

Finding the Law: Enhancing Statutory Article Retrieval via Graph Neural Networks

Antoine Louis Gijs van Dijck Gerasimos Spanakis

Law & Tech Lab, Maastricht University

{a.louis, gijs.vandijck, jerry.spanakis}@maastrichtuniversity.nl

Abstract

Statutory article retrieval (SAR), the task of retrieving statute law articles relevant to a legal question, is a promising application of legal text processing. In particular, high-quality SAR systems can improve the work efficiency of legal professionals and provide basic legal assistance to citizens in need at no cost. Unlike traditional ad-hoc information retrieval, where each document is considered a complete source of information, SAR deals with texts whose full sense depends on complementary information from the topological organization of statute law. While existing works ignore these domain-specific dependencies, we propose a novel graph-augmented dense statute retriever (G-DSR) model that incorporates the structure of legislation via a graph neural network to improve dense retrieval performance. Experimental results show that our approach outperforms strong retrieval baselines on a real-world expert-annotated SAR dataset.¹

1 Introduction

Today, the high cost of legal expertise prevents less fortunate people from understanding and reacting to legal issues that may arise (Ponce et al., 2019). In recent years, an increasing number of works have focused on legal text processing (Zhong et al., 2020) with the intent to assist legal practitioners and citizens while reducing legal costs and improving equal access to justice for all. Statutory article retrieval (SAR), the task of retrieving statute law articles relevant to a legal question, marks the first and one of the most crucial steps in any legal aid process. Our goal is to help reduce the gap between people and the law by improving SAR systems that could provide citizens with the first component of a free professional legal aid service.

Prior work has addressed SAR with standard information retrieval approaches such as term-based

¹Our source code is available at <https://github.com/maastrichtlawtech/gdsr>.

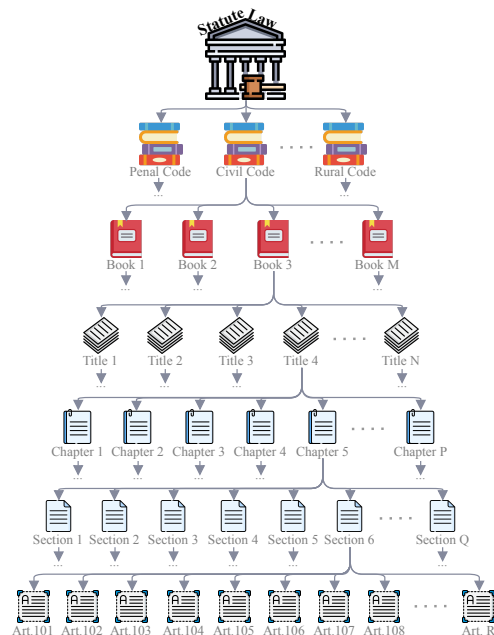


Figure 1: Illustration of the hierarchical organization of statute law. Each law code is structured into books, titles, chapters, and sections. The deeper the divisions, the closer the legal concepts of the articles below them.

models or dense embedding-based models (Kim et al., 2019; Nguyen et al., 2021). While good performance has been achieved, these approaches rely on the flawed assumption that articles are complete and independent sources of information. In reality, statute law is an ensemble of *interdependent* rules meticulously organized into different codes, books, titles, chapters, and sections, as illustrated in Figure 1. Each level in the structure of legislation comes with a unique heading that informs about the content of the articles below it. An article takes on its whole meaning only when considered at its rightful place in the structure with the complementary information from its neighboring articles.

This work shows that such a structure can be highly beneficial for retrieving statutes. We propose a graph-augmented dense statute retriever (G-

DSR) model that leverages the topological structure of legislation to enhance the article content information. Specifically, the proposed model extends the document encoder of a dense retriever with a graph neural network to learn knowledge-rich cross-article representations. Similar to previous work, we adopt a contrastive learning strategy to optimize the similarity between the representations of relevant query-article pairs.

The contributions of this paper are threefold:

- We propose a graph-augmented dense retriever model for statutory article retrieval that explicitly utilizes the topological organization of statute law to enrich the article information.
- We conduct empirical evaluations on our model and demonstrate improvements over strong retrieval baselines.
- We perform ablation studies on various model components and training strategies to understand the impact of several design options on the effectiveness of our model.

2 Preliminaries

In this section, we formally introduce the task of *statutory article retrieval* and discuss the specific difficulties associated with it. We then explain how we identify the structure of legislation as an essential consideration in SAR.

Problem formulation. Given a simple legal question, such as "*Who should pay for the construction of the common wall?*", SAR aims to return one or several relevant articles from the legislation. Formally speaking, a SAR system can be expressed as a function $R : (q, \mathcal{C}) \mapsto \mathcal{F}$ that takes as input a question q along with a corpus of articles $\mathcal{C} = \{a_1, a_2, \dots, a_N\}$, and returns a much smaller filter set $\mathcal{F} \subset \mathcal{C}$ of the supposedly relevant articles, ranked by decreasing order of relevance. For a fixed $k = |\mathcal{F}| \ll |\mathcal{C}|$, the retriever can be evaluated in isolation with multiple rank-based metrics. Most modern retriever systems follow a two-stage retrieval approach (Guo et al., 2016; Hui et al., 2017; McDonald et al., 2018), where a *pre-fetcher* first aims to return all relevant documents in the filter set \mathcal{F} and a *re-ranker* then attempts to make more relevant documents in \mathcal{F} appear before less relevant ones. In this work, we focus our research on improving the pre-fetcher component for SAR.

Challenges. SAR comes with two core challenges that make the task unique compared to tra-

ditional information retrieval. First, the statutes to be retrieved are written in a language that dramatically differs from the ordinary plain language used in the questions. The legal language uses a specialized jargon known for its frequent and deliberate use of formal words, Latin phrases, lengthy sentences, and expressions with flexible meanings (Charrow and Crandall, 1990). Second, statutory articles are long text sequences that may reach several thousand words. This implies overcoming the maximum input length limit of 512 tokens imposed by BERT-based models, which have recently become the standard in neural information retrieval due to their effectiveness.

Structure of legislation. The legislation comes with a well-thought-out organization of its written rules to facilitate access to provisions covering a given subject (Onoge, 2015). This organization is established in a hierarchical manner, where higher-level divisions cover broad legal domains while lower-level divisions deal with specific legal concepts. To examine the importance of this hierarchy in the SAR process, we conduct a preliminary investigation in which we study the reasoning legal experts follow when performing the task. We summarize these experts' approach in Appendix A.1. We observe that legal experts rely heavily on the structure of law when retrieving relevant articles to a legal question, which indicates that the different divisions' headings in the legislation carry valuable information that retrieval systems should consider. Additionally, we analyze the degree to which neighboring articles cover related subjects in Appendix A.2 and find high levels of similarities, which suggests that information from neighboring articles should be considered to capture an article's whole meaning.

3 Approach

In this section, we present a new general approach for SAR that learns to retrieve relevant statutes by using both the textual semantic information from articles and the structural graph information from the legislation. Our model, called graph-augmented dense statute retriever (G-DSR), consists of two main building blocks, as depicted in Figure 2, that are trained independently with the same objective. We first describe the dense retriever component of our approach in Section 3.1 and then explain how our legislative graph encoder builds upon it in Section 3.2.

