

# LeCoPCR: Legal Concept-guided Prior Case Retrieval for European Court of Human Rights cases

Santosh T.Y.S.S, Isaac Misael Olguín Nolasco, Matthias Grabmair

School of Computation, Information, and Technology;

Technical University of Munich, Germany

## Abstract

Prior case retrieval (PCR) is crucial for legal practitioners to find relevant precedent cases given the facts of a query case. Existing approaches often overlook the underlying semantic intent in determining relevance with respect to the query case. In this work, we propose LeCoPCR, a novel approach that explicitly generate intents in the form of legal concepts from a given query case facts and then augments the query with these concepts to enhance models understanding of semantic intent that dictates relevance. To overcome the unavailability of annotated legal concepts, we employ a weak supervision approach to extract key legal concepts from the reasoning section using Determinantal Point Process (DPP) to balance quality and diversity. Experimental results on the ECtHR-PCR dataset demonstrate the effectiveness of leveraging legal concepts and DPP-based key concept extraction.

## 1 Introduction

Prior Case Retrieval (PCR) involves finding relevant precedent cases for an undecided case query, assisting legal practitioners in citing pertinent precedents to establish applicable laws and craft their arguments (Turtle, 1995). The exponential growth of legal cases has made navigating the corpus for relevant prior cases increasingly challenging, attracting significant attention from both the legal and IR communities (Locke and Zucco, 2022). Compared to traditional text retrieval, PCR poses unique challenges due to longer, more complex queries, and the abstract nature of relevance in legal tasks (Van Opijnen and Santos, 2017).

Several PCR datasets have been constructed recently, such as COLIEE (Kano et al., 2019), IRLed (Mandal et al., 2017), IL-PCR (Joshi et al., 2023) in English from Canada and India, CAIL-SCM (Xiao et al., 2019) and LeCaRD (Ma et al., 2021) from China in Chinese. However, Santosh et al.

(2024) pointed out that the existing PCR datasets from common law systems do not simulate a realistic setting, as their queries use complete case documents while only masking references to prior cases, resulting in queries containing legal reasoning which is not yet available for an undecided case. To address this, they curated a new PCR dataset based on European Court of Human Rights (ECtHR) jurisdiction, which explicitly separate facts from reasoning and uses only facts of the case as query, treating cited cases in the reasoning section of the query case as relevant precedents.

Earlier methods for PCR relied on expert knowledge based methods with ontologies (Saravanan et al., 2009), thesaurus (CA and KLEIN, 2006), legal issue decomposition (Zeng et al., 2005), citation networks (López Monroy et al., 2013; Minocha et al., 2015). Recent years have seen the emergence of deep learning approaches with text representations (Shao et al., 2020; Rossi and Kanoulas, 2019), domain-aware pre-training (Xiao et al., 2021), incorporating structural information (Li et al., 2023; Ma et al., 2023), summarization to account for longer length (Askari et al., 2021; Tran et al., 2019; Zhou et al., 2023), hybrid methods combining text and citation networks (Bhattacharya et al., 2022). A key challenge in PCR lies in determining the relevance of a prior case, which typically is to substantiate arguments based on shared legal principles or norms and to strengthen the reasoning in the query case. Capturing this semantic intent—the conceptual relationship between the prior and query cases—is essential for accurate retrieval. However, existing methods often rely on the assumption that models can implicitly infer this intent from data, leading to suboptimal performance.

Identifying relevant prior cases to be cited in the query case requires forward-looking anticipation. Users identify the legal principles central to their arguments and search for prior cases that reinforce them. The most authoritative cases, which

discuss or clarify the required legal concepts, are then selected for citation. While anticipatory retrieval aims to maximize recall by identifying all potentially relevant cases, determining authoritativeness ensures precision by focusing on the most impactful cases.

In this work, we focus on improving recall in prior case retrieval by explicitly capturing the semantic intent based on specific legal concepts or norms that dictate the relevance of a case to the query. Since the reasoning which provides this semantic intent is not explicitly available at inference time for the query case, unlike prior methods that implicitly expect the model to learn this intent, we make the model deduce the reasoning required for the query case and leverage that reasoning to identify relevant prior cases. However, generating complete reasoning poses two challenges: (i) the lengthy reasoning sections in legal cases exceed the length limitations of current seq2seq models, and (ii) much of the detailed argumentation is irrelevant to the task of citing a specific case, acting as noise. To address this, we hypothesize that condensing the reasoning into key concepts provides a stronger signal for relevance. Therefore, we focus on training a seq2seq model to generate these key concepts from the query case facts as input.

To address the lack of annotated data for legal concepts, we adopt a weak supervision approach to extract key concepts from the reasoning section. We use Determinantal Point Process (DPP) (Borodin, 2009) to select key concepts from a list of candidates, balancing quality and diversity. DPP has proven effective in various diversity promoting tasks, such as recommendation (Wilhelm et al., 2018), answer retrieval (Nandigam et al., 2022), text summarization (Cho et al., 2019a,b; Li et al., 2019), video summarization (Zhang et al., 2016; Sharghi et al., 2017). These concepts are then used to train a seq2seq model to generate them from given query case facts during inference. We then augment the query with these generated concepts to retrieve similar precedent cases, akin to query expansion method (Rocchio Jr, 1971; Lavrenko and Croft, 2001; Liu et al., 2022). Our experiments on the ECtHR-PCR dataset demonstrate the effectiveness of explicitly leveraging legal concepts to enhance the model’s understanding of relevance. Further, DPP-based concept extraction outperforms other approaches in retrieval performance validated through oracle ablations, attributed to the selection of representative diverse concepts.

## 2 Our Method: LeCoPCR

LeCoPCR initially generates legal concepts from the factual description of a query case and then uses these concepts along with the query to retrieve relevant prior cases. LeCoPCR enhances retrieval by deducing necessary legal concepts required to establish precedence and provide argumentation for the query case. These legal concepts serve as the intents behind citing a specific prior case in the reasoning section of a query case.

### 2.1 Legal Concept Generation

We use a seq2seq model to generate legal concepts from case facts descriptions. Due to the lack of annotated datasets with legal concepts, we adopt a weak supervision method, wherein we extract key concepts from the reasoning section of cases as silver data for training. It’s worth noting that the query comprises only factual descriptions, excluding the reasoning section, which is unavailable prior to the verdict. To extract key legal concepts from the reasoning section, we first identify candidate concepts expressed as noun phrases using part-of-speech tagging and noun phrase chunking with NLTK package (Bird et al., 2009) using the regular expression  $\langle \text{NN} . * | \text{JJ} \rangle * \langle \text{NN} . * \rangle$  from Bennani-Smires et al. (2018). Subsequently, we use determinantal point process (DPP) (Borodin, 2009) to select a subset from these candidate concepts, ensuring a balance between relevance and diversity. *DPP for key concept selection:* Let  $S = 1, \dots, |S|$  denote a finite set of extracted candidate concepts. Formally, DPP defines a probability distribution over an exponential number of sets (all  $2^{|S|}$  subsets of  $S$ ), parameterized by a single  $|S| \times |S|$  positive semi-definite kernel matrix, denoted as  $L$ . If  $k$  is a random set of elements drawn from  $S$ , then the probability of selecting that subset is given by determinants of sub-matrices of  $L$ :

$$p(k; L) = \frac{\det(L_k)}{\det(L+I)} \quad \because \sum_k \det(L_k) = \det(L+I)$$

Here,  $\det(\cdot)$  is the determinant of a matrix,  $I$  is the identity matrix and  $L_k$  is a sub-matrix of  $L$  containing only entries indexed by elements of  $k \subseteq S$ . Kulesza et al. (2012) provide a decomposition of the  $L$ -ensemble matrix as a gram matrix, allowing the modeling of relevance and dissimilarity independently and combining them into a single unified formulation with  $L_{ij} = q_i \cdot s_{ij} \cdot q_j$ , where  $q_i \in \mathbb{R}^+$  is a positive real number indicating the

quality/relevance of concept candidate  $i$ , and  $s_{ij}$  captures the similarity between concepts  $i$  and  $j$ .

