



LexGenie: Automated Generation of Structured Reports for European Court of Human Rights Case Law

Santosh T.Y.S.S¹, Mahmoud Aly¹, Oana Ichim², Matthias Grabmair¹

¹School of Computation, Information, and Technology;
Technical University of Munich, Germany

²Graduate Institute of International and Development Studies, Geneva, Switzerland
{santosh.tokala, mahmoud.aly, matthias.grabmair}@tum.de
oana.ichim@graduateinstitute.ch

Abstract

Analyzing large volumes of case law to uncover evolving legal principles, across multiple cases, on a given topic is a demanding task for legal professionals. Structured topical reports provide an effective solution by summarizing key issues, principles, and judgments, enabling comprehensive legal analysis on a particular topic. While prior works have advanced query-based individual case summarization, none have extended to automatically generating multi-case structured reports. To address this, we introduce LexGenie, an automated LLM-based pipeline designed to create structured reports using the entire body of case law on user-specified topics within the European Court of Human Rights jurisdiction. LexGenie retrieves, clusters, and organizes relevant passages by topic to generate a structured outline and cohesive content for each section. Expert evaluation confirms LexGenie’s utility in producing structured reports that enhance efficient, scalable legal analysis.

1 Introduction

Court judgments, beyond resolving individual cases, play a critical role in developing, clarifying, and safeguarding legal principles, ensuring the consistent application of law within a given jurisdiction (Farzindar, 2004; Saravanan et al., 2006; T.y.s.s et al., 2025a). Consequently, legal professionals face the challenging task of analyzing and synthesizing large volumes of complex case law to extract relevant legal precedents, understand the application of laws, and inform their legal strategies (Bhattacharya et al., 2019; Tyss et al., 2024c; Santosh et al., 2025). In response to this growing demand, recent efforts have focused on automatic summarization of individual cases, which condense the content of a single case, making it easier for legal professionals to quickly grasp key points (Zhong et al., 2019; Shukla et al., 2022;

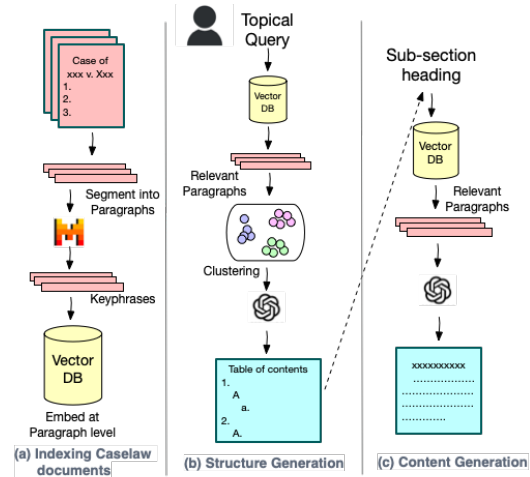


Figure 1: Overview of our approach, LexGenie.

Deroy et al., 2023; Santosh et al., 2024d). In practice, a single case may include multiple documents of varied types, such as complaints, opinions, motions, briefs, settlements, affidavits, and discovery materials—often totaling hundreds of pages per case, leading to exploration of multi-document legal summarization systems that process and distill information across multiple legal texts (Shen et al., 2022). Furthermore, a one-size-fits-all approach that produces a single, generic summary may not sufficiently address the diverse and specific needs of legal professionals and this limitation has spurred the development of aspect or query-focused case summarization systems, which provide tailored summaries based on users’ specific information needs, allowing for a more customized and relevant output (Tyss et al., 2024a).