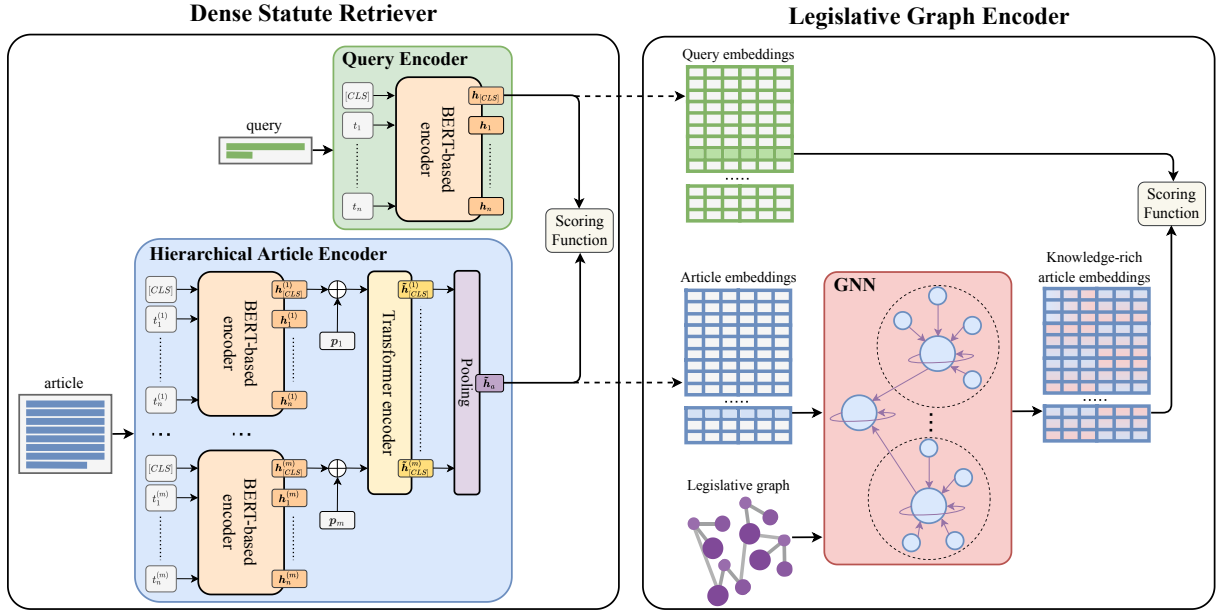


Figure 2: **An illustration of the graph-augmented dense statute retriever (G-DSR) model.** G-DSR consists of two main building blocks that are trained independently. **Left:** The dense statute retriever (DSR) first learns high-quality low-dimensional embedding spaces for both the queries and articles such that relevant query-article pairs appear closer than irrelevant ones in those vector spaces. **Right:** The legislative graph encoder (LGE) then learns to enrich the article representations by aggregating information from the organization of statute law.

3.1 Dense Statute Retriever

Our approach’s first component, called dense statute retriever (DSR), aims to learn high-quality low-dimensional embedding spaces for questions and articles so that relevant question-article pairs appear closer than irrelevant ones in those spaces. Below, we review the overall architecture of the retriever and detail the design of its query and article encoders. We then describe the contrastive learning strategy we employ and choice of negative pairs.

Bi-encoder. We use the widely adopted bi-encoder architecture (Bromley et al., 1993) as the foundation of our dense retriever. The latter maps queries and articles into dense vector representations and calculates a relevance score $s : (q, a) \mapsto \mathbb{R}_+$ between query q and article a by the similarity of their embeddings, i.e.,

$$s(q, a) = \text{sim} \left(E_Q^\theta(q), E_A^\phi(a) \right), \quad (1)$$

where $E_Q^\theta(q), E_A^\phi(a) \in \mathbb{R}^d$ denote the query and article embeddings respectively, and $\text{sim} : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$ is a similarity function such as cosine or dot-product.

Query encoder. To encode the queries, we feed them into a BERT-based (Devlin et al., 2019) model

$E_Q^\theta : \mathcal{W}^n \mapsto \mathbb{R}^d$ with weights θ , that maps an input sequence of n tokens from vocabulary \mathcal{W} to d -dimensional real-valued token embeddings. We take the last layer’s [CLS] token representation as the query embedding, i.e.,

$$E_Q^\theta(q) = \text{BERT}_{[\text{CLS}]}(q). \quad (2)$$

Hierarchical article encoder. Since statutory articles may be longer than the maximum input length of a standard BERT-based encoder, we use a hierarchical variation that can process longer textual sequences (Pappagari et al., 2019; Zhang et al., 2019; Yang et al., 2020a). Each article a is first split into smaller text passages $[p_1, p_2, \dots, p_m]$, where a passage p_i is a sequence of tokens $[t_1^{(i)}, t_2^{(i)}, \dots, t_n^{(i)}]$ with $n \leq 512$. These passages are then independently passed through a shared BERT-based model to extract a list of *context-unaware* passage representations using the respective [CLS] token embeddings, as illustrated in Figure 2. Next, the hierarchical model sums the [CLS] token representations of each passage with learnable passage position embeddings and feeds the resulting representations into a small Transformer encoder to make them aware of the surrounding passages. The final article representation is computed through a pooling operation over the

context-aware passage representations, i.e.,

$$E_A^\phi(a) = \text{pool} \left(\left[\tilde{\mathbf{h}}_{[\text{CLS}]}^{(1)}, \dots, \tilde{\mathbf{h}}_{[\text{CLS}]}^{(m)} \right] \right), \quad (3)$$

where $\tilde{\mathbf{h}}_{[\text{CLS}]}^{(i)} \in \mathbb{R}^d$ is the contextualized embedding of passage p_i , and $\text{pool} : \mathbb{R}^{m \times d} \mapsto \mathbb{R}^d$ is either mean or max pooling.

Contrastive learning. The training objective of the bi-encoder is to learn effective embedding functions $E_Q^\theta(\cdot)$ and $E_A^\phi(\cdot)$ such that relevant question-article pairs have a higher similarity than irrelevant ones. Let $\mathcal{D} = \{\langle q_i, a_i^+ \rangle\}_{i=1}^N$ be the training data where each of the N instances consists of a query q_i associated with a relevant article a_i^+ . By sampling a set of negative articles \mathcal{A}_i^- for each question q_i , we can create a training set $\mathcal{T} = \{\langle q_i, a_i^+, \mathcal{A}_i^- \rangle\}_{i=1}^N$. For each training instance in \mathcal{T} , we contrastively optimize the negative log-likelihood of the positive article against the negative ones, i.e.,

$$L(q_i, a_i^+, \mathcal{A}_i^-) = -\log \frac{e^{s(q_i, a_i^+)/\tau}}{\sum_{a \in \mathcal{A}_i^- \cup \{a_i^+\}} e^{s(q_i, a)/\tau}}, \quad (4)$$

where $\tau > 0$ is a temperature parameter to be set.

Negatives. We consider two types of negative examples: (i) in-batch (Chen et al., 2017; Henderson et al., 2017), i.e., articles paired with the other questions from the same mini-batch, and (ii) BM25, i.e., top articles returned by BM25 that are not relevant to the question.

3.2 Legislative Graph Encoder

Our approach’s second component, called Legislative Graph Encoder (LGE), aims to enrich article representations given by the trained retriever’s article encoder by fusing information from a legislative graph. Below, we elaborate on the legislative graph construction and the graph training process.

Graph construction. To leverage the hierarchical organization of statute law, we formalize the latter as a tree structure consisting of two types of node: (i) *section* nodes, which are titled structural units that represent the consecutive divisions in codes of law (i.e., the headings of the books, titles, chapters, and sections), and (ii) *article* nodes, which are textual content units that represent the different statutory articles. As illustrated in Figure 1, the edges represent the hierarchical connections between section and article nodes. Formally, such a tree can be represented as a directed acyclic

graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with \mathcal{V} as the node set and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ as the edge set.

Node feature initialization. Nodes in \mathcal{V} are commonly associated with d -dimensional features. We apply the article encoder $E_A^\phi(\cdot)$ from the trained bi-encoder to encode the semantic information of nodes (i.e., section headings and article contents) offline and use the resulting embeddings as the initial node features $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$.

Node feature update. To fuse the information of node features using the graph structure, we use a graph neural network (GNN). Such a model consists of a stack of neural network layers, where each layer aggregates local neighborhood information (i.e., features of neighbors) around each node and then passes this aggregated information on to the next layer. Generally speaking, a GNN takes as inputs the feature matrix \mathbf{X} and the graph’s adjacency matrix $\mathbf{A} \in \mathbb{R}_+^{|\mathcal{V}| \times |\mathcal{V}|}$, with $\mathbf{A}_{i,j}$ as the edge weight between nodes i and j , and produces a node-level output $\mathbf{Z} \in \mathbb{R}^{|\mathcal{V}| \times d}$ that captures each node’s structural properties. Every GNN layer can be written as a non-linear function

$$\mathbf{H}^{(l+1)} = f(\mathbf{H}^{(l)}, \mathbf{A}), \quad (5)$$

with $\mathbf{H}^0 = \mathbf{X}$ and $\mathbf{H}^L = \mathbf{Z}$, L being the number of layers. In its simplest form, the layer-wise propagation rule is such that

$$f(\mathbf{H}^{(l)}, \mathbf{A}) = \sigma(\mathbf{A}\mathbf{H}^{(l)}\mathbf{W}^{(l)}), \quad (6)$$

where $\mathbf{W}^{(l)}$ is the input linear transformation’s weight matrix for the l -th neural network layer and $\sigma(\cdot)$ is a non-linear activation function. We propose to use a 3-layer GATv2 network (Brody et al., 2022), a variant of GAT (Velickovic et al., 2018) that has the ability to learn the strength of connection between neighboring nodes through a *dynamic* attention mechanism. Formally, a GATv2 layer updates a node’s hidden state as follows

$$\mathbf{h}_i^{(l+1)} = \sigma \left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(l)} \mathbf{W}^{(l)} \mathbf{h}_j^{(l)} \right), \quad (7)$$

where $\mathcal{N}(i)$ is the set of first-order neighbors of node i , and $\alpha_{ij}^{(l)}$ are normalized attention coefficients indicating the importance of node j ’s features to i in the l -th layer. The latter are computed based on the features of the connected nodes using an attention function $\text{att} : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$ such that

$$\alpha_{ij}^{(l)} = \text{softmax} \left(\sigma(\text{att}(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)})) \right). \quad (8)$$

Learning process. To optimize the GNN parameters, we adopt the same contrastive learning strategy used to train the bi-encoder. Since graph \mathcal{G} can be relatively large, performing an update of all the node features in \mathcal{G} at every training iteration would incur high computational costs. Besides, most of these computations would be of no use as only the updated representations of nodes $\{\mathcal{A}_i^- \cup a_i^+\}_{i=1}^{|\mathcal{B}|}$ from batch \mathcal{B} are needed to update the model parameters. Therefore, we build a sub-graph \mathcal{G}^{sub} at each training step that only contains the article nodes from batch \mathcal{B} as well as their L -hop neighbors (where L is decided by the number of GNN layers). We then pass that sub-graph to the graph network and use the resulting article representations to compute the loss in Equation (4). Comparably to the node features, the query embeddings are pre-computed offline before training by the query encoder $E_Q^\theta(\cdot)$ of our trained bi-encoder.

4 Experimental Setup

In this section, we present the basic setup for experiments. In particular, Section 4.1 describes the dataset we conduct our experiments on, Section 4.2 details our model implementation, Section 4.3 reviews the different baselines we use for comparison, and Section 4.4 reports the evaluation metrics.

4.1 Dataset

We conduct experiments on the publicly available Belgian Statutory Article Retrieval Dataset (Louis and Spanakis, 2022, BSARD).² To the best of our knowledge, BSARD is the only SAR dataset that provides the lists of consecutive division headings each article belongs to, which is crucial for building the graph of the legislative structure. The dataset consists of 1,100+ *French* native questions on various legal topics, as shown in Table 1, labeled by skilled experts with references to relevant statutory articles from the Belgian legislation. The retrieval corpus comprises 22,600+ articles collected from 32 Belgian codes covering numerous legal domains. The questions are relatively short and might have several relevant legal articles. We refer readers to the original paper for further data collection and analysis details.

4.2 Implementation Details

Model. We use the publicly released CamemBERT (Martin et al., 2020) checkpoint to initialize

²<https://huggingface.co/datasets/antoiloui/bsard>








Topic	Train	Dev	Test
 Family	216	56	67
 Housing	203	38	66
 Money	103	35	36
 Justice	96	25	30
 Foreigners	41	9	13
 Social security	27	8	6
 Work	23	6	4
Total	709	177	222

Table 1: Topic distribution of questions in BSARD.

DSR’s query encoder. Due to the specificity of the legal language the article encoder has to deal with, we follow prior work on domain adaptation (Gururangan et al., 2020; Jørgensen et al., 2021) and continue pre-training CamemBERT on BSARD statutory articles for 50k gradient steps to adapt it to the target legal domain. We use the resulting domain-specific checkpoint to warm-start the article’s first-level encoder. The second-level encoder is a two-layer Transformer encoder of 14M parameters with a similar configuration (i.e., 768-hidden, 3072-intermediate, 12-heads, 0.1 dropout, GeLU). We use max-pooling to aggregate the final chunk representations and cosine as the decomposable similarity function.

Data augmentation. Due to the recent success in using synthetic query generation to improve dense retrieval performance (Liang et al., 2020; Ma et al., 2021; Thakur et al., 2021), we propose to augment BSARD with synthetic domain-targeted queries. We use a mT5 model (Raffel et al., 2020) fine-tuned on general domain data from mMARCO (Bonifacio et al., 2021) to synthesize queries for our target statutory articles.³ We generate five queries per article, which results in a total of around 118k synthetic queries. We combine the latter with the gold BSARD train samples and obtain an augmented training set of around 122.5k question-article pairs.

Optimization. We train DSR for 15 epochs with a batch size of 24 using AdamW (Loshchilov and Hutter, 2017) with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-7$, weight decay of 0.01, and learning rate warm up along the first 5% of the training steps to a maximum value of $2e-5$, after which linear decay is applied. We then optimize LGE parameters for 10

³<https://huggingface.co/doc2query/msmarco-french-mt5-base-v1>

epochs with a batch size of 512 using AdamW with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-7$, weight decay of 0.1, and a constant learning rate of $2e-4$. We use 16-bit automatic mixed precision to accelerate training and save memory. Details on our hyperparameter tuning process are given in Appendix B.

Hardware & schedule. Training is performed on a single 32 GB NVIDIA V100 GPU hosted on a server with a dual 20-core Intel Xeon E5-2698 v4 CPU @2.20GHz and 512 GB of RAM. It takes around 1 day to train DSR and 35 minutes for LGE.

Libraries. We implement, train, and tune our models using Transformers (Wolf et al., 2020), PyTorch (Paszke et al., 2019), PyTorch-Geometric (Fey and Lenssen, 2019), PyTorch-Lightning (Falcon, 2019), W&B Sweeps (Biewald, 2020), and DeepSpeed (Rasley et al., 2020).

4.3 Baselines

We compare our approach against three strong retrieval systems. As a sparse baseline model, we follow prior work and consider BM25 (Robertson et al., 1994),⁴ a popular bag-of-words retrieval function based on exact term matching. We then examine the document expansion technique docT5query (Nogueira and Lin, 2019), which augments each article with a pre-defined number of synthetic queries generated by a finetuned mT5 model,³ and then uses a traditional BM25 lexical index from the augmented articles for retrieval. Last, we include the results of a supervised dense passage retrievers (Karpukhin et al., 2020, DPR) pre-finetuned on more than 90.5k question-context pairs from a combination of three French QA datasets.⁵

4.4 Evaluation

We evaluate model performance using three commonly used ranking measures (Manning et al., 2008), namely the macro-averaged recall at different cutoffs ($R@k$), mean average precision (mAP), and mean r-precision (mRP). Those metrics are further defined in Appendix D. We deliberately omit to report the $\text{precision}@k$ given that questions in BSARD have a variable number of relevant articles, which implies that questions with r relevant articles would always have $\text{P}@k < 1$ if $k > r$. Similarly, the mean reciprocal rank (mRR) is not appropriate

⁴We use $k_1 = 2.5$ and $b = 0.2$. Details on BM25 hyperparameters tuning are given in Appendix C.

⁵https://huggingface.co/etalab-ia/dpr-question_encoder-fr_qa-camembert

for BSARD as only the first relevant article would be considered. As some questions might have up to 100 relevant articles, we use $k \in \{100, 200, 500\}$ for the $\text{recall}@k$.

5 Experiments

In this section, we empirically evaluate the effectiveness of our proposed approach against competitive baselines and discuss the main results in Section 5.1. Next, we provide an ablation study in Section 5.2 to understand how different design and training options affect our model’s performance.

5.1 Main Results

Table 2 shows retrieval performance on BSARD test set. Although we report model performance on two rank-aware metrics (i.e., mAP and mRP), we emphasize that our approach is specifically aimed at improving the pre-fetching component of a retriever (Zhang et al., 2021a) and therefore focuses on optimizing rank-unaware metrics (i.e., $R@k$).

First, we compare the performance of our proposed G-DSR model⁽⁸⁾ against other well-known retrieval approaches and find it significantly outperforms all of them on SAR. In particular, it improves over the sparse retrieval methods^(1,2) by around 30% on $\text{recall}@k$ and by more than 25% on mAP and mRP. It also outperforms a competitive pre-finetuned DPR model⁽⁴⁾ by 6% on $R@100$, 9% on $R@200$, and 5% on $R@500$. However, the latter shows a better performance on rank-aware metrics compared to our DSR models, which we speculate might be due to its extensive pre-finetuning step on three domain-general retrieval datasets, leading the model to a deeper knowledge of the task at hand.

Next, we investigate the influence of different training strategies on the rank-unaware results of our base dense retriever.⁽⁵⁾ We find that DSR’s performance is improved when adapting the article text encoder to the legal domain before finetuning on the target data.⁽⁶⁾ Besides, training DSR on a larger dataset containing synthetic domain-targeted queries improves its performance even more.⁽⁷⁾

Finally, our results show that using a GNN model on top of DSR allows to enrich the article representations and leads to the best overall performance.⁽⁸⁾ Interestingly, G-DSR also significantly improves the rank-aware performance of our best performing DSR model by $\sim 12\%$, suggesting that a GNN could act as an effective re-ranker for SAR.

Model	#Params	R@100 (↑)	R@200 (↑)	R@500 (↑)	mAP (↑)	mRP (↑)
Baselines						
1 BM25	-	49.3	57.3	63.0	16.8	13.6
2 docT5query	-	51.7	59.4	65.8	18.7	15.0
4 DPR	220M	77.9	81.3	88.2	45.4	39.1
Ours						
5 DSR	234M	77.1	81.8	86.7	35.6	28.8
6 DSR w. domain-adaptive pre-training	234M	79.8	83.9	88.9	39.5	31.3
7 DSR w. data augmentation	234M	82.7	88.7	92.8	35.3	27.5
8 G-DSR	262M	84.3	90.4	93.1	47.1	40.2

Table 2: Retrieval performance on BSARD Test set. The best results are marked in bold.

5.2 Ablation Study

To further understand how different design choices and training strategies affect the results, we conduct several additional experiments and discuss our findings below.

Alternative pre-trained LMs. In addition to CamemBERT, we experiment with several other French or multilingual pre-trained language models to initialize the first-level text encoders in DSR – namely, mBERT (Devlin et al., 2019), XLM-R (Conneau et al., 2020), and ELECTRA-fr (Clark et al., 2020).⁶ We fine-tune the different warm-started models on BSARD training set and report dev results in Table 3. We find that a CamemBERT-initialized DSR model performs best.

Alternative GNNs. Additionally to GATv2, we explore different GNN architectures to perform the node feature update – namely, GCN (Kipf and Welling, 2017), GraphSAGE (Hamilton et al., 2017), GAT (Velickovic et al., 2018), and k-GNN (Morris et al., 2019) – and summarize the results in Table 4. Our experiments show that using an alternative GNN model does not affect performance much, which suggests that the act of fusing information from neighboring nodes is more important than the way the aggregation is performed.

Similarity and loss functions. Besides cosine for scoring pairs of query-article representations, we also experiment with dot-product and Euclidean distance and find both inferior to cosine. As an alternative to negative log-likelihood, we test the triplet loss (Burgess et al., 2005) and observe that the latter significantly decreases model performance. More details can be found in Appendix E.

⁶<https://huggingface.co/dbmdz/electra-base-french-europeana-cased-discriminator>

Pre-trained LM	#Params	R@100 (↑)	R@500 (↑)
XLM-R	278M	59.8	78.7
mBERT	177M	69.2	86.6
ELECTRA-fr	110M	57.6	73.7
CamemBERT	110M	75.4	88.5

Table 3: BSARD Dev results of DSR warm-started with different pre-trained word embedding models.

Model	#Params	R@100 (↑)	R@500 (↑)
GCN	14M	84.4	92.0
GAT	14M	84.4	92.1
GraphSAGE	28M	84.5	92.9
k-GNN	28M	83.8	93.0
GATv2	28M	84.8	92.3

Table 4: BSARD Dev results of LGE with different (3-layer) GNN architectures.

6 Related Work

Our work operates at the intersection of several research areas, including long document modeling, dense information retrieval, graph neural networks, and legal NLP.

Long document modeling. The emergence of deep neural networks for language processing brought new challenges to text encoding, one of which is learning high-quality representations of long documents. For example, Tang et al. (2015) employ a bottom-up approach using CNN and BiLSTM-based hierarchical networks, where sentences are first encoded into vectors, which are then combined to form a single document vector. Similarly, Yang et al. (2016) build a document vector by aggregating important words into sentence vectors and then aggregating important sentence vectors to document vectors using attention mech-

anisms. More recently, hierarchical variants of Transformer-based models have been explored for various language tasks, including document classification (Mulyar et al., 2019; Pappagari et al., 2019; Chalkidis et al., 2019; Wu et al., 2021), summarization (Zhang et al., 2019), semantic matching (Yang et al., 2020a), and question answering (Liu et al., 2022). In addition to hierarchical attention Transformer-based (HAT) models, several sparse attention Transformers (SAT) have been introduced to reduce the computational complexity of the model, thus allowing to process sequences longer than 512 tokens (Child et al., 2019; Beltagy et al., 2020; Zaheer et al., 2020). However, Chalkidis et al. (2022) show that a pre-trained HAT model performs comparably or better than an equally-sized SAT model across several downstream tasks while being substantially faster and less memory intensive. Recently, other non-Transformer-based approaches have been proposed for efficient long sequence processing based on structured state spaces (Gu et al., 2022; Gupta et al., 2022).

Dense information retrieval. Traditionally, lexical approaches such as TF-IDF and BM25 (Robertson et al., 1994) have been the *de facto* standard for textual information retrieval due to their robustness and efficiency. However, these approaches suffer from the lexical gap problem (Berger et al., 2000) and can only retrieve documents containing keywords present in the query. To overcome this limitation, recent work relies on neural-based architectures to capture semantic relationships between pairs of texts (Lee et al., 2019; Karpukhin et al., 2020; Yang et al., 2020b; Xiong et al., 2021). These models map queries and documents into dense vector representations and calculate a relevance score by the similarity of the vectors (Gillick et al., 2018), which allows the document representations to be pre-computed and indexed offline for inference. The dense retrieval approach was recently extended by hybrid lexical-dense methods, which aim to combine the strengths of both approaches (Seo et al., 2019; Gao et al., 2021; Luan et al., 2021). We refer the readers to Yates et al. (2021) for a survey on neural information retrieval.

Graph neural networks. Graph neural networks (GNNs) capture the topological relationships among the nodes of a graph using an information diffusion mechanism that propagates node features according to the underlying graph-structured data

(Scarselli et al., 2009). These models have shown their effectiveness and flexibility in a wide variety of NLP tasks, including text classification (Lin et al., 2021; Yu et al., 2022), relation extraction (Zhang et al., 2018; Li et al., 2020; Carbonell et al., 2020), and question answering (Cao et al., 2019; Xu et al., 2021b). Recently, GNNs have been employed for document retrieval to enhance the vector representations by leveraging the topological structure of the documents, where nodes are passages from a document and edges are relations between these passages (Xu et al., 2021a; Zhang et al., 2021b; Albaredo et al., 2022).

