) = \sum_{p \in P_q} -\log \frac{\exp(f(q, p)/\tau)}{\sum_{c \in \{p\} \cup P'_q} \exp(f(q, c)/\tau)}$$

where hyperparameter τ is a scalar temperature. Following Karpukhin et al. (2020), we consider two types of negatives: (i) in-batch -articles paired with the other queries in the the same batch, and (ii) BM25- top articles returned by BM25 that are not relevant to the query.

2.2 Graph Encoder

To capture many-to-many interactions between query and articles effectively, we construct a query-article bipartite graph from training data and utilize graph attention network, to enrich the representations of both queries and articles through their multiple interactions simultaneously.

Graph Construction: We construct a bipartite graph utilizing all queries from the training set and all the articles from the corpus as nodes, establishing edges between queries and their corresponding relevant articles. Additionally, we augment our bipartite graph with the hierarchical organizational topology of statute structure by introducing additional nodes to represent sections, chapters, titles, and books and creating edges to denote the hierarchical connections between these structural units and the article nodes. This enables the article nodes to learn the complementary information from neighbouring elements as articles are not completely independent of each other but an ensemble of interdependent rules organized into different codes, books, titles, chapters, and sections (Louis et al., 2023). We also label edges based on the type of nodes they connect (i.e., Query-Article, Section-Article etc).

Node Initialization: We use the article encoder to obtain embeddings for structural units in the legislative topology graph such as section, article nodes and query encoder to obtain embeddings for query nodes based on their textual content.

Graph Attention Network (GAT): A graph neural network layer updates every node representation by aggregating its neighbors representations. GAT (Veličković et al., 2018; Brody et al., 2021) inject the graph structure into the attention mechanism by performing masked attention using neighborhood of a node. It employs a K multi-headed attention mechanism with residual connections and then their features are concatenated, resulting in the updated node feature representation as follows:

$$\mathbf{x}'_i = \parallel_{k=1}^K \sigma \left(\alpha_{i,j}^k \mathbf{W}_s^k \mathbf{x}_i + \sum_{j \in N(i)} \left(\alpha_{i,j}^k \mathbf{W}_t^k \mathbf{x}_j \right) \right)$$

where \parallel indicates concatenation, σ indicates non-linearity, W_s, W_t indicates learnable weight matrices. $\alpha_{i,j}^k$ denote attention weight computed using the node features x_i, x_j and edge connecting them $e_{i,j}$, which in our case indicate the node types.

$$\alpha_{i,j}^k = \text{Softmax} \left((\mathbf{a}^k)^T \text{LeakyReLU} \left(\left[W_s^k \mathbf{x}_i \parallel W_t^k \mathbf{x}_j \parallel W_e^k e_{i,j} \right] \right) \right)$$

To learn GAT parameters, we adopt the same contrastive learning used to train bi-encoder by obtaining the query and article representations from the graph nodes. We only sample sub-graph with L-hop neighbours (L denote number of GAT layers) based on the current batch of query, articles to save computational cost and pass them into GAT to extract node features for loss computation. These bi-partite interactions, lead to enriched representations for both the articles and queries representations capturing multiple aspects covered in them, However, during inference, we can not use query representations from graph as we encounter unseen queries that are not present in the constructed graph, resulting in an inductive learning setting on graphs.

We employ knowledge distillation (Hinton et al., 2015) to facilitate representation transfer for queries from the graph network to the query encoder. This involves distilling the relevance scores assessed with the query representations from the graph (acting as the teacher model) into the query bi-encoder (student model). Formally, given a query q and a list of candidate articles $P = \{p_1, p_2, \dots, p_m\}$, we obtain article representations from graph network p^g and query representations from query bi-encoder q^b and graph network q^g . We compute two relevance scores between query and article representations and convert them into probability distributions of the scores over candidate articles. We apply Knowledge distillation to mimic the scoring distribution of a more expressive graph network based relevance scores $s(q^g, p^g)$ with the distribution of a bi-encoder $s(q^b, p^g)$ measured by KL divergence as follows

$$L_{KD} = \sum_{q \in Q, p \in P} s(q^g, p^g) \cdot \log \frac{s(q^g, p^g)}{s(q^b, p^g)}$$

$$s(q^g, p^g) = \frac{e^{f(q^g, p^g)}}{\sum_{p' \in P} e^{f(q^g, p'^g)}}$$

Our final loss for the second-stage training is the sum of contrastive loss using graph representations and the knowledge distillation loss. This joint training drives the bi-encoder to update in tandem with the graph representations, which are also initialized with bi-encoder representations during the start.

3 Experiments

3.1 Dataset & Baselines

We use BSARD (Louis and Spanakis, 2022) containing 1108 french legal questions, with references

Method	R@			MAP	MRP
	100	200	500		
BM25	49.3	57.3	63.0	16.8	13.6
BE w/o Hier.	77.9	81.8	88.1	36.4	30.1
BE	79.9	83.3	88.7	39.2	31.6
BE+GE-Stat.	82.3	85.1	89.9	42.6	34.8
QABISAR	83.7	87.9	91.3	43.1	35.6

Table 1: Comparison of QABISAR with prior works.

to relevant articles from a corpus of 22,600 Belgian legal articles. We derive following baselines from prior works of [Louis and Spanakis \(2022\)](#); [Louis et al. \(2023\)](#): Sparse retriever such as (i) BM25 ([Robertson et al., 1995](#)), Dense retrievers such as (ii) Bi-encoder with LegalCamemBERT as query and article encoder (BE w/o Hier.) (iii) Bi-encoder with hierarchical variant for article encoding (BE) (iv) BE along with graph encoder applied on statute structure graph (BE+GE-Stat.) Implementation details can be found in App. A.

3.2 Performance comparison

Following previous work, we evaluate the retriever’s performance using Recall@k (R@K) (k=100,200,500), Mean Average Precision (MAP) and Mean R-Precision (MRP). R@K measures the proportion of relevant articles in the top k candidates, with results averaged across all instances. MAP and MRP provide the mean of average precision and R-Precision scores for each query where average precision is the average of Precision@k scores for every rank position of each relevant document and Precision@k represents the proportion of relevant documents in the top k candidates. R-Precision indicates proportion of the relevant articles in the top-k ranked ones where k is exact number of relevant articles for that query. Higher scores in these metrics indicate better performance.

From Table 1, QABISAR consistently outperforms prior works across all metrics. This validates the superiority of two key aspects: (a) enriched multi-faceted representations of articles shaped by interactions with queries, going beyond BE+GE-Stat. which only considers cross-article dependencies from statute topology graph. (b) transfer of enriched query representations from the graph network to the query bi-encoder, enabling it to grasp the complex semantics embedded, even without direct graph-based guidance during inference.

Ablation Study on QABISAR: We examine the effectiveness of each component in QABISAR

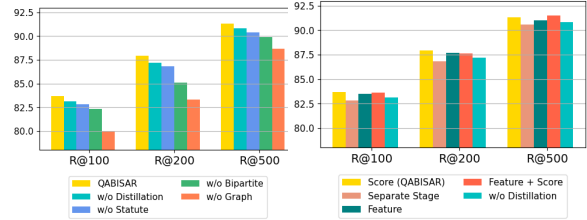


Figure 1: Ablation Study on QABISAR

Figure 2: Effect of various distillation strategies

through these ablations: (1) removing the distillation KD loss during training; (2) excluding bipartite graph interactions and using only the statute topology graph; (3) eliminating statute topology graph and relying solely on the bipartite graph; (4) removing the entire graph encoder. From Fig. 1, we observe that distillation enables effective transfer of query representations into the bi-encoder, rendering them as expressive as graph representations. The inclusion of both the statute structure graph and the statute-query bipartite graph proves more effective, indicating the limitations of modeling relevance of each query-article pair in isolation. Between these two graph views, removing bipartite interactions has a more impact on performance, suggesting that they facilitate the effective capture of multi-faceted information present in articles by leveraging their simultaneous interactions with different queries, as well as other related articles connected by 1-hop neighbor bridges via queries.

Effect of Distillation: We investigate various strategies for distilling query representations from the graph into the bi-encoder: (i) utilizing relevance score via KL divergence loss (ii) feature distillation through L_2 distance loss between both the query representations ([Heo et al., 2019](#)) (iii) combining both feature and score distillation. From Fig. 2, we observe that employing either or both distillation methods proves effective for transferring query representations compared to not using KD. Score distillation outperforms feature distillation, likely due to the limited number of queries in the training set, leading to overfitting losing generalizability with feature distillation alone. Moreover, score distillation enables the representations to become as expressive as graph representations for computing relevance scores, rather than aiming for exact replication as in feature distillation with L_2 loss. The combination of both methods does not yield as effective results as using score distillation alone, reinforcing overfitting effect. We explore the im-

pact of joint training in QABISAR by applying KD after training the graph encoder separately. It underperforms compared to training without KD, as joint training allows for a more gradual steering of representations through the training, compared to the abrupt adjustment in separate stages.