While these solutions represent significant progress, they remain limited to single-case contexts, often missing the broader perspective necessary to track the evolution of legal principles across multiple cases. For more strategic analysis, legal professionals require cross-case insights that reveal how precedents and interpretations develop

over time. To meet this need, structured reports are typically prepared, focusing on specific Articles or Transversal Themes. These reports summarize key principles, issues, and judgments drawn from multiple cases by identifying relevant cases, recognizing common legal patterns, and organizing the information into a multidimensional framework—akin to a detailed table of contents, with insights structured under each dimension. Yet, manually generating these structured, multi-case reports is labor-intensive and time-consuming. As case law grows in volume and legal issues increase in complexity, the demand for automatically creating these reports has become pressing. Our work addresses this challenge by moving beyond single-case summarization toward an extreme summarization approach that synthesizes patterns and principles across entire body of case law. We explore the utility of current technologies, such as large language models (LLMs), to assist in generating structured reports that support a comprehensive, cross-case understanding of key legal issues.

We develop LexGenie, a fully automated pipeline leveraging LLMs to generate structured reports based on the topical queries issued by the user, focusing on European Court of Human Rights (ECHR) Jurisdiction, which adjudicates complaints by individuals against states about alleged violations of their rights as enshrined in the European Convention of Human Rights. LexGenie employs a two-stage pipeline: in the first stage, it retrieves relevant passages according to the user’s query, optimizing for recall and performs clustering to create a topic-based outline for the report. In the second stage, LexGenie generates content for each section by sourcing precise paragraphs for that sub-topic. We validate LexGenie’s effectiveness through a small-scale evaluation conducted by an ECHR legal expert, demonstrating its ability to produce accurate and relevant report structure and content. Additionally, we examine whether LLMs can assist in assessing output quality, finding a positive correlation with expert annotations.

2 LexGenie

We present the methodology behind LexGenie, for generating structured reports from ECHR case law judgments based on user-issued topical queries. LexGenie structurally organizes each report into a coherent dimensions related to the topic, enabling users to navigate and understand a thematic legal

area, supported by references to relevant case law judgments. We then describe our user interface, designed for accessibility and ease of use.

2.1 Approach

LexGenie’s workflow comprises three main steps: (i) indexing case law documents at the paragraph level, offline, into a vector datastore for efficient query-based retrieval, (ii) structure generation module, which retrieves relevant paragraphs based on the query, organizes them into hierarchical thematic clusters to finally generate a coherent outline with headings and sub-headings and (iii) content generation, where relevant content is sourced and expanded upon for each subsection of the outline.

2.1.1 Indexing Case law documents

We gather the complete ECHR case law collection from the latest version of Santosh et al. (2024a), sourced from HUDOC, the public ECHR database. Each judgment is organized by paragraph numbers, which serve as the primary unit for cross-referencing within the ECHR writing style (Tyss et al., 2024b).

Rather than indexing the raw paragraph text as embeddings, we use a keyphrase-based approach to represent each paragraph’s main themes. This focus on keyphrases enhances the embeddings by centering them around key legal concepts, while minimizing the inclusion of case-specific details that would otherwise arise with full-text embeddings, thus facilitating accurate, thematic matches with user queries. To obtain these keyphrases, we prompt the Mistral-7B-Instruct model (Jiang et al., 2023)¹ using each paragraph’s text. Appendix A.1 provides the prompt and an example of paragraph-level keyphrase generation. To improve efficiency, we use batch prompting (Cheng et al., 2023), running inference in groups of paragraphs sourced from the same judgment document rather than one at a time. This approach reduces token and processing time costs while contextualizing each paragraph within the broader scope of the case. Once generated, these paragraph keywords are concatenated and embedded using OpenAI’s text-embedding-3-small model. We store the resulting dense vector embeddings in a FAISS database, integrated via the LangChain framework², which enables efficient, semantically-similar retrieval.

¹<https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3>

²<https://www.langchain.com/langchain>

2.1.2 Structure Generation

This module analyzes the entire body of case law to extract relevant concepts related to user queries and organizes them into a coherent table of contents. By structuring sub-topics effectively, it enhances the user’s understanding of key legal themes and facilitates navigation through complex subjects. The process involves four main steps: retrieving relevant paragraphs, hierarchically clustering them based on shared themes, generating topical headings for each cluster, and organizing these headings into a cohesive narrative flow.

First, we retrieve relevant paragraphs based on Maximal Marginal Relevance (MMR) (Carbonell and Goldstein, 1998) using the LangChain framework. MMR balances relevance (semantic similarity with query) and diversity (semantic similarity between retrieved items), ensuring that the selected paragraphs encompass a broad spectrum of themes related to the query topic. Next, we apply BERTopic (Grootendorst, 2022) to cluster the retrieved paragraphs, which helps in identifying and organizing common themes. Utilizing the text-embedding-3-small model in conjunction with HDBSCAN, we generate hierarchical topical clusters (McInnes et al., 2017). To create topic headings for each cluster, we prompt the GPT-4o-mini model with five representative paragraphs from each cluster, as detailed in Appendix A.2. Once the topic names are generated for each cluster individually, we finally prompt GPT-4o-mini to refine all the headings and subheadings into a cohesive, ordered structure. This can involve re-ordering, merging and organizing topics to ensure logical flow across all (sub-)clusters, resulting in a well-structured report outline. Detailed prompt is provided in App. A.2.

2.1.3 Content Generation

In this phase, we generate content for each sub-section (leaf node) in the established table of contents. First, we construct a query by concatenating the sub-heading with the headings along its hierarchical path from the root node and is used to retrieve the top relevant paragraphs from the datastore. This augmented query, providing contextual relevance enables the retrieval of more precise paragraphs targeted towards the specific sub-section.

Next, we generate the content for each sub-section using the retrieved paragraphs following an iterative incremental updating approach using the GPT-4o-mini model (Chang et al., 2023), to han-

dle cases where the length of relevant paragraphs exceeds the model’s prompt length. In the initial iteration, the model is prompted with 25 relevant paragraphs to generate content for the specified sub-section, while also including references to the corresponding paragraphs. In subsequent iterations, the model receives the content generated up to that point along with the next set of 25 relevant paragraphs, prompting it to modify the previously generated content by integrating any additional insights from the latest paragraphs. Appendix A.3 provides the detailed prompts.

2.2 User Interface

LexGenie is accessible as a web app, which can be run locally and is available at <https://tinyurl.com/2a9jhrpu>. A video demonstration of LexGenie is available at <https://tinyurl.com/585h53cj>. An user inputs a search query to initiate the retrieval process. Adjustable parameters, such as the number of judgments retrieved and a similarity threshold, allow users to control the scope of retrieved content. In the initial retrieval step, relevant paragraphs are displayed as judgments, with paragraph numbers linked to the original HUDOC case law documents for easy reference. Users can further refine these results by adjusting the ranked list of retrieved items, removing, or adding new passages including additional passages from other cases in the datastore can be incorporated through a fuzzy search-based dropdown menu. Based on these refined paragraphs, a structured table of contents is generated through clustering and organization using LLM calls.

Users can review and edit the generated table of contents before proceeding to content generation. The table of contents is displayed in a side navigation panel, allowing users to navigate through the hierarchy of headings and subheadings. Then users can generate content for individual (sub-)sections or for the entire table of contents by clicking the appropriate buttons. The generated content is post-processed to include citations linking each segment of the report to the respective ECHR documents, based on the references provided by the LLM model. The final report is available for download in PDF format. LexGenie’s UI is designed to reflect the underlying pipeline, making each step in the report creation process transparent and customizable. This design enables feedback collection from users, allowing us to assess the interface’s effectiveness at each stage of the process.

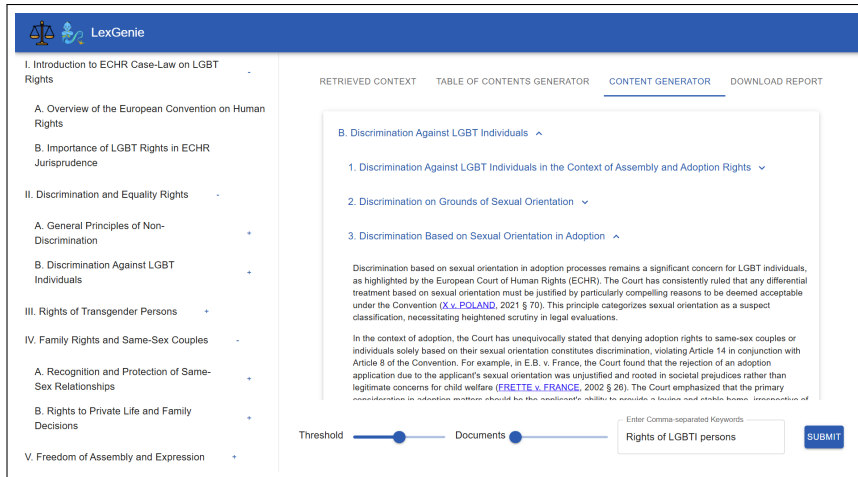


Figure 2: LexGenie interface. Given a legal topic as query, it automatically retrieves relevant documents and generates a table of content structure for the report. Finally, content for each sub-section in report is populated and the whole report is available for download.

3 Evaluation

3.1 Report Structure

We assess the quality of the generated report structure across the following dimensions: (a) **Topical Relevance**: Emphasizes how closely the generated headings and subheadings align with the user’s query. (b) **Subtopic Consistency**: Focuses on the alignment of subtopics under each parent heading, ensuring intra-cluster consistency. (c) **Cluster Distinction**: Highlights the uniqueness of each topic cluster, ensuring clear differentiation and minimal inter-cluster redundancy. (d) **Narrative Flow**: Evaluates the logical progression of the structure, ensuring it guides the reader smoothly through the topics. (e) **Comprehensiveness of Topics**: Measures the extent to which the headings and subheadings encompass all critical aspects of the query, avoiding any significant gaps.

We investigate the effect of each design choice in LexGenie: (i) **Keyphrase vs. Paragraph-Based Indexing**: While LexGenie employs a keyphrase-based approach to index each paragraph for retrieval and clustering, we modify it to use the raw paragraph text for indexing and further clustering. (ii) **Retrieval Strategy**: LexGenie uses the Maximal Marginal Relevance (MMR) criterion to balance relevance and diversity. We replace it with a traditional relevance-based criterion. (iii) **Impact of Reorganization**: LexGenie employs an LLM call to order the generated headings and subheadings into a cohesive structure. We remove this call and concatenate the cluster-based individually generated headings to form the final structure.

Human Evaluation We randomly select 20 queries covering broad topics such as Articles and Themes from existing ECHR case law guides³. We then ask a legal expert, the third author of this paper, to manually evaluate the quality of the generated report structures using a 1-5 scale, where a higher score indicates better quality. The evaluation is based on each of the five dimensions outlined above for LexGenie and the three ablation systems.

From Table 1, we observe that the report structure generated by LexGenie is highly rated by our legal expert, reflecting topically relevant headings, well-grouped sub-topics with clear delineation, logically organized narrative flow, and comprehensive coverage of relevant aspects. The keyphrase-based approach significantly outperforms the paragraph-based approach across all metrics, particularly in sub-topic consistency and cluster distinction. This suggests that keyphrase generation effectively steers the model to focus on core legal concepts, while paragraph embeddings tend to capture additional case-specific details, which may dilute relevance in retrieval and clustering.

When diversity criterion in retrieval (w/o MMR) is removed, we observe the appearance of similar sub-topics among the top retrieved results, leading to gaps in topic coverage, as reflected in lower comprehensiveness scores. The reduced cluster distinction can be attributed to the lack of sub-topic diversity, which complicates clear separation be-

³Case-law guides are structured reports maintained by courts registry, accessible on the ECHR Knowledge Sharing Platform at <https://ks.echr.coe.int/web/echr-ks/all-case-law-guides>.

Model	Topical Rel.		Subtopic Con.		Cluster Dist.		Narr. Flow		Comprehen.	
	Human	Auto	Human	Auto	Human	Auto	Human	Auto	Human	Auto
LexGenie	3.95	4.48	3.75	4.49	3.70	3.91	3.75	4.27	3.25	3.27
Paragraph-based	3.85	4.38	3.50	4.33	3.45	3.40	3.60	4.18	3.10	3.16
w/o MMR	3.90	4.31	3.70	4.09	3.55	3.53	3.85	4.22	3.0	2.84
w/o Reorganization	3.85	4.45	3.55	3.97	3.45	2.69	3.55	3.37	3.20	3.06

Table 1: Human and Automatic Evaluation Results for Report Structure Quality.

	Topical Relevance		Content Org.		Citation Faith.		Comprehen.	
	Human	Auto	Human	Auto	Human	Auto	Human	Auto
Single	4.6	4.87	4.5	4.47	3.9	4.29	4.0	4.23
Incremental	4.9	4.94	4.5	4.59	4.5	4.36	4.5	4.55

Table 2: Human and Automatic Evaluation Results for Content Quality.

tween clusters. Although narrative flow improves slightly due to less diverse sub-topics, this comes at the cost of thematic variety. Lastly, omitting the LLM reorganization step results in declines across narrative flow, sub-topic consistency and cluster distinction. Without reorganization, the structure lacks coherence and topics are less clearly differentiated, ultimately hindering thematic clarity.

Automatic Evaluation We evaluate the capabilities of LLMs to conduct automated assessments across the five dimensions using the same set of 20 queries selected for human evaluation. We employ the G-Eval (Liu et al., 2023) framework, which prompts LLMs with chain-of-thought and a form-filling paradigm, to assess the quality of generated outputs. For all metrics, we provide the detailed instruction, generated report structure along with the query to provide an assessment in scale of 1-5. For comprehensiveness of topics evaluation, which requires additional external knowledge to understand the topical coverage, we also provide the model with table of content structure from original case law guides as reference context. This allows the model to compare the generated content structure against this reference to identify the missing aspects, to assess comprehensiveness.

From Table 1, we observe that LexGenie achieves high scores across automated metrics, aligning closely with human expert evaluations. Notably, the automated metrics reveal lower comprehensiveness scores for the approach without MMR, attributed to reduced sub-topic diversity in the retrieval process—an observation mirrored in the expert assessments. Likewise, the absence of reorganization adversely impacts narrative flow and both intra- and inter-cluster consistency. Additionally, the paragraph-based approach underperforms

relative to the keyphrase-based approach, both in retrieval and clustering, suggesting that keyphrase-based representations better capture core topics enhancing intra- and inter- cluster consistency.

3.2 Content Generation

We assess the quality of the generated content under each (sub-)heading across the following dimensions: (a) Topical Relevance: measures how well the generated content aligns with the (sub-)section heading. (b) Content Organization: evaluates the logical flow and coherence of the content throughout. (c) Citation Faithfulness: assesses the extent to which the generated content is supported by appropriate and reliable citations. (d) Comprehensiveness: examines whether all relevant aspects of the section topic are comprehensively addressed, ensuring no critical information is overlooked. While an incremental prompting is used in LexGenie, we compare it with a single prompting approach where content is generated using all retrieved passages provided to the model simultaneously.

Human Evaluation We randomly select 10 (sub-)headings from existing ECHR case law guides, which serve as leaf nodes and generate corresponding content for each. A legal expert manually evaluates the quality of the generated content for each heading on a 1-5 scale across the four dimensions, with higher scores indicating better quality. As shown in Table 2, both the incremental and single prompting approaches maintain a coherent narrative structure. However, the incremental prompting generates content that is more firmly grounded in the provided heading and retrieved paragraphs, with appropriate citations, in contrast to the single prompting approach. The lower performance of the single setup can be attributed to the overwhelming

amount of content presented to the LLMs, which complicates the distillation of important information across multiple paragraphs. This suggests that the model is better able to focus on relevant aspects when given smaller batches of paragraphs rather than handling all the retrieved context at once. This phenomenon aligns with the well-known "lost in the middle" problem (Liu et al., 2024), wherein models struggle to access relevant information situated in the middle of long contexts, even for models designed to handle long contexts. Consequently, this results in lower comprehensiveness scores, as some relevant information is overlooked despite using the same retrieved paragraphs in both setups.

Automatic Evaluation We conduct an automatic assessment using the G-Eval framework across the four dimensions with the 10 sampled headings. The LLM is prompted with the generated content and headings, along with specific instructions tailored for each metric. To evaluate citation faithfulness, we include the original paragraphs from the cited case law judgments within the generated content. For comprehensiveness, we provide the actual content corresponding to each heading from the case law guide. As shown in Table 2, the automated assessments correlate closely with human evaluations across these dimensions. While expert assessment remains essential for gauging the quality and utility of structured reports, our findings indicate that automated LLM-based evaluations using the G-Eval framework can deliver rapid insights, offering a cost-effective alternative to expert assessments.

3.3 Qualitative Case Study

We conduct a qualitative case study on the LexGenie-generated report focusing on the 'Rights to LGBTI Persons'. A complete generated report is provided in <https://tinyurl.com/43f86jw8>. The most compelling aspect identified is the detailed treatment of discrimination and equality rights, particularly the focus on intersectionality under sub-topic II.A2. This section effectively illustrates how LGBTI rights, though not explicitly enumerated in the European Convention on Human Rights, have been progressively built through interpretations of various articles, notably Article 8 (private life) and Article 14 (discrimination). These provisions have been instrumental in advancing LGBTI protections, including adoption rights, succession rights, marriage equality, and pension benefits. The system's ability to highlight these key substantive aspects captures the ECHR's approach

to addressing discrimination against LGBTI individuals. Sections II and III provide the most insightful overviews, offering well-supported legal protections, and references to relevant case citations, supporting those claims.

Despite these strengths, the model has notable shortcomings. It overlooks crucial contextual insights, such as the role of states' duties and positive obligations, which are vital for understanding discrimination cases. Additionally, it fails to address significant areas like migration issues, which span Articles 3, 8, and 5, and hate crime protections under Articles 3, 10, and 11. These omissions undermine a comprehensive understanding of the ECHR's jurisprudence. Structurally, the absence of transitional sub-topics or thematic connectors disrupts the logical flow, making it difficult to grasp the interconnected nature of topics like freedom of assembly (V) and LGBTI rights. This limitation stems from the current content generation pipeline, which focuses on isolated subsections without addressing cross-section redundancies or integrating detailed contextual links. Bridging these structural and contextual gaps could greatly enhance the usability and coherence of these generated guides.

4 Conclusion

In this paper, we introduce LexGenie, an automated LLM-based pipeline designed to generate structured report based on user-specified query from extensive case law, specifically within the ECHR jurisdiction. LexGenie's two-stage pipeline first retrieves and organizes relevant passages according to user-defined topical queries, creating a structured outline that captures core legal issues and patterns. In the second stage, it generates cohesive, contextually accurate content for each section, providing a nuanced understanding of complex legal matters. Expert evaluations confirm LexGenie's effectiveness in delivering relevant, well-organized reports, illustrating its potential to enable scalable, high-quality legal analysis. Additionally, initial automated evaluations using LLMs indicate a promising alternative to traditional expert reviews. Despite its strengths, challenges such as improving context integration and addressing structural flow remain. Future work can expand to other jurisdictions and integrate multi-case analysis tools, such as temporal trend identification, to further support legal professionals in dynamic legal landscapes.

Limitations

One key limitation lies in the quality of the retrieved passages and their clustering. Although the system aims to organize content into meaningful outlines, errors in retrieval or clustering can result in misaligned or overly broad sections that dilute the coherence of the report. This issue is particularly pronounced when dealing with ambiguous or overlapping topics, where the system may fail to distinguish fine-grained distinctions between related legal principles. Additionally, the pipeline does not currently incorporate mechanisms for ranking retrieved content by legal importance or authoritativeness (T.y.s.s et al., 2025b), which can lead to the inclusion of peripheral or temporally outdated information (Santosh et al., 2024b,c).

Another limitation is the lack of advanced contextual linking across sections. LexGenie generates content for individual subsections in isolation, which often results in a disjointed narrative that fails to capture the interconnected nature of legal issues. This fragmentation can hinder a comprehensive understanding of the broader legal landscape and reduce the utility of the generated reports for complex legal analyses.

Ethics

We utilize case law data from HUDOC, the official database of the European Court of Human Rights. This publicly available data includes the real names of individuals involved, as the judgments are not anonymized. However, our work engages with this data solely for research purposes, without any intent or functionality that could exacerbate harm beyond the inherent exposure of the data's public availability.

LexGenie is developed as a tool to assist legal professionals by automating the generation of structured reports from case law, enhancing the efficiency of legal research. The system is intended to augment human expertise rather than replace it. While LexGenie provides valuable insights, its outputs may contain errors, such as hallucinated or misinterpreted legal references, which necessitate careful review and validation by qualified professionals. Users are explicitly advised against relying solely on the system for critical legal decisions. By ensuring the tool's transparency and openly sharing its methodology, we aim to promote responsible use while underscoring the need for human oversight in all applications.

Additionally, the reliance on pre-trained large language models introduces the risk of perpetuating biases present in the training data. Legal judgments often reflect historical biases or systemic inequities, and there is a potential for these to be inadvertently amplified in LexGenie's outputs. To address these challenges, we advocate for continuous monitoring, user feedback and iterative improvements to the system. This includes efforts to identify and mitigate any biases, ensuring that the tool aligns with ethical standards.

References

- 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.
- Jaime Carbonell and Jade Goldstein. 1998. The use of mmr, diversity-based reranking for reordering documents and producing summaries. In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 335–336.
- Yapei Chang, Kyle Lo, Tanya Goyal, and Mohit Iyyer. 2023. Boookscore: A systematic exploration of book-length summarization in the era of llms. *arXiv preprint arXiv:2310.00785*.
- Zhoujun Cheng, Jungo Kasai, and Tao Yu. 2023. Batch prompting: Efficient inference with large language model apis. *arXiv preprint arXiv:2301.08721*.
- 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*.
- 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.
- Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.

- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. G-eval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.
- Leland McInnes, John Healy, and Steve Astels. 2017. hdbscan: Hierarchical density based clustering. *Journal of open source software*, 2(11):205.
- TYSS Santosh, Rashid Gustav Haddad, and Matthias Grabmair. 2024a. Ecthr-pcr: A dataset for precedent understanding and prior case retrieval in the european court of human rights. *arXiv preprint arXiv:2404.00596*.
- TYSS Santosh, Kevin D Ashley, Katie Atkinson, and Matthias Grabmair. 2024b. Towards supporting legal argumentation with nlp: Is more data really all you need? In *Proceedings of the Natural Legal Language Processing Workshop 2024*, pages 404–421.
- TYSS Santosh, Chen Jia, Patrick Goroncy, and Matthias Grabmair. 2025. Relexed: Retrieval-enhanced legal summarization with exemplar diversity. *arXiv preprint arXiv:2501.14113*.
- TYSS Santosh, Tuan-Quang Vuong, and Matthias Grabmair. 2024c. Chronoslex: Time-aware incremental training for temporal generalization of legal classification tasks. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3022–3039.
- TYSS Santosh, Cornelius Weiss, and Matthias Grabmair. 2024d. Lexsumm and lext5: Benchmarking and modeling legal summarization tasks in english. In *Natural Legal Language Processing Workshop 2024*, volume 2024, pages 381–403.
- 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.
- Santosh Tyss, Mahmoud Aly, and Matthias Grabmair. 2024a. Lexabsumm: Aspect-based summarization of legal decisions. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 10422–10431.
- Santosh T.y.s.s, Youssef Farag, and Matthias Grabmair. 2025a. CoPERLex: Content planning with event-based representations for legal case summarization. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 1016–1032, Albuquerque, New Mexico. Association for Computational Linguistics.
- Santosh Tyss, Elvin A Quero Hernandez, