

RELexED: Retrieval-Enhanced Legal Summarization with Exemplar Diversity

Santosh T.Y.S.S, Chen Jia, Patrick Goroncy, Matthias Grabmair

School of Computation, Information, and Technology;

Technical University of Munich, Germany

Abstract

This paper addresses the task of legal summarization, which involves distilling complex legal documents into concise, coherent summaries. Current approaches often struggle with content theme deviation and inconsistent writing styles due to their reliance solely on source documents. We propose RELexED, a retrieval-augmented framework that utilizes exemplar summaries along with the source document to guide the model. RELexED employs a two-stage exemplar selection strategy, leveraging a determinantal point process to balance the trade-off between similarity of exemplars to the query and diversity among exemplars, with scores computed via influence functions. Experimental results on two legal summarization datasets demonstrate that RELexED significantly outperforms models that do not utilize exemplars and those that rely solely on similarity-based exemplar selection.

1 Introduction

Legal summarization is an essential task, aimed at distilling complex legal documents, such as court rulings and case judgments, into concise summaries (Farzindar, 2004; Saravanan et al., 2006). These summaries assist legal professionals in quickly grasping the core information, saving time and enhancing decision-making efficiency (Grover et al., 2003a,b). However, summarizing legal texts is uniquely challenging due to the highly structured and formalized nature of legal writing. Legal summaries must maintain factual accuracy, adhere to specific terminologies, and follow established patterns that reflect the reasoning and argumentation typical in legal documents (Bhattacharya et al., 2019, 2021; Deroy et al., 2023).

Current approaches to legal summarization often rely solely on the content of the source document, leading to problems such as content deviation, where generated summaries stray from the

central themes of the case and inconsistent writing style, failing to match the rigor and formality required in legal contexts (Shukla et al., 2022; Moro and Ragazzi, 2022; Elaraby and Litman, 2022; Shen et al., 2022; Santosh et al., 2024b,d). This occurs because existing models are expected to implicitly learn the templatization process during training, which is prevalent in legal writing where they follow specific templates or writing structures. To address this, it is essential to introduce writing templates as exemplars that can bring not only guidance of writing format but also additional background knowledge necessary (Oya et al., 2014; Gao et al., 2019). Rather than manually creating templates for each scenario which turns expensive to obtain from domain experts, one scalable way is to use reference summaries in the training corpus as exemplars (An et al., 2021; Wang et al., 2022).

To the best of our knowledge, there has been no exploration of retrieval-enhanced exemplars for legal summarization task. This approach consists of two modules: Retriever and Summarizer. Given a query document, the Retriever is asked to retrieve the most related exemplars from the training corpus. The second module Summarizer generates the summary not only relying on source document but also the retrieved exemplars. This method bears resemblance to recent works of in-context learning (Brown, 2020; Dong et al., 2022) which appends a few demonstration examples to the query to form a prompt which is then fed into the language model for prediction. In this work, we are mainly interested in supervised learning where the model parameters are fine-tuned to learn from given examples than in-context learning which does not perform parameter updates and expects the model to learn the pattern using the demonstration examples in the prompt. The success of retrieval-enhanced methods in supervised learning demonstrates that, even with vast numbers of parameters, models are unable to memorize every pattern from the train-

ing data and hence by retrieving relevant examples and explicitly providing them to the model can enhance its performance (Wang et al., 2022; Tyss et al., 2024). This approach suggests that rather than constantly scaling models to larger sizes, we can achieve high-quality results by equipping a moderately sized model with training data that closely matches the current instance, thereby not only boosts performance but also significantly reduces computational costs (An et al., 2021).

The main crucial link of retrieval-enhanced summarization are the chosen exemplars. Previous approaches generally rely on retrieving examples based on the similarity between the input and the examples based on traditional lexical-based BM25 (Cao et al., 2018; Wang et al., 2022) or embeddings (An et al., 2021). Despite the improved performance, we posit that similarity based exemplar selection do not always take into account the inter-relationship between different exemplars. For instance, the ignorance of redundancy among examples can result in almost identical examples, providing no additional information to the model.

In this paper, we propose a two-stage exemplar selection strategy for legal summarization, RELexED which balances both the quality and the diversity among the selected exemplars. In the first stage, we rely on semantic similarity to filter out unrelated examples and then in the second stage, instead of selecting each example independently, RELexED captures the inter-relationship between selected examples. To model the joint probability of selected set of examples given a specific input, we leverage determinantal point process (DPP) (Kulesza et al., 2012) that learns to select relevant yet the most diverse example set. To measure the quality or similarity criterion, we resort to recent instance-based explanation methods (Koh and Liang, 2017; Pruthi et al., 2020) which aim to quantify how a training example affects the prediction of a test example after training. Our experiments on two legal summarization datasets, SuperSCOTUS and CivilSum demonstrate the effectiveness of retrieval-enhanced summarization and specifically, our proposed exemplar selection method, RELexED which goes beyond using similarity-based selection only.

2 RELexED: Our Method

We first introduce the framework for retrieval-enhanced summarization and the background of the

Determinantal Point Process (DPP). Then, we introduce our two-stage exemplar selection strategy, RELexED, balancing quality and diversity.

2.1 Preliminaries

Retrieval-enhanced Summarization First, we index the training corpus of document-summary pairs into a list of key-value pairs where the document serves as the key and its summary as the value. Given the input x , the retriever returns the k exemplars from the training corpus. Then, these retrieved results are combined with the input x to feed into the summarization model to generate the summary. Given that legal documents tend to be lengthy, we only incorporate the values (summaries) from the exemplars alongside the query input to the summarization model.

DPP DPP is a probabilistic model to express interactions between items and could help us to select a representative subset while keeping high diversity among different items (Kulesza et al., 2012; Chen et al., 2018). Formally, let $S = \{1, \dots, |S|\}$ denote a finite set of items and 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-matrix of L :

$$p(k; L) = \frac{\det(L_k)}{\det(L+I)} \quad \sum_k \det(L_k) = \det(L+I) \quad (1)$$

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 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 is a positive real number indicating the quality/relevance of item i , and s_{ij} captures the similarity between items i and j . To understand why DPP serve as a trade-off between quality and diversity, please refer to the detailed explanation in App. A. The MAP inference for DPP involves sub-modular maximization, which is NP-hard. Therefore we use greedy algorithm for faster inference (Chen et al., 2018). It begins with an empty set and iteratively adds each item to the selected set Y . The chosen item i in each iteration is the one that maximizes the determinant value

when added to the current selected set.

$$c = \arg \max_{i \in S-Y} [f(Y \cup \{i\}) - f(Y)]$$

where $f(Y) = \log \det(L_Y)$

2.2 Two-stage Exemplar Selection

In the first stage, we leverage semantic similarity and utilize the BM25 to retrieve the top k_1 documents similar to the query from the training corpus. This step helps to eliminate irrelevant training instances that is not related to the given input and also aids to reduce the size of the candidate set, which simplifies the subsequent computations of the DPP matrix elements. The second stage involves employing DPP to select the top k exemplars from these k_1 , balancing relevance and diversity. The core question that arises in this stage is how to compute the quality/relevance score for each individual item with respect to the query item and the similarity score for pairs of items.

We leverage influence functions, which aims to trace a model’s predictions back to the most responsible training examples (Koh and Liang, 2017). For a model f with parameters θ and loss function $l(f_\theta, \cdot)$, the gradient $g(\theta, \cdot)$ for a sample z is given by $g(f_\theta, z) = \nabla l(f_\theta, z)$. Pruthi et al. (2020) provides TracIn which is a gradient-only alternative to influence function approximation and the influence of a sample z on sample z' is given by

$$\text{TracIn}(f_\theta, z, z') = g(f_\theta, z) \cdot g(f_\theta, z')$$

We use an auxiliary model (similar to summarization model) fine-tuned to generate a summary directly from the input alone without exemplars. Given computational expensive nature to compute these, we follow Thakkar et al. (2023) to leverage the layer agnostic nature of TracIn and use only first layer of the encoder to compute the influence score. For quality/relevance score q_i we compute the influence score of item i on the query item x . For similarity score s_{ij} , we compute the influence score of item i on item j . Once we populate the kernel matrix L , we greedily select items iteratively till the pre-defined number of exemplars that fit into the summarization model and concatenate them in the same order of selection.

3 Experiments

3.1 Datasets & Metrics

We experiment with two legal summarization datasets. **SuperSCOTUS** (Fang et al., 2023) is a

multi-source dataset of U.S. Supreme Court (SCOTUS) cases. We use the summarization dataset provided, which uses the majority opinion of the court as the input to generate the Syllabus, which provides a summary of case. It consists of 4058 case-summary pairs, split into 3246/406/406 for training, validation, and test with an average length of document and summary as 5405.46 and 902.23 tokens respectively. **CivilSum** (Malik et al., 2024) contains 23350 court decisions from the Supreme Court and High Courts of India paired with human-written summaries. It is split into 21015/1168/1167 for training, validation and test. It has an average document and summary length of 2639.49 and 131.99 tokens respectively.

We evaluate the quality of the generated summaries using ROUGE-1,2,L (Lin, 2004) for lexical overlap with the reference paragraph, BERTScore (Zhang et al., 2019) for semantic similarity between generated and reference summary. For faithfulness, we report AlignScore (Zha et al., 2023) for factual consistency based on a unified alignment function between the input context and generated text. We also report coherence and fluency scores using UniEval metric (Zhong et al., 2022) based on the generated summary.

We use longformer encoder-decoder (Beltagy et al., 2020) as our summarization model to account for longer documents. We use 4 and 8 exemplars for SuperSCOTUS and CivilSum respectively. Implementation details can be found in App. B.

3.2 Results

We use the following models for comparison. The *w/o exemplars* model serves as the baseline, trained without incorporating any exemplars. All other approaches involve exemplar selection. The *BM25* model uses exemplars retrieved based on BM25 semantic similarity. The *BM25 + DPP (BM25)* model implements a two-stage exemplar selection process where diversity is introduced in the second stage using DPP, but both quality and similarity scores are computed with BM25 scores. Our proposed method, ReLexED (BM25 + DPP (IF)), uses influence function scores to compute quality and similarity during the DPP-based selection process.

From Table 1, we observe that incorporating exemplars significantly improves performance across both datasets, particularly in the UniEval metrics of coherence and fluency. Compared to similarity-based (BM25) exemplar selection, we find that introducing diversity through DPP yields additional

	R-1	R-2	R-L	BS	AS	Coh	Flu
SuperSCOTUS							
w/o exemplars	51.17	24.62	28.24	64.12	56.12	69.12	66.20
BM25	51.87	25.18	28.91	65.28	59.74	72.42	69.12
BM25 + DPP (BM25)	53.12	25.86	30.03	65.47	60.16	73.16	71.29
BM25 + DPP (IF)	54.04	26.46	30.84	65.84	61.15	74.86	72.26
CivilSum							
w/o exemplars	39.74	17.29	29.75	57.33	59.92	63.34	74.40
BM25	40.66	18.76	30.57	59.13	62.15	65.13	76.21
BM25 + DPP (BM25)	41.18	19.43	31.15	59.82	62.72	66.15	78.12
BM25 + DPP (IF)	41.68	20.25	32.06	59.68	63.82	66.53	77.75

Table 1: Results on SuperSCOTUS and CivilSum. BS, AS, Coh and Flu denote BERTScore, AlignScore, Coherence and Fluency respectively. Our method **RELexED** refers to BM25 + DPP (IF). RELexED achieves statistically significant improvements over the baseline (w/o exemplars) performance, using the Wilcoxon signed-rank test with a 95% confidence interval.

	SuperSCOTUS		CivilSum	
	EQ	IE	EQ	IE
BM25	0.77	0.81	0.75	0.86
+ DPP (BM25)	0.75	0.79	0.73	0.84
+ DPP (IF)	0.75	0.74	0.72	0.76

Table 2: Comparison of exemplar selection approaches. EQ, IE denotes similarity between Exemplars-Query and Inter-Exemplars.

gains, underscoring the redundancy of using overly similar examples and emphasizing the importance of considering inter-relationships among exemplars. Overall, computing similarity and quality scores using influence scores ensures a more diverse and informative set of exemplars. Gradients offer a more nuanced notion of similarity, as similar examples often result in nearly identical gradients. This is something that BM25-based scores in DPP fail to capture, leading to a less effective selection of diverse examples. Consequently, our influence-based approach captures deeper relationships among examples, improving summarization quality. Overall, we observe huge improvements in both style-based and faithfulness metrics, with notable or comparable gains in lexical-based measures.