4 Conclusion

We introduced QABISAR, a novel framework for SAR leveraging bipartite interactions between queries and articles to facilitate learning of multi-faceted representations and employ knowledge distillation to transfer the enriched query representations from graph into the bi-encoder. Through comprehensive ablation studies over publicly available BSARD dataset, we demonstrated its effectiveness. We hope this work to inspire more investigations on the different interaction schemes for capturing many-to-many relationship effectively.

Limitations

Our experimental contributions are centered around the BSARD dataset, which is based on the French language and built on the Belgian legal system. It’s important to note that the BSARD dataset introduces a linguistic bias as Belgium is a multilingual country with French, Dutch, and German speakers, yet the legal questions and provisions provided are only available in French (Louis and Spanakis, 2022). While our approach, QABISAR, demonstrates promising results within this specific context, its performance may vary when applied to legal jurisdictions with different legislative structures and languages. Expanding the application of QABISAR to other jurisdictions remains an avenue for future exploration, highlighting the need for concerted efforts to construct SAR datasets from diverse legal systems.

Our work primarily focuses on optimizing recall in the first stage of the retrieval system. For practical utility, a complementary re-ranking component is necessary to improve precision by identifying the most relevant statutes for each query, which are subsequently fed into a QA system to answer legal queries posed by individuals. Moreover, for this QA system to be truly accessible, it should not only retrieve relevant articles but also possess the capability to simplify legal texts, making them comprehensible to laypeople.

Ethics Statement

We conduct experiments using the publicly available SAR dataset, BSARD (Louis and Spanakis, 2022). While leveraging pre-trained encoders enhances our model’s performance, we acknowledge the inherent risk of inheriting biases embedded within these encoders. Consequently, it is imperative to subject our models to thorough scrutiny, particularly concerning equal treatment imperatives related to their performance, behavior, and intended use. Moreover, we recognize the potential for false information through such automated systems, which can have profound implications. Therefore, one needs to remain vigilant in consuming information from automated systems. Additionally, we are mindful of the broader societal impact of our technology, particularly its influence on marginalized communities. We advocate actively engaging with legal stakeholders to ensure ethical and responsible development and deployment of any SAR system.

References

- Shaked Brody, Uri Alon, and Eran Yahav. 2021. How attentive are graph attention networks? In *International Conference on Learning Representations*.
- Jaekool Choi, Euna Jung, Jangwon Suh, and Wonjong Rhee. 2021. Improving bi-encoder document ranking models with two rankers and multi-teacher distillation. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pages 2192–2196.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Bidirectional encoder representations from transformers. *arXiv preprint arXiv:1810.04805*.
- Randy Goebel, Yoshinobu Kano, Mi-Young Kim, Juliano Rabelo, Ken Satoh, and Masaharu Yoshioka. 2023. Summary of the competition on legal information, extraction/entailment (coliee) 2023. In *Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law*, pages 472–480.
- Byeongho Heo, Jeesoo Kim, Sangdoon Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. 2019. A comprehensive overhaul of feature distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1921–1930.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Sebastian Hofstätter, Sophia Althammer, Michael Schröder, Mete Sertkan, and Allan Hanbury. 2020.

- Improving efficient neural ranking models with cross-architecture knowledge distillation. *arXiv preprint arXiv:2010.02666*.
- 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)*, pages 6769–6781.
- Mi-Young Kim, Juliano Rabelo, and Randy Goebel. 2019. Statute law information retrieval and entailment. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, pages 283–289.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From word embeddings to document distances. In *International conference on machine learning*, pages 957–966. PMLR.
- Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Antoine Louis and Gerasimos Spanakis. 2022. A statutory article retrieval dataset in french. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6789–6803.
- Antoine Louis, Gijts Van Dijk, and Gerasimos Spanakis. 2023. Finding the law: Enhancing statutory article retrieval via graph neural networks. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2753–2768.
- Yuxiang Lu, Yiding Liu, Jiayang Liu, Yunsheng Shi, Zhengjie Huang, Shikun Feng, Yu Sun, Hao Tian, Hua Wu, Shuaiqiang Wang, Dawei Yin, et al. 2022. Ernie-search: Bridging cross-encoder with dual-encoder via self on-the-fly distillation for dense passage retrieval. *arXiv preprint arXiv:2205.09153*.
- Raghavendra Pappagari, Piotr Zelasko, Jesús Villalba, Yishay Carmiel, and Najim Dehak. 2019. Hierarchical transformers for long document classification. In *2019 IEEE automatic speech recognition and understanding workshop (ASRU)*, pages 838–844. IEEE.
- Alejandro Ponce, Sarah Chamness Long, Elizabeth Andersen, Camilo Gutierrez Patino, Matthew Harman, Jorge A Morales, Ted Piccone, Natalia Rodriguez Cajamarca, Adriana Stephan, Kirssy Gonzalez, et al. 2019. Global insights on access to justice 2019: Findings from the world justice project general population poll in 101 countries. *World Justice Project*, page 1.
- Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2021. Rocketqa: An optimized training approach to dense passage retrieval for open-domain question answering. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.
- Juliano Rabelo, Randy Goebel, Mi-Young Kim, Yoshinobu Kano, Masaharu Yoshioka, and Ken Satoh. 2022. Overview and discussion of the competition on legal information extraction/entailment (coliee) 2021. *The Review of Socionetwork Strategies*, 16(1):111–133.
- Juliano Rabelo, Mi-Young Kim, Randy Goebel, Masaharu Yoshioka, Yoshinobu Kano, and Ken Satoh. 2021. Coliee 2020: methods for legal document retrieval and entailment. In *New Frontiers in Artificial Intelligence: JSAI-isAI 2020 Workshops, JURISIN, LENLS 2020 Workshops, Virtual Event, November 15–17, 2020, Revised Selected Papers 12*, pages 196–210. Springer.
- Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gattford, et al. 1995. Okapi at trec-3. *Nist Special Publication Sp*, 109:109.
- TYS Santosh, Rashid Gustav Haddad, and Matthias Grabmair. 2024a. Ecthr-pcr: A dataset for precedent understanding and prior case retrieval in the european court of human rights. *arXiv preprint arXiv:2404.00596*.
- TYS Santosh, Elvin Quero Hernandez, and Matthias Grabmair. 2024b. Query-driven relevant paragraph extraction from legal judgments. *arXiv preprint arXiv:2404.00595*.
- TYS Santosh, Kristina Kaiser, and Matthias Grabmair. 2024c. Cusines: Curriculum-driven structure induced negative sampling for statutory article retrieval. *arXiv preprint arXiv:2404.00590*.
- Trevor Strohman, Donald Metzler, Howard Turtle, and W Bruce Croft. 2005. Indri: A language model-based search engine for complex queries. In *Proceedings of the international conference on intelligent analysis*, volume 2, pages 2–6. Washington, DC.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *International Conference on Learning Representations*.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems*, 33:5776–5788.
- Masaharu Yoshioka, Yoshinobu Kano, Naoki Kiyota, and Ken Satoh. 2018. Overview of japanese statute law retrieval and entailment task at coliee-2018. In *The Proceedings of the 12th International Workshop on Juris-Informatics (JURISIN2018)*, pages 117–128.

A Implementation Details

Following [Louis et al. \(2023\)](#), we initialize the second-level encoder in the hierarchical article encoder with a two-layer transformer encoder featuring a hidden dimension of 768, an intermediate dimension of 3072, 12 heads, a dropout rate of 0.1, and the GeLU non-linearity function. Our training process for the dense encoder spans 15 epochs with a batch size of 24, employing the AdamW optimizer ([Loshchilov and Hutter, 2018](#)) with hyperparameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-7$, a weight decay of 0.01, and a learning rate warm-up for the first 5% of training steps, reaching a maximum value of $2e-5$, after which linear decay is applied. For the graph encoder, we conduct 20 epochs of training with a batch size of 512, using the AdamW optimizer with a learning rate of $2e-4$. We use a learnable embedding layer for edge types with dimension equal to node features. We assign weights of 0.7 and 0.3 to the contrastive loss and KD loss, respectively, in the second stage training.