To understand why  $\det(L_k)$  serves as a balanced measure of quality and diversity for a selected set, consider a subset  $Y = \{i, j\}$  of elements. The probability of choosing this subset is given as:

$$\begin{aligned} P(Y; L) &\propto \det(L_Y) \\ &= \begin{bmatrix} q_i \cdot s_{ii} \cdot q_i & q_j \cdot s_{ij} \cdot q_j \\ q_j \cdot s_{ji} \cdot q_i & q_j \cdot s_{jj} \cdot q_j \end{bmatrix} \\ &= q_i^2 \cdot q_j^2 \cdot (1 - s_{ij}^2) \end{aligned}$$

If candidate elements are highly relevant, any subset containing them will have a high probability. Conversely, if two candidate elements are similar, any set containing both will have a low probability. Geometrically, this can be interpreted as the squared volume of the space spanned by candidate concept vectors of  $Y$ , where quality indicates vector length and similarity represents the angle between vectors. This determinant expression turns more complex for larger matrices but a similar intuition holds there. In our case, considering each element as a candidate concept, the final subset achieving the highest probability will include a set of highly relevant concepts while maintaining diversity among them via pairwise repulsion.

*Deriving similarity and relevance values:* We use LegalBERT (Chalkidis et al., 2020) to obtain the [CLS] representation of each concept and compute their similarities using cosine similarity. Inspired by Zhang et al. (2022), we avoid computing relevance as cosine similarity directly between a candidate concept and its source paragraph due to potential sequence length mismatch. Instead, we adopt a paragraph-paragraph relevance computation approach. Here, we mask the candidate concept from its source paragraph and compute the similarity score between the masked paragraph and the normal paragraph using LegalBERT and take the complement of this score as the relevance score. This approach assumes that the semantic meaning of the masked document remains largely unchanged if only a trivial concept is masked. Additionally, we consider the position of concepts relative to the citation marker in the reasoning section as indicative of their importance, potentially signaling derivation from the cited case. To incorporate this information, we introduce position regularization  $\rho_i$  for each candidate concept, inspired by Florescu and Caragea (2017) which augments the relevance

score by multiplying with  $\rho_i$ :

$$r_i = \hat{r}_i \cdot \rho_i \quad \text{where} \quad \rho_i = \text{softmax}(e^{1/k})$$

where  $k$  indicates the distance between the  $i$ -th candidate concept and the nearest citation marker. *DPP Inference* The MAP inference for DPP involves sub-modular maximization, which is NP-hard (Ko et al., 1995). Therefore we use greedy algorithm for faster inference (Chen et al., 2018). This algorithm begins with an empty set and iteratively adds one concept to the selected set. The chosen concept  $c$  in each iteration is the one that maximizes the determinant value when added to the current selected set.

$$\begin{aligned} c &= \arg \max_{i \in S-Y} [f(Y \cup \{i\}) - f(Y)] \\ \text{where} \quad f(Y) &= \log \det(L_Y) \end{aligned}$$

## 2.2 Retrieval Model

Once legal concepts are generated for a given query case, we concatenate them with the facts to create an effective query input for the retriever. LeCoPCR, being model-agnostic, is demonstrated using both the lexical-based BM25 (Robertson et al., 2004) and neural-based bi-encoder (Karpukhin et al., 2020) as retrievers. In the bi-encoder approach, we transform the query and candidate cases into dense representations independently and compute relevance scores using the dot product between them. We train the dense model using contrastive loss, which pulls together the query and relevant cases while pushing away irrelevant ones. Irrelevant cases are sampled randomly from the candidate pool which are not relevant to the query. During training, we utilize extracted concepts for query augmentation. However, during inference, we use generated legal concepts, which may exhibit lower quality compared to the extracted ones, leading to exposure bias in the retriever. To mitigate this issue, we employ a hybrid training setup where we augment queries with noisy concepts sampled from other documents, along with the golden extracted concepts to make model robust to noisy generations during inference.

## 3 Experiments

We use the ECtHR-PCR dataset (Santosh et al., 2024), consisting of 15,204 cases, chronologically split into training (9.7k, 1960–2014), development (2.1k, 2015–2017), and test (3.2k, 2018–2022) sets. All the cases preceding the date of query

case serve as candidate cases, resulting in an average of 14,101.2 candidates per test query. On average, there are 12 relevant cases per test query. We report Recall@k and Mean Average Precision (MAP). Recall@k measures the proportion of relevant documents ranked in the top-k candidates, with k values of 50, 100, 500 and 1000. We report the average Recall@k across all instances. MAP calculates the mean of the Average Precision scores for each instance, where Average Precision is the average of Precision@k scores for every rank position of each relevant document. Higher scores indicate better performance. We employ LongT5 (Guo et al., 2022) for legal concept generation and Longformer (Beltagy et al., 2020) for dense retriever, accounting for longer case documents. Detailed implementation details are provided in App. A.

### 3.1 Results

From Table 1, it’s evident that LeCoPCR enhances performance for both BM25 and Longformer, suggesting that legal concepts serve as a bridge to learn relevance signals. Additionally, we present oracle values, which use the extracted concepts from the reasoning section of the query case, forming the upper bound. While the oracle concepts exhibit significant potential for improvement, the modest enhancement in LeCoPCR can be attributed to the inclusion of noisy concepts generated during inference, which may provide incorrect relevance signals. Training with noisy samples under the hybrid setup (HT) with Longformer yields better performance than LeCoPCR, indicating that exposure to concepts of varying quality enhances the model’s robustness. These results also indicate scope for further improvement in concept generation models.

	R50	R100	R500	R1K	MAP
BM25	21.84	27.49	48.62	60.88	9.54
+ LeCoPCR	22.52	28.42	49.86	62.22	9.78
+ Oracle	31.26	38.39	60.79	69.05	14.19
Longformer	23.98	33.97	63.41	76.10	11.59
+ LeCoPCR	25.22	35.16	64.86	78.23	12.46
+ LeCoPCR-HT	26.47	38.62	67.14	79.39	13.68
+ Oracle	32.89	44.20	72.68	82.66	16.22

Table 1: Performance on ECHR-PCR dataset. RK, HT indicate Recall@k and Hybrid training respectively.

**Ablation on Concept Extraction:** We compare our DPP-based method to different legal concept extraction techniques such as word-based TF-IDF and phrase-based methods like TextRank (Mihalcea and Tarau, 2004), KeyBERT (Grootendorst,

2020), and MDERank (Zhang et al., 2022). We ablate DPP method by removing position regularization and MDE rank-based relevance computation, computing it based on cosine similarity between concept and document. Performance is reported using BM25, where the query is augmented with these extracted oracle concepts in Table 2. Using concepts helps identify relevance signals better, as these legal concepts serve as an intent to cite relevant prior cases. TextRank outperforms TF-IDF, indicating that span-based legal concepts provide better signals. KeyBERT performs better than TextRank, highlighting the importance of semantic relatedness in extracting better concepts. MDERank surpasses KeyBERT, suggesting that document-to-document matching is more effective. Even in our method, with the removal of MDE, performance decreases. Removing position regularization also reduces performance, supporting our hypothesis that concepts closer to citation markers are crucial.

	R50	R100	R500	R1K	MAP
-	21.84	27.49	48.62	60.88	9.54
TF-IDF	22.93	28.50	50.59	61.94	11.04
TextRank	23.46	29.83	51.26	62.24	11.36
KeyBERT	24.25	31.43	52.36	63.18	11.52
MDERank	26.05	32.34	53.11	64.75	11.69
DPP	31.26	38.39	60.79	69.05	14.19
w/o MDE	26.39	34.27	55.58	66.72	12.75
w/o Position	29.42	35.29	58.26	68.18	13.22

Table 2: Performance with Oracle concepts extracted from different methods, using BM25 as retriever.

### Ablation on legal concept generation models :

We compare different long-context models such as LongT5 (Guo et al., 2022), SLED-BART (Ivgi et al., 2023), and LED (Beltagy et al., 2020). We calculate the coverage of the generated concepts with respect to the reasoning section at both the word-level and concept (phrase) level in Tab. 3. We notice LongT5 outperforming LED, which is semi-pretrained and initialized by repeatedly copying BART without undergoing end-to-end pre-training for longer sequences, thereby hindering its performance in long-context scenarios. Similarly, SLED employs a short-range pre-trained model like BART and uses a chunking approach on the encoder side to handle longer lengths, underscoring the necessity for long-context pre-training to effectively capture context for this task.



	Word Cov.	Concept Cov.
LongT5	51.56	39.24
SLED-BART	46.12	34.46
LED	44.29	30.29

Table 3: Coverage of generated concepts with respect to reasoning section for different generation models.

### 3.2 Case Study

**Example 1: The Case of O.P. v. THE REPUBLIC OF MOLDOVA<sup>1</sup>:** Some of the key concepts extracted by our DPP framework are *speediness of the requisite judicial controls, reasonableness of the suspicion, speedy judicial decision, lawfulness of detention, reasonable will, deprivation of liberty, non-pecuniary damage*. The concepts generated by LongT5 include *right to liberty and security, lawful arrest or detention, reasonable suspicion, speediness of review*. These extracted and generated concepts certainly capture semantic intents which can facilitate to retrieve relevant prior cases for this query, demonstrating the effectiveness of LeCoPCR.

**Example 2: The Case of BANCŞÓK AND LÁSZLÓ MAGYAR v. HUNGARY<sup>2</sup>:** The extracted key concepts are *inhuman or degrading treatment or punishment, constitutional complaint, effective remedy, unprecedented delay, potential effectiveness of the remedy, punitive element of punishment, de facto reducible, State’s margin of appreciation, general pardon procedure, apply for parole, prospect of release, progress towards rehabilitation, legitimate penological grounds, non-pecuniary damage*. While the generated ones are *Life imprisonment, degrading punishment, periodic review, penological grounds, compassionate grounds, conditional release of the prisoner, medical assistance, rehabilitation*. These both concept lists reaffirm our hypothesis of LeCoPCR to boost retrieval performance.

## 4 Conclusion

We introduced LeCoPCR, a novel approach for PCR that leverages legal concepts to enhance the model’s understanding of semantic intent that dictates relevance. We addressed the challenge of unavailability of annotated legal concepts by employing a weak supervision approach to extract key concepts from the reasoning section using DPP with

masked-document-to-document matching and citation proximity, for quality assessment. Our results on the ECtHR-PCR dataset validate LeCoPCR’s effectiveness. While our approach tackles relevance intent at a coarse-grained level for each query case, future research could explore fine-grained intent analysis for each cited case within a query case.

### Limitations

We assess LeCoPCR using the ECtHR-PCR dataset, which focuses on European Court of Human Rights judgments. While our approach may be generally applicable to other legal jurisdictions, its suitability across diverse legal systems, citation practices and semantic complexities warrants further investigation. Additionally, while LeCoPCR enhances retrieval explainability through generated legal concepts, our current evaluation, reliant on exact match-based coverage of the reasoning section, may not fully capture semantic nuances. Hence, a thorough human study is necessary to comprehensively gauge LeCoPCR’s explainability.

A limitation of LeCoPCR is its exclusive focus on textual content, overlooking potential insights from the citation network’s interconnected relationships, which offer a broader global perspective on case law. Future research should explore integrating citation structures to enhance retrieval performance. Another challenge lies in the temporal evolution of case law, reflecting changes in norms, societal attitudes, and the temporal nature of precedents, which may be overruled, thereby impacting the law’s scope. Consequently, future approaches should account for the temporal aspects of precedents to model relevance effectively.

It’s crucial to recognize that PCR datasets are constructed based on cases cited in the reasoning section of the query case, potentially leading to biases such as selective citation or the omission of relevant precedents. Thus, the effectiveness of LeCoPCR could be influenced by the dataset’s representation, necessitating human evaluation to gauge the utility of retrieved precedents. However, conducting such evaluations poses challenges due to the lengthy and complex legal text, requiring annotators with deep expertise in ECHR jurisprudence, thus making it resource-intensive.

Furthermore, while the models explored in this study prioritize recall by acting as pre-fetchers to ensure retrieving all relevant cases, end-users typically expect a high-precision retrieval system that

<sup>1</sup><https://hudoc.echr.coe.int/?i=001-212690>

<sup>2</sup><https://hudoc.echr.coe.int/?i=001-212669>

precisely identifies a smaller set of relevant cases. Achieving this requires an additional re-ranking step in the retrieval pipeline to optimize the precision of the ranked list. In this study, we focus on the initial retrieval step, leaving the development of a re-ranker component for future research.

## Ethics Statement

Our experiments utilize a dataset of ECHR decisions, which is publicly available and are sourced from the public court database HUDOC<sup>3</sup>, which includes real names and lacks anonymization. While there are concerns about the non-anonymized nature of this information, we do not foresee any significant harm beyond the disclosure. We believe that developing effective PCR systems is vital for supporting legal professionals in managing the growing caseload, underscoring the need for research to prioritize enhancing legal services and democratizing law (Tsarapatsanis and Aletras, 2021). However, it is crucial to address various shortcomings and ensure the responsible and ethical deployment of legal-oriented technologies.

Furthermore, it is important to recognize that using historical data to train retrieval models may introduce biases, and leveraging pre-trained encoders can inherit biases encoded within them. Ensuring that PCR models intended for practical deployment do not inadvertently perpetuate or exacerbate existing biases in the legal system, such as racial or gender bias, is imperative. These models must undergo scrutiny against applicable equal treatment imperatives regarding their performance, behavior, and intended use, emphasizing the significance of fairness and accountability in the development and deployment of legal AI technologies.

## References

- Arian Askari, Suzan Verberne, O Alonso, S Marchesin, M Najork, and G Silvello. 2021. Combining lexical and neural retrieval with longformer-based summarization for effective case law retrieval. In *DESIRES*, pages 162–170.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Kamil Bennani-Smires, Claudiu Musat, Andreea Hossmann, Michael Baeriswyl, and Martin Jaggi. 2018. Simple unsupervised keyphrase extraction using sentence embeddings. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 221–229.
- Paheli Bhattacharya, Kripabandhu Ghosh, Arindam Pal, and Saptarshi Ghosh. 2022. Legal case document similarity: You need both network and text. *Information Processing & Management*, 59(6):103069.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc."
- Alexei Borodin. 2009. Determinantal point processes. *arXiv preprint arXiv:0911.1153*.
- Michel CA and Wouter VAN KLEIN. 2006. Thesaurus-based retrieval of case law. In *Legal Knowledge and Information Systems: JURIX 2006: the Nineteenth Annual Conference*, volume 152, page 61. IOS Press.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. Legal-bert: The muppets straight out of law school. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2898–2904.
- Laming Chen, Guoxin Zhang, and Eric Zhou. 2018. Fast greedy map inference for determinantal point process to improve recommendation diversity. *Advances in Neural Information Processing Systems*, 31.
- Sangwoo Cho, Logan Lebanoff, Hassan Foroosh, and Fei Liu. 2019a. Improving the similarity measure of determinantal point processes for extractive multi-document summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1027–1038.
- Sangwoo Cho, Chen Li, Dong Yu, Hassan Foroosh, and Fei Liu. 2019b. Multi-document summarization with determinantal point processes and contextualized representations. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 98–103.
- Corina Florescu and Cornelia Caragea. 2017. A position-biased pagerank algorithm for keyphrase extraction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Maarten Grootendorst. 2020. Keybert: Minimal keyword extraction with bert.
- Mandy Guo, Joshua Ainslie, David C Uthus, Santiago Ontanon, Jianmo Ni, Yun-Hsuan Sung, and Yinfei Yang. 2022. Longt5: Efficient text-to-text transformer for long sequences. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 724–736.
- Maor Ivgi, Uri Shaham, and Jonathan Berant. 2023. Efficient long-text understanding with short-text models. *Transactions of the Association for Computational Linguistics*, 11:284–299.

<sup>3</sup><https://hudoc.echr.coe.int>

- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547.
- Abhinav Joshi, Akshat Sharma, Sai Kiran Tanikella, and Ashutosh Modi. 2023. U-creat: Unsupervised case retrieval using events extraction. *arXiv preprint arXiv:2307.05260*.
- Yoshinobu Kano, Mi-Young Kim, Masaharu Yoshioka, Yao Lu, Juliano Rabelo, Naoki Kiyota, Randy Goebel, and Ken Satoh. 2019. Coliee-2018: Evaluation of the competition on legal information extraction and entailment. In *New Frontiers in Artificial Intelligence: JSAI-isAI 2018 Workshops, JURISIN, AIBiz, SKL, LENLS, IDAA, Yokohama, Japan, November 12–14, 2018, Revised Selected Papers*, pages 177–192. Springer.
- 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.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Chun-Wa Ko, Jon Lee, and Maurice Queyranne. 1995. An exact algorithm for maximum entropy sampling. *Operations Research*, 43(4):684–691.
- Alex Kulesza, Ben Taskar, et al. 2012. Determinantal point processes for machine learning. *Foundations and Trends® in Machine Learning*, 5(2–3):123–286.
- Victor Lavrenko and W Bruce Croft. 2001. Relevance-based language models.
- Haitao Li, Qingyao Ai, Jia Chen, Qian Dong, Yueyue Wu, Yiqun Liu, Chong Chen, and Qi Tian. 2023. Sailer: structure-aware pre-trained language model for legal case retrieval. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1035–1044.
- Lei Li, Wei Liu, Marina Litvak, Natalia Vanetik, and Zuying Huang. 2019. In conclusion not repetition: Comprehensive abstractive summarization with diversified attention based on determinantal point processes. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 822–832.
- Linqing Liu, Minghan Li, Jimmy Lin, Sebastian Riedel, and Pontus Stenetorp. 2022. Query expansion using contextual clue sampling with language models. *arXiv preprint arXiv:2210.07093*.
- Daniel Locke and Guido Zuccon. 2022. Case law retrieval: problems, methods, challenges and evaluations in the last 20 years. *arXiv preprint arXiv:2202.07209*.
- Alfredo López Monroy, Hiram Calvo, Alexander Gelbukh, and Georgina García Pacheco. 2013. Link analysis for representing and retrieving legal information. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 380–393. Springer.
- Yixiao Ma, Yunqiu Shao, Yueyue Wu, Yiqun Liu, Ruizhe Zhang, Min Zhang, and Shaoping Ma. 2021. Lecard: a legal case retrieval dataset for chinese law system. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pages 2342–2348.
- Yixiao Ma, Yueyue Wu, Qingyao Ai, Yiqun Liu, Yunqiu Shao, Min Zhang, and Shaoping Ma. 2023. Incorporating structural information into legal case retrieval. *ACM Transactions on Information Systems*, 42(2):1–28.
- Arpan Mandal, Kripabandhu Ghosh, Arnab Bhattacharya, Arindam Pal, and Saptarshi Ghosh. 2017. Overview of the fire 2017 ired track: Information retrieval from legal documents. In *FIRE (Working Notes)*, pages 63–68.
- Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, pages 404–411.
- Akshay Minocha, Navjyoti Singh, and Arjit Srivastava. 2015. Finding relevant indian judgments using dispersion of citation network. In *Proceedings of the 24th international conference on World Wide Web*, pages 1085–1088.
- Poojitha Nandigam, Nikhil Rayaprolu, and Manish Shrivastava. 2022. Diverse multi-answer retrieval with determinantal point processes. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2220–2225.
- Stephen Robertson, Hugo Zaragoza, and Michael Taylor. 2004. Simple bm25 extension to multiple weighted fields. In *Proceedings of the thirteenth ACM international conference on information and knowledge management*, pages 42–49.
- Joseph John Rocchio Jr. 1971. Relevance feedback in information retrieval. *The SMART retrieval system: experiments in automatic document processing*.
- Juline Rossi and Evangelos Kanoulas. 2019. Legal information retrieval with generalized language models. *Proceedings of the 6th Competition on Legal Information Extraction/Entailment. KanoE*.
- T.Y.S.S Santosh, Rashid Gustav Haddad, and Matthias Grabmair. 2024. Ecthr-pcr: A dataset for precedent understanding and prior case retrieval in the european court of human rights. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING)*.

- Manavalan Saravanan, Balaraman Ravindran, and Shivani Raman. 2009. Improving legal information retrieval using an ontological framework. *Artificial Intelligence and Law*, 17:101–124.
- Yunqiu Shao, Jiaxin Mao, Yiqun Liu, Weizhi Ma, Ken Satoh, Min Zhang, and Shaoping Ma. 2020. Bert-pli: Modeling paragraph-level interactions for legal case retrieval. In *IJCAI*, pages 3501–3507.
- Aidean Sharghi, Jacob S Laurel, and Boqing Gong. 2017. Query-focused video summarization: Dataset, evaluation, and a memory network based approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4788–4797.
- Vu 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.
- 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.
- Howard Turtle. 1995. Text retrieval in the legal world. *Artificial Intelligence and Law*, 3:5–54.
- Marc Van Opijnen and Cristiana Santos. 2017. On the concept of relevance in legal information retrieval. *Artificial Intelligence and Law*, 25:65–87.
- Mark Wilhelm, Ajith Ramanathan, Alexander Bonomo, Sagar Jain, Ed H Chi, and Jennifer Gillenwater. 2018. Practical diversified recommendations on youtube with determinantal point processes. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 2165–2173.
- Chaojun Xiao, Xueyu Hu, Zhiyuan Liu, Cunchao Tu, and Maosong Sun. 2021. Lawformer: A pre-trained language model for chinese legal long documents. *AI Open*, 2:79–84.
- Chaojun Xiao, Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Zhiyuan Liu, Maosong Sun, Tianyang Zhang, Xianpei Han, Zhen Hu, Heng Wang, et al. 2019. Cail2019-scm: A dataset of similar case matching in legal domain. *arXiv preprint arXiv:1911.08962*.
- Yiming Zeng, Ruili Wang, John Zeleznikow, and Elizabeth Kemp. 2005. Knowledge representation for the intelligent legal case retrieval. In *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*, pages 339–345. Springer.
- Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman. 2016. Video summarization with long short-term memory. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII 14*, pages 766–782. Springer.
- Linhan Zhang, Qian Chen, Wen Wang, Chong Deng, ShiLiang Zhang, Bing Li, Wei Wang, and Xin Cao. 2022. Mderank: A masked document embedding rank approach for unsupervised keyphrase extraction. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 396–409.
- Youchao Zhou, Heyan Huang, and Zhijing Wu. 2023. Boosting legal case retrieval by query content selection with large language models. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, pages 176–184.

## A Implementation Details

For legal concept generation model, we set the maximum input and output sequence length to 4096 and 256. We train with a learning rate of 1e-4 with a scheduler that warms up from zero during the first 10% of the steps and then linearly decays back to zero for the remaining steps. The whole model is trained for 10 epochs. We use four beams and repetition penalty of 2.0. For BM25 retriever, we set the value of K1 and b to 1.2 and 0.75 respectively. For dense retriever, we use the maximum sequence length of 4096 and use 7 negatives per positive. We use learning rate of 1e-4 and is trained for 3 epochs. We use FAISS (Johnson et al., 2019), an open-source library for building efficient datastore construction for similarity search during dense retrieval. All models are optimized end-to-end using Adam (Kingma and Ba, 2014).