Application to the legal domain. In recent years, the legal domain has attracted much interest in the NLP community, both for its challenging characteristics and massive volumes of textual data (Chalkidis and Kampas, 2019; Zhong et al., 2020). Researchers see it as an opportunity to develop novel automated methodologies that can reduce heavy and redundant tasks for legal professionals while providing a reliable, affordable form of legal support for laypeople (Bommasani et al., 2021). Earlier techniques for legal information retrieval were mainly based on term-matching approaches (Kim and Goebel, 2017; Tran et al., 2018). Recently, a growing number of works have used neural networks to enhance retrieval performance, including word embedding models (Landthaler et al., 2016), doc2vec models (Sugathadasa et al., 2018), CNN-based models (Tran et al., 2019), and BERT-based models (Nguyen et al., 2021; Chalkidis et al., 2021; Althammer et al., 2022). To the best of our knowledge, we are the first to exploit the structure of statute law with GNNs to improve the performance of dense retrieval models.

7 Conclusion

In this paper, we introduce G-DSR, a novel approach for statutory article retrieval (SAR) that leverages the topological structure of legislation to improve retrieval performance. Specifically, G-DSR enriches the article representations of a dense retriever designed for long document retrieval by employing a graph neural network that uses the organization of statute law to learn knowledge-rich cross-article embeddings. Experiments show that G-DSR outperforms competitive baselines on a real-world expert-annotated SAR dataset. We also include a detailed analysis to motivate our design choices and training strategies.

Limitations

While our approach performs well on statutory article retrieval, it comes with several limitations that provide avenues for future work.

First, experimental results are based on questions and labels drafted by legal professionals. It is possible that other legal professionals would draft the questions differently or, less likely yet possible, that they would deem different statutory provisions relevant. This raises the question of to what extent similar results would be obtained if the model were trained on a different dataset, for instance, based on other experts or domains, hence testing the approach's generalizability. The main challenge in this regard is obtaining data, as organizations are unlikely to share or even collect similar data.

Second, our proposed methodology was evaluated exclusively on the Belgian legislation, whose laws are organized in a hierarchical manner where the deeper the divisions, the more closely related the legal concepts of the articles under them. Although we believe our approach could be applied to most, if not all, jurisdictions that rely on statute law (including both civil and common law countries), different jurisdictions may have different organizations of their legal provisions, which could potentially affect the model's performance. It is also worth mentioning that the dataset used for evaluation comes with a linguistic bias as Belgium is a multilingual country with French, Dutch, and German speakers, but the provided provisions are only available in French. Studying the applicability and impact of the present work to other jurisdictions and languages is an exciting research direction that is challenging in practice due to the scarcity of high-quality multilingual statute retrieval datasets.

Then, our approach currently considers the topological structure of legislation for modeling the inter-article dependencies, which implies that information is aggregated between direct neighboring articles only while those from more distant sections are completely ignored. Nevertheless, it is common for articles to cite other articles from different sections or even different statutes. Therefore, we believe that considering richer legal graph structures, especially legal citation networks, could increase effectiveness even more. However, building such citation networks from raw texts requires a considerable text-processing effort.

Finally, although G-DSR shows promise for statutory article retrieval, it is not yet ready for

practical use in the real world. One issue is that our model is designed to be an effective pre-fetcher, optimizing recall such that all articles relevant to a question appear in an unordered filter set of size k (k being relatively large). However, in practice, users would expect a high-quality retrieval system to not only find these relevant articles but also to sort them by decreasing order of importance, requiring an adequate re-ranker. Then, it is essential to recognize that while access to relevant legal provisions is a necessary step in helping the general public solve their legal issues, it is not a sufficient condition on its own as laypeople may still struggle to understand the legal jargon and apply the provisions to their specific situations. Ideally, the tool to be made accessible to the public should consist of a two-stage framework: (i) a legal provision retriever, which selects a small subset of relevant legal articles in response to a given question, and (ii) a legal-to-natural translator or summarizer, which examines the retrieved articles and generates an answer in natural language. In the present work, we chose to focus on the first stage of this framework and leave the second for future work.

Ethics Statement

The scope of this work is to provide a new methodology along with extensive experiments to drive research forward in statutory article retrieval. We believe the latter is an important application field where more research should be conducted to improve legal aid services and access to justice for all. We do not foresee situations where the use of our methodology would lead to harm (Tsarapatsanis and Aletras, 2021). Nevertheless, although our goal is to improve the understanding of the law by those who suffer from legal information asymmetry, it cannot be excluded that the technology presented here could exacerbate inequality if states, companies, or lawyers benefit more from its use than the intended beneficiaries (i.e., citizens, consumers, or employees).

Acknowledgments

This research is partially supported by the Sector Plan Digital Legal Studies of the Dutch Ministry of Education, Culture, and Science. In addition, this research was made possible, in part, using the Data Science Research Infrastructure (DSRI) hosted at Maastricht University.

References

- Lucas Albarede, Philippe Mulhem, Lorraine Goeuriot, Claude Le Pape-Gardeux, Sylvain Marie, and Trinidad Chardin-Segui. 2022. [Passage retrieval on structured documents using graph attention networks](#). In *Proceedings of the 44th European Conference on Information Retrieval Research*, volume 13186 of *Lecture Notes in Computer Science*, pages 13–21. Springer. [Page 8]
- Sophia Althammer, Sebastian Hofstätter, Mete Sertkan, Suzan Verberne, and Allan Hanbury. 2022. [PARM: A paragraph aggregation retrieval model for dense document-to-document retrieval](#). In *Proceedings of the 44th European Conference on Information Retrieval Research*, volume 13185 of *Lecture Notes in Computer Science*, pages 19–34. Springer. [Page 8]
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *CoRR*, abs/2004.05150. [Page 8]
- Adam L. Berger, Rich Caruana, David Cohn, Dayne Freitag, and Vibhu O. Mittal. 2000. [Bridging the lexical chasm: statistical approaches to answer-finding](#). In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 192–199. ACM. [Page 8]
- Lukas Biewald. 2020. [Experiment tracking with weights and biases](#). Software available from wandb.com. [Page 6]
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen Creel, Jared Quincy Davis, Dorottya Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kuditipudi, and et al. 2021. [On the opportunities and risks of foundation models](#). *CoRR*, abs/2108.07258. [Page 8]
- Luiz Henrique Bonifacio, Israel Campiotti, Roberto de Alencar Lotufo, and Rodrigo Nogueira. 2021. [mmarco: A multilingual version of MS MARCO passage ranking dataset](#). *CoRR*, abs/2108.13897. [Page 5]
- Shaked Brody, Uri Alon, and Eran Yahav. 2022. [How attentive are graph attention networks?](#) In *Proceedings of the 10th International Conference on Learning Representations*. OpenReview.net. [Page 4]
- Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. 1993. [Signature verification using a siamese time delay neural network](#). *Advances in Neural Information Processing Systems*, 6:737–744. [Page 3]
- Christopher J. C. Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Gregory N. Hullender. 2005. [Learning to rank using gradient descent](#). In *Proceedings of the 22nd International Conference on Machine Learning*, volume 119 of *ACM International Conference Proceeding Series*, pages 89–96. ACM. [Page 7]
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2019. [Question answering by reasoning across documents with graph convolutional networks](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, pages 2306–2317. Association for Computational Linguistics. [Page 8]
- Manuel Carbonell, Pau Riba, Mauricio Villegas, Alicia Fornés, and Josep Lladós. 2020. [Named entity recognition and relation extraction with graph neural networks in semi structured documents](#). In *Proceedings of the 25th International Conference on Pattern Recognition*, pages 9622–9627. IEEE. [Page 8]
- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. [Neural legal judgment prediction in english](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 4317–4323. Association for Computational Linguistics. [Page 8]
- Ilias Chalkidis, Xiang Dai, Manos Fergadiotis, Prodromos Malakasiotis, and Desmond Elliott. 2022. [An exploration of hierarchical attention transformers for efficient long document classification](#). *CoRR*, abs/2210.05529. [Page 8]
- Ilias Chalkidis, Manos Fergadiotis, Nikolaos Manginas, Eva Katakalous, and Prodromos Malakasiotis. 2021. [Regulatory compliance through doc2doc information retrieval: A case study in EU/UK legislation where text similarity has limitations](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3498–3511. Association for Computational Linguistics. [Pages 8 and 15]
- Ilias Chalkidis and Dimitrios Kampas. 2019. [Deep learning in law: early adaptation and legal word embeddings trained on large corpora](#). *Artificial Intelligence and Law*, 27(2):171–198. [Page 8]
- Veda R Charrow and Jo Ann Crandall. 1990. [Legal language: What is it and what can we do about it?](#) [Page 2]
- Ting Chen, Yizhou Sun, Yue Shi, and Liangjie Hong. 2017. [On sampling strategies for neural network-based collaborative filtering](#). In *Proceedings of the 23rd ACM SIGKDD International Conference on*

- Knowledge Discovery and Data Mining*, pages 767–776. [Page 4]
- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. [Generating long sequences with sparse transformers](#). *CoRR*, abs/1904.10509. [Page 8]
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: pre-training text encoders as discriminators rather than generators](#). *CoRR*, abs/2003.10555. [Page 7]
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451. Association for Computational Linguistics. [Page 7]
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, pages 4171–4186. Association for Computational Linguistics. [Pages 3 and 7]
- William Falcon. 2019. [PyTorch Lightning](#). [Page 6]
- Matthias Fey and Jan Eric Lenssen. 2019. [Fast graph representation learning with PyTorch Geometric](#). In *ICLR Workshop on Representation Learning on Graphs and Manifolds*. [Page 6]
- Luyu Gao, Zhuyun Dai, Tongfei Chen, Zhen Fan, Benjamin Van Durme, and Jamie Callan. 2021. [Complement lexical retrieval model with semantic residual embeddings](#). In *Proceedings of the 43rd European Conference on Information Retrieval Research*, volume 12656 of *Lecture Notes in Computer Science*, pages 146–160. Springer. [Page 8]
- Daniel Gillick, Alessandro Presta, and Gaurav Singh Tomar. 2018. [End-to-end retrieval in continuous space](#). *CoRR*, abs/1811.08008. [Page 8]
- Albert Gu, Karan Goel, and Christopher Ré. 2022. [Efficiently modeling long sequences with structured state spaces](#). In *Proceedings of the 10th International Conference on Learning Representations*. OpenReview.net. [Page 8]
- Jiafeng Guo, Yixing Fan, Qingyao Ai, and W. Bruce Croft. 2016. [A deep relevance matching model for ad-hoc retrieval](#). In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*, pages 55–64. ACM. [Page 2]
- Ankit Gupta, Albert Gu, and Jonathan Berant. 2022. [Diagonal state spaces are as effective as structured state spaces](#). *CoRR*, abs/2203.14343. [Page 8]
- Suchin Gururangan, Ana Marasovic, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360. Association for Computational Linguistics. [Page 5]
- William L. Hamilton, Zhitao Ying, and Jure Leskovec. 2017. [Inductive representation learning on large graphs](#). *Advances in Neural Information Processing Systems*, 30:1024–1034. [Page 7]
- Matthew L. Henderson, Rami Al-Rfou, Brian Strope, Yun-Hsuan Sung, László Lukács, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. 2017. [Efficient natural language response suggestion for smart reply](#). *CoRR*, abs/1705.00652. [Page 4]
- Kai Hui, Andrew Yates, Klaus Berberich, and Gerard de Melo. 2017. [PACRR: A position-aware neural IR model for relevance matching](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1049–1058. Association for Computational Linguistics. [Page 2]
- Rasmus Kær Jørgensen, Mareike Hartmann, Xiang Dai, and Desmond Elliott. 2021. [mdapt: Multilingual domain adaptive pretraining in a single model](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3404–3418. Association for Computational Linguistics. [Page 5]
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. 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*, pages 6769–6781. Association for Computational Linguistics. [Pages 6 and 8]
- Mi-Young Kim and Randy Goebel. 2017. [Two-step cascaded textual entailment for legal bar exam question answering](#). In *Proceedings of the 16th edition of the International Conference on Artificial Intelligence and Law*, pages 283–290. ACM. [Page 8]
- Mi-Young Kim, Juliano Rabelo, and Randy Goebel. 2019. [Statute law information retrieval and entailment](#). In *Proceedings of the 6th Competition on Legal Information Retrieval and Entailment Workshop in association with the Seventeenth International Conference on Artificial Intelligence and Law*, pages 283–289. ACM. [Page 1]
- Thomas N. Kipf and Max Welling. 2017. [Semi-supervised classification with graph convolutional networks](#). In *Proceedings of the 5th International Conference on Learning Representations*. OpenReview.net. [Page 7]
- Jörg Landthaler, Bernhard Waltl, Patrick Holl, and Florian Matthes. 2016. [Extending full text search for legal document collections using word embeddings](#).

- In *Proceedings of the 29th International Conference on Legal Knowledge and Information Systems*, volume 294 of *Frontiers in Artificial Intelligence and Applications*, pages 73–82. IOS Press. [Page 8]
- Kenton Lee, Ming-Wei Chang, and Kristina Toutanova. 2019. [Latent retrieval for weakly supervised open domain question answering](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 6086–6096. Association for Computational Linguistics. [Page 8]
- Bo Li, Wei Ye, Zhonghao Sheng, Rui Xie, Xiangyu Xi, and Shikun Zhang. 2020. [Graph enhanced dual attention network for document-level relation extraction](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1551–1560. International Committee on Computational Linguistics. [Page 8]
- Davis Liang, Peng Xu, Siamak Shakeri, Cícero Nogueira dos Santos, Ramesh Nallapati, Zhiheng Huang, and Bing Xiang. 2020. [Embedding-based zero-shot retrieval through query generation](#). *CoRR*, abs/2009.10270. [Page 5]
- Yuxiao Lin, Yuxian Meng, Xiaofei Sun, Qinghong Han, Kun Kuang, Jiwei Li, and Fei Wu. 2021. [Bertgcn: Transductive text classification by combining GCN and BERT](#). *CoRR*, abs/2105.05727. [Page 8]
- Yang Liu, Jiayang Liu, Li Chen, Yuxiang Lu, Shikun Feng, Zhida Feng, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2022. [ERNIE-SPARSE: learning hierarchical efficient transformer through regularized self-attention](#). *CoRR*, abs/2203.12276. [Page 8]
- Ilya Loshchilov and Frank Hutter. 2017. [Decoupled weight decay regularization](#). In *Proceedings of the 7th International Conference on Learning Representations*. [Page 5]
- 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. Association for Computational Linguistics. [Pages 5, 14, and 15]
- Yi Luan, Jacob Eisenstein, Kristina Toutanova, and Michael Collins. 2021. [Sparse, dense, and attentional representations for text retrieval](#). *Transactions of the Association for Computational Linguistics*, 9:329–345. [Page 8]
- Ji Ma, Ivan Korotkov, Yinfei Yang, Keith B. Hall, and Ryan T. McDonald. 2021. [Zero-shot neural passage retrieval via domain-targeted synthetic question generation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1075–1088. Association for Computational Linguistics. [Page 5]
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to information retrieval*. Cambridge University Press. [Page 6]
- Louis Martin, Benjamin Müller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamé Seddah, and Benoît Sagot. 2020. [Camembert: a tasty french language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219. Association for Computational Linguistics. [Page 5]
- Ryan T. McDonald, George Brokos, and Ion Androutsopoulos. 2018. [Deep relevance ranking using enhanced document-query interactions](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1849–1860. Association for Computational Linguistics. [Page 2]
- Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. 2019. [Weisfeiler and leman go neural: Higher-order graph neural networks](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence*, pages 4602–4609. AAAI Press. [Page 7]
- Andriy Mulyar, Elliot Schumacher, Masoud Rouhizadeh, and Mark Dredze. 2019. [Phenotyping of clinical notes with improved document classification models using contextualized neural language models](#). *CoRR*, abs/1910.13664. [Page 8]
- Ha-Thanh Nguyen, Phuong Minh Nguyen, Thi-Hai-Yen Vuong, Quan Minh Bui, Chau Minh Nguyen, Tran Binh Dang, Vu Tran, Minh Le Nguyen, and Ken Satoh. 2021. [JNLP team: Deep learning approaches for legal processing tasks in COLIEE 2021](#). *CoRR*, abs/2106.13405. [Pages 1 and 8]
- Rodrigo Nogueira and Jimmy Lin. 2019. [From doc2query to docttttquery](#). [Page 6]
- Elohor Onoge. 2015. [Structure of legislation: A paradigm for accessibility and effectiveness](#). *European Journal of Law Reform*, 17:440–470. [Page 2]
- 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*, pages 838–844. IEEE. [Pages 3 and 8]
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). *Advances in Neural Information Processing Systems*, 32:8024–8035. [Page 6]
- Alejandro Ponce, Sarah Chamness Long, Elizabeth Andersen, Camilo Gutierrez Patino, Matthew Harman, Jorge A Morales, Ted Piccone, Natalia Rodriguez Cajarca, Adriana Stephan, Kirssy Gonzalez, Jennifer VanRiper, Alicia Evangelides, Rachel Martin,

- Priya Khosla, Lindsey Bock, Erin Campbell, Emily Gray, Amy Gryskiewicz, Ayyub Ibrahim, Leslie Solis, Gabriel Hearn-Desautels, and Francesca Tinucci. 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]
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. *Exploring the limits of transfer learning with a unified text-to-text transformer*. *Journal of Machine Learning Research*, 21:140:1–140:67. [Page 5]
- Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. 2020. *DeepSpeed: System optimizations enable training deep learning models with over 100 billion parameters*. In *The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3505–3506. [Page 6]
- Stephen E. Robertson, Steve Walker, Susan Jones, Micheline Hancock-Beaulieu, and Mike Gatford. 1994. *Okapi at TREC-3*. In *Proceedings of the 3rd Text REtrieval Conference*, volume 500-225 of *NIST Special Publication*, pages 109–126. National Institute of Standards and Technology (NIST). [Pages 6 and 8]
- Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2009. *The graph neural network model*. *IEEE Transactions on Neural Networks*, 20(1):61–80. [Page 8]
- Min Joon Seo, Jinhyuk Lee, Tom Kwiakowski, Ankur P. Parikh, Ali Farhadi, and Hannaneh Hajishirzi. 2019. *Real-time open-domain question answering with dense-sparse phrase index*. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 4430–4441. Association for Computational Linguistics. [Page 8]
- Keet Sugathadasa, Buddhi Ayesha, Nisansa de Silva, Amal Shehan Perera, Vindula Jayawardana, Dimuthu Lakmal, and Madhavi Perera. 2018. *Legal document retrieval using document vector embeddings and deep learning*. In *Science and information conference*, pages 160–175. Springer. [Page 8]
- Duyu Tang, Bing Qin, and Ting Liu. 2015. *Document modeling with gated recurrent neural network for sentiment classification*. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1422–1432. The Association for Computational Linguistics. [Page 7]
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. *BEIR: A heterogenous benchmark for zero-shot evaluation of information retrieval models*. *CoRR*, abs/2104.08663. [Page 5]
- Vu Tran, Son Truong Nguyen, and Minh Le Nguyen. 2018. *JNLP group: legal information retrieval with summary and logical structure analysis*. In *12th International Workshop on Juris-informatics*. [Page 8]
- Vu D. Tran, Minh Le Nguyen, and Ken Satoh. 2019. *Building legal case retrieval systems with lexical matching and summarization using A pre-trained phrase scoring model*. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, pages 275–282. ACM. [Page 8]
- Dimitrios Tsarapatsanis and Nikolaos Aletras. 2021. *On the ethical limits of natural language processing on legal text*. In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021*, pages 3590–3599. Association for Computational Linguistics. [Page 9]
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. *Graph attention networks*. In *Proceedings of the 6th International Conference on Learning Representations*. OpenReview. [Pages 4 and 7]
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. *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. Association for Computational Linguistics. [Page 6]
- Chuhan Wu, Fangzhao Wu, Tao Qi, and Yongfeng Huang. 2021. *Hi-transformer: Hierarchical interactive transformer for efficient and effective long document modeling*. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, pages 848–853. Association for Computational Linguistics. [Page 8]
- Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N. Bennett, Junaid Ahmed, and Arnold Overwijk. 2021. *Approximate nearest neighbor negative contrastive learning for dense text retrieval*. In *Proceedings of the 9th International Conference on Learning Representations*. OpenReview. [Page 8]
- Peng Xu, Xinchu Chen, Xiaofei Ma, Zhiheng Huang, and Bing Xiang. 2021a. *Contrastive document representation learning with graph attention networks*. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3874–3884. Association for Computational Linguistics. [Page 8]
- Yichong Xu, Chenguang Zhu, Ruochen Xu, Yang Liu, Michael Zeng, and Xuedong Huang. 2021b. *Fusing context into knowledge graph for commonsense question answering*. In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021*, pages 1201–1207. Association for Computational Linguistics. [Page 8]

- Liu Yang, Mingyang Zhang, Cheng Li, Michael Bendersky, and Marc Najork. 2020a. [Beyond 512 tokens: Siamese multi-depth transformer-based hierarchical encoder for long-form document matching](#). In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*, pages 1725–1734. ACM. [Pages 3 and 8]
- Yinfei Yang, Daniel Cer, Amin Ahmad, Mandy Guo, Jax Law, Noah Constant, Gustavo Hernández Ábrego, Steve Yuan, Chris Tar, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2020b. [Multilingual universal sentence encoder for semantic retrieval](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 87–94. Association for Computational Linguistics. [Page 8]
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alexander J. Smola, and Eduard H. Hovy. 2016. [Hierarchical attention networks for document classification](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1480–1489. The Association for Computational Linguistics. [Page 7]
- Andrew Yates, Rodrigo Nogueira, and Jimmy Lin. 2021. [Pretrained transformers for text ranking: BERT and beyond](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2666–2668. ACM. [Page 8]
- Donghan Yu, Chenguang Zhu, Yiming Yang, and Michael Zeng. 2022. [JAKET: joint pre-training of knowledge graph and language understanding](#). In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, pages 11630–11638. AAAI Press. [Page 8]
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. [Big bird: Transformers for longer sequences](#). *Advances in Neural Information Processing Systems*, 33:17283–17297. [Page 8]
- Xingxing Zhang, Furu Wei, and Ming Zhou. 2019. [HIB-ERT: document level pre-training of hierarchical bidirectional transformers for document summarization](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 5059–5069. Association for Computational Linguistics. [Pages 3 and 8]
- Yue Zhang, Chengcheng Hu, Yuqi Liu, Hui Fang, and Jimmy Lin. 2021a. [Learning to rank in the age of muppets: Effectiveness-efficiency tradeoffs in multi-stage ranking](#). In *Proceedings of the Second Workshop on Simple and Efficient Natural Language Processing*, pages 64–73. Association for Computational Linguistics. [Page 6]
- Yufeng Zhang, Jinghao Zhang, Zeyu Cui, Shu Wu, and Liang Wang. 2021b. [A graph-based relevance matching model for ad-hoc retrieval](#). In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, pages 4688–4696. AAAI Press. [Page 8]
- Yuhao Zhang, Peng Qi, and Christopher D. Manning. 2018. [Graph convolution over pruned dependency trees improves relation extraction](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2205–2215. Association for Computational Linguistics. [Page 8]
- Haoxi Zhong, Chaojun Xiao, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2020. [How does NLP benefit legal system: A summary of legal artificial intelligence](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5218–5230. Association for Computational Linguistics. [Pages 1 and 8]

Appendix

A Preliminary Studies

A.1 How important is the structure of law for statutory article retrieval?

To better understand the reasoning skilled humans follow when retrieving statutory articles, we ask several legal experts to retrieve relevant articles to questions sampled from the Belgian Statutory Article Retrieval Dataset (Louis and Spanakis, 2022, BSARD). We deliberately choose experts unfamiliar with Belgian law and thus have no past knowledge of the location of articles covering a particular subject. In what follows, we summarize the approach these experts use.

First, they determine whether the issue involves either *public* or *private* law. To distinguish between the two, it is necessary to identify to whom the rules apply (i.e., the parties involved in the issue that hold the rights or duties). Generally speaking, public law deals with issues that affect the general public or state (i.e., society as a whole), whereas private law deals with issues that affect individuals, families, and businesses. This first step allows the experts to make an initial selection among the codes of law, which generally relate to only one of either public or private law.

Next, the experts refine their search by determining the field of law (e.g., contract law), followed by the sub-field (e.g., tenant law), and so on, until a set of potentially relevant codes is created in the question’s domains. The experts then focus on the table of contents of one of the selected codes and undertake a hierarchical search that starts from the book’s headings and progressively extends to its titles, chapters, and sections. This step makes it possible to filter out many irrelevant articles by

analyzing the connection between the question’s subject and the different sections’ headings.

Finally, the experts explore the articles within the sections deemed potentially relevant to the question in search of the expected answer. If the experts realize that the chosen direction is a dead end, they return to the previous higher level of the structure, choose another potentially relevant direction, and narrow their search from there.

From this study, we conclude that legal experts rely heavily on the structure of law when retrieving relevant articles to a legal question, which indicates that the different divisions’ headings carry valuable information that retrieval systems should consider.

A.2 How related are neighboring articles in statute law?

In statute law, the sense of a given article is not necessarily self-contained by itself but instead spans across different articles from the same or even different sections. To confirm this, we study to what extent consecutive articles (as they appear in the statute books) address similar subjects.

We consider the Belgian Civil Code, which is the book whose articles are most cited in BSARD, and randomly sample sets of 200 consecutive articles out of it. We then normalize the articles by lowercasing, lemmatizing, and removing stop-words, punctuation, and numbers. Finally, we compute the cosine similarities between the TF-IDF representations of all articles from a given set. Figure 3 shows a heatmap of article similarities for such a set. We see that consecutive articles do indeed cover similar topics, suggesting that the information in a given article is likely to be complementary to that in its neighboring articles. Therefore, we assume that neighboring articles should be considered to capture an article’s whole meaning.

B G-DSR Hyperparameter Tuning

We conduct hyperparameter tuning using Bayes search based on performance on BSARD development set, measured with the macro-averaged $R@200$. Due to limited computational resources, we train our models on BSARD training set only – which takes approximately 1 hour and 15 minutes for DSR and around 5 minutes for LGE – and use the constrained search spaces described below.

DSR grid search space:

- batch size: {8, 16, 24, 32}
- learning rate: {5e-5, 4e-5, 3e-5, 2e-5, 1e-5}




Figure 3: Cosine similarities between TF-IDF representations of 200 consecutive articles from the Belgian Civil Code, taken from the Belgian Statutory Article Retrieval Dataset (Louis and Spanakis, 2022, BSARD).

- weight decay: {0, 0.1, 0.01, 0.001}
- max chunk length: {64, 128, 256, 512}
- max document length: {1024, 2048}
- pooling strategy: {mean, max}
- similarity function: {dot-product, cosine, L2}
- temperature: {0, 0.1, 0.01, 0.001}

LGE grid search space:

- batch size: {8, 16, 32, 64, 128, 256, 512}
- learning rate: {2e-2, 2e-3, 2e-4, 2e-5, 2e-6}
- weight decay: {0, 0.1, 0.01, 0.001}
- #layers: {1, 2, 3, 4}

In total, we run 100 hyperparameter search trials for both DSR and LGE. The optimal hyperparameters, shown in Table 5, are used to re-train the models combining both train and development sets for a final evaluation on the test set.

C BM25 Hyperparameter Tuning

Following Chalkidis et al. (2021), who show that BM25 performance is highly dependent on adequately choosing the (k_1, b) values for the task at hand, we perform a hyperparameter grid search on BSARD development set and plot the results in Figure 4. We observe that, in the case of SAR, the best performance is obtained with $k_1 = 2.5$ and $b = 0.2$. Therefore, we use these values for the final evaluation on BSARD test set.

Hyperparam	DSR	LGE
Model		
Maximum Chunk Length	128	-
Maximum Document Length	1024	-
Pooling Strategy	max	-
#Layers	-	3
Loss		
Temperature	0.01	0.01
Similarity	cos	cos
Training		
Batch Size	24	512
Weight Decay	0.01	0.1
Max Epochs	15	10
Peak Learning Rate	2e-5	2e-4
Learning Rate Decay	Linear	Constant
Warmup ratio	0.05	0.0
AdamW ϵ	1e-7	1e-7
AdamW β_1	0.9	0.9
AdamW β_2	0.999	0.999
Gradient Clipping	1.0	1.0

Table 5: Training hyperparameters for DSR and LGE.

D Evaluation Metrics

Let $\text{rel}_q(a) \in \{0, 1\}$ be the binary relevance label of article a for question q , and $\langle i, a \rangle \in \mathcal{F}_q$ a result tuple (article a at rank i) from the filter set $\mathcal{F}_q \subset \mathcal{C}$ of ranked articles retrieved for question q .

Recall. The *recall* is the fraction of relevant articles retrieved for query q w.r.t. the total number of relevant articles in the corpus \mathcal{C} , i.e.,

$$R_q = \frac{\sum_{\langle i, a \rangle \in \mathcal{F}_q} \text{rel}_q(a)}{\sum_{a \in \mathcal{C}} \text{rel}_q(a)}. \quad (9)$$

When computed for a filter set of size $k = |\mathcal{F}_q| \ll |\mathcal{C}|$, i.e., at a certain cutoff and not on the entire list of articles in \mathcal{C} , we report the metrics with the suffix “@ k ”.

R-Precision The R-Precision is the proportion of the top- R retrieved articles that are relevant to query q , where R is the total number of relevant articles for q , i.e.,

$$\text{RP}_q = \frac{\sum_{\langle i, a \rangle \in \{\mathcal{F}_q\}_{i=1}^R} \text{rel}_q(a)}{R}. \quad (10)$$

Average Precision. The *average precision* is the mean of the precision value obtained after each relevant article is retrieved, that is

$$\text{AP}_q = \frac{\sum_{\langle i, a \rangle \in \mathcal{F}_q} P_{q,i} \times \text{rel}_q(a)}{\sum_{a \in \mathcal{C}} \text{rel}_q(a)}, \quad (11)$$




Figure 4: Heatmap showing BM25 results on BSARD Dev set for different values of k_1 and b .

Loss	Similarity	R@100 (\uparrow)	R@200 (\uparrow)	R@500 (\uparrow)
Cross-entropy	Cosine	75.0	80.7	86.7
	Dot	21.6	37.1	56.0
	Euclidean	43.9	58.1	71.6
Triplet	Cosine	5.4	9.9	17.2
	Dot	4.3	6.4	10.3
	Euclidean	9.0	14.3	25.2

Table 6: BSARD Dev results of DSR trained using different similarity and loss functions.

where $P_{q,j}$ is the *precision* computed at rank j for query q , i.e., the fraction of relevant articles retrieved for query q w.r.t. the total number of articles in the retrieved set $\{\mathcal{F}_q\}_{i=1}^j$:

$$P_{q,j} = \frac{\sum_{\langle i, a \rangle \in \{\mathcal{F}_q\}_{i=1}^j} \text{rel}_q(a)}{|\{\mathcal{F}_q\}_{i=1}^j|}. \quad (12)$$

We report the macro-averaged *recall* at various cut-offs (R@ k), *mean Average Precision* (mAP), and *mean R-Precision* (mRP), which are the average values over a set of n queries.

E Ablation Details

Besides cosine similarity and negative log-likelihood (NLL) loss, we also test the dot-product and Euclidean (inverse of distance is taken as similarity measure) as well as the triplet loss. The temperature for the NLL loss is set to 0.01, and the margin value of the triplet loss is set to 1. We report the results on BSARD development set in Table 6. For a fair comparison, all models are trained for 15 epochs with a batch size of 24, weight decay of 0.01, warm-up proportion of 0.05, an initial learning rate of 2e-5, and a linear decay learning rate schedule.