Balancing similarity-diversity of Selected Exemplars: We analyse exemplar selection methods by measuring (i) how similar are the summaries of selected exemplars to the reference summary of the query instance (EQ) (ii) how similar are the exemplar summaries among themselves (IE). We report the average cosine similarity computed across each pair of summaries, utilizing embeddings from the LegalBERT model (Chalkidis et al., 2020).

As shown in Tab. 2, incorporating a diversity criterion with DPP, reduces inter-exemplar similarity (IE) promoting diversity. This also reduces similarity between the exemplar summaries and the query summary. But as evidenced in Tab. 1, these diverse exemplars lead to greater performance improvements compared to the model relying solely on exemplars selected via BM25 indicating that highly similar exemplars tend to focus on the same aspects, leading to redundancy and limiting the introduction of new information. In contrast, diversifying exemplars allows to capture a broader range of relevant details. While BM25-based score computation in DPP may show term variation, but they often correspond to similar underlying concepts, resulting in the selection of redundant exemplars. Using influence scores (computed through gradients) enables to reliably compute similar examples and applying a diversity criterion suppresses overly similar exemplars, reflected in lower IE scores in Tab. 2, leading to enhanced performance compared to BM25-based DPP.

3.3 Case Study

In *United States Steel Corp. v. Multistate Tax Commission* from SuperSCOTUS dataset, the case presented diverse legal aspects, including compliance with the Compact Clause, challenges under the Commerce Clause, taxpayer rights under the Fourteenth Amendment, and dissenting concerns regarding federal oversight. However, BM25-based exemplar selection focused narrowly on the Compact Clause, retrieving examples that emphasized the Compact’s non-interference with federal supremacy, while neglecting other key aspects,

such as burdens on interstate commerce and dissenting perspectives. This narrow focus led to summaries that lacked nuance and comprehensive coverage. By contrast, our RELexED framework addressed this limitation through a two-stage exemplar selection strategy that balances relevance and diversity. RELexED retrieved exemplars covering a wider range of legal dimensions, such as the operational impact of the Compact, constitutional concerns, and dissenting opinions. This diversity enriched the summarization process, resulting in summaries that are more coherent, well-rounded, and reflective of the multifaceted nature of the case.

4 Conclusion

In this work, we propose RELexED, a retrieval-augmented framework for legal summarization that introduces a two-stage exemplar selection strategy, accounting for the inter-relationships between exemplars. We leverage determinantal point process to balance relevance and diversity among the retrieved examples, with scores computed using model gradient-based influence functions. Our experiments on two legal datasets demonstrate that RELexED outperforms traditional similarity-based exemplar selection, paving the way for designing more effective methods to incorporate informative exemplars into legal summarization tasks.

Limitations

Our experiments are conducted on two specific legal datasets from different jurisdictions. The extent to which our observations hold true for a broader range of legal systems remains an open question. The legal domain is vast and varied, and different jurisdictions may exhibit unique characteristics that impact the generalizability of summarization models and warrant further investigation for understanding cross-jurisdictional adaptation of these models (Santosh et al., 2024b).

Additionally, our current approach primarily focuses on exemplars based on semantic similarity and semantic diversity. However, it would be beneficial to explore additional features, such as the temporal nature of legal documents (Santosh et al., 2024c) and multi-aspect considerations (Santosh et al., 2024a). Understanding different facets of legal cases, such as procedural history, legal arguments and jurisdictional relevance, could lead to more tailored exemplar selections for each aspect, enhancing the quality of generated summaries.

Our evaluation primarily relies on established summarization metrics which provide a quantitative measure of summarization quality, they may not fully capture the nuanced legal content, context and intricacies essential for legal professionals. This limitation suggests a potential avenue for further research into developing additional legal domain-specific evaluation metrics that can more accurately reflect the complexities inherent in legal documents. Another significant limitation of our study is the absence of direct participation or validation by legal experts in assessing the summarization outputs. We could not conduct expert evaluations due to a lack of access to legal professionals. Engaging legal experts in future evaluations would provide valuable insights into the practical relevance of the summaries.

Ethics Statement

All datasets utilized in this study are publicly available and were sourced from legal repositories. While the case documents included in these datasets are not anonymized, we do not anticipate any harm arising from their availability, as the datasets are intended for research purposes and contribute to advancing legal technology.

However, we acknowledge that the datasets may contain inherent biases that reflect the existing legal system, societal norms, and historical case law trends. These biases can manifest in various forms, such as the overrepresentation or underrepresentation of certain legal issues, demographics, or jurisdictions. Consequently, these biases could influence the training process, potentially leading to summaries that reinforce existing disparities within the legal system. Future work should focus on curating datasets that encompass a wider range of legal contexts, ensuring inclusivity across various jurisdictions and case types. Additionally, ongoing efforts to audit and refine the datasets will be crucial in identifying and correcting biases that may adversely affect the model’s performance. We encourage further research into bias detection and mitigation strategies to enhance the integrity of legal summarization systems.

References

Chenxin An, Ming Zhong, Zhichao Geng, Jianqiang Yang, and Xipeng Qiu. 2021. Retrievalsum: A retrieval enhanced framework for abstractive summarization. *arXiv preprint arXiv:2109.07943*.

- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Paheli Bhattacharya, Kaustubh Hiware, Subham Rajgaria, Nilay Pochhi, Kripabandhu Ghosh, and Saptarshi Ghosh. 2019. A comparative study of summarization algorithms applied to legal case judgments. In *Advances in Information Retrieval: 41st European Conference on IR Research, ECIR 2019, Cologne, Germany, April 14–18, 2019, Proceedings, Part I 41*, pages 413–428. Springer.
- Paheli Bhattacharya, Soham Poddar, Koustav Rudra, Kripabandhu Ghosh, and Saptarshi Ghosh. 2021. Incorporating domain knowledge for extractive summarization of legal case documents. In *Proceedings of the eighteenth international conference on artificial intelligence and law*, pages 22–31.
- Tom B Brown. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Ziqiang Cao, Wenjie Li, Sujian Li, and Furu Wei. 2018. Retrieve, rerank and rewrite: Soft template based neural summarization. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 152–161.
- 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.
- Aniket Deroy, Kripabandhu Ghosh, and Saptarshi Ghosh. 2023. How ready are pre-trained abstractive models and llms for legal case judgement summarization? *arXiv preprint arXiv:2306.01248*.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey on in-context learning. *arXiv preprint arXiv:2301.00234*.
- Mohamed Elaraby and Diane Litman. 2022. Arglegal-sum: Improving abstractive summarization of legal documents with argument mining. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6187–6194.
- Biaoyan Fang, Trevor Cohn, Timothy Baldwin, and Lea Frermann. 2023. Super-scotus: A multi-sourced dataset for the supreme court of the us. In *Proceedings of the Natural Legal Language Processing Workshop 2023*, pages 202–214.
- Atefeh Farzindar. 2004. Atefeh farzindar and guy lapalme, letsum, an automatic legal text summarizing system in t. gordon (ed.), legal knowledge and information systems. jurix 2004: The seventeenth annual conference. amsterdam: Ios press, 2004, pp. 11–18. In *Legal knowledge and information systems: JURIX 2004, the seventeenth annual conference*, volume 120, page 11. IOS Press.
- Shen Gao, Xiuying Chen, Piji Li, Zhangming Chan, Dongyan Zhao, and Rui Yan. 2019. How to write summaries with patterns? learning towards abstractive summarization through prototype editing. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3741–3751.
- Claire Grover, Ben Hachey, Ian Hughson, and Chris Korycinski. 2003a. Automatic summarisation of legal documents. In *Proceedings of the 9th international conference on Artificial intelligence and law*, pages 243–251.
- Claire Grover, Ben Hachey, and Chris Korycinski. 2003b. Summarising legal texts: Sentential tense and argumentative roles. In *Proceedings of the HLT-NAACL 03 Text Summarization Workshop*, pages 33–40.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In *International conference on machine learning*, pages 1885–1894. PMLR.
- Alex Kulesza, Ben Taskar, et al. 2012. Determinantal point processes for machine learning. *Foundations and Trends® in Machine Learning*, 5(2–3):123–286.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Manuj Malik, Zheng Zhao, Marcio Fonseca, Shrisha Rao, and Shay B Cohen. 2024. Civilsum: A dataset for abstractive summarization of indian court decisions. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2241–2250.
- Gianluca Moro and Luca Ragazzi. 2022. Semantic self-segmentation for abstractive summarization of long documents in low-resource regimes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11085–11093.
- Tatsuro Oya, Yashar Mehdad, Giuseppe Carenini, and Raymond Ng. 2014. A template-based abstractive meeting summarization: Leveraging summary and source text relationships. In *Proceedings of the 8th International Natural Language Generation Conference (INLG)*, pages 45–53.

Garima Pruthi, Frederick Liu, Satyen Kale, and Mukund Sundararajan. 2020. Estimating training data influence by tracing gradient descent. *Advances in Neural Information Processing Systems*, 33:19920–19930.

TYS Santosh, Mahmoud Aly, and Matthias Grabmair. 2024a. Lexabsumm: Aspect-based summarization of legal decisions. *arXiv preprint arXiv:2404.00594*.

TYS Santosh, Vatsal Venkatkrishna, Saptarshi Ghosh, and Matthias Grabmair. 2024b. Beyond borders: Investigating cross-jurisdiction transfer in legal case summarization. *arXiv preprint arXiv:2403.19317*.

TYS Santosh, Tuan-Quang Vuong, and Matthias Grabmair. 2024c. Chronoslex: Time-aware incremental training for temporal generalization of legal classification tasks. *arXiv preprint arXiv:2405.14211*.

TYSS Santosh, Cornelius Weiss, and Matthias Grabmair. 2024d. Lexsumm and lext5: Benchmarking and modeling legal summarization tasks in english. *arXiv preprint arXiv:2410.09527*.

Murali Saravanan, Balaraman Ravindran, and Shivani Raman. 2006. Improving legal document summarization using graphical models. *Frontiers in Artificial Intelligence and Applications*, 152:51.

Zejiang Shen, Kyle Lo, Lauren Yu, Nathan Dahlberg, Margo Schlanger, and Doug Downey. 2022. Multi-lexsum: Real-world summaries of civil rights lawsuits at multiple granularities. *Advances in Neural Information Processing Systems*, 35:13158–13173.

Abhay Shukla, Paheli Bhattacharya, Soham Poddar, Rajdeep Mukherjee, Kripabandhu Ghosh, Pawan Goyal, and Saptarshi Ghosh. 2022. Legal case document summarization: Extractive and abstractive methods and their evaluation. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing*, pages 1048–1064.

Megh Thakkar, Tolga Bolukbasi, Sriram Ganapathy, Shikhar Vashishth, Sarath Chandar, and Partha Talukdar. 2023. Self-influence guided data reweighting for language model pre-training. *arXiv preprint arXiv:2311.00913*.

Santosh Tyss, Hassan Sarwat, Ahmed Mohamed Abdelaal Abdou, and Matthias Grabmair. 2024. Mind your neighbours: Leveraging analogous instances for rhetorical role labeling for legal documents. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 11296–11306.

Shuohang Wang, Yichong Xu, Yuwei Fang, Yang Liu, Siqu Sun, Ruo Chen Xu, Chenguang Zhu, and Michael Zeng. 2022. Training data is more valuable than you think: A simple and effective method by retrieving from training data. *arXiv preprint arXiv:2203.08773*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 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.

Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. Alignscore: Evaluating factual consistency with a unified alignment function. *arXiv preprint arXiv:2305.16739*.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.

Ming Zhong, Yang Liu, Da Yin, Yuning Mao, Yizhu Jiao, Pengfei Liu, Chenguang Zhu, Heng Ji, and Jiawei Han. 2022. Towards a unified multi-dimensional evaluator for text generation. *arXiv preprint arXiv:2210.07197*.

A DPP Explanation

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:

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

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.

B Implementation Details

We implement our code using Huggingface Transformers library (Wolf et al., 2020). We use a learning rate of $1e-4$ and select the best model based on the R-1 score on the validation set. The

model is trained for 10 epochs with mixed precision, gradient clipping with a maximum norm of 1.0, early stopping patient set of 3 and optimized end-to-end with Adam optimizer (Kingma and Ba, 2014), $\beta_1 = 0.9$, and $\beta_2 = 0.999$. We set the maximum query length to 10240/4096; rest of encoder length (16384) is filled with 4/8 exemplars and the maximum output lengths are set to 4096/256 for SuperSCOTUS and CivilSum respectively. We set to filter out top 40 exemplars from first stage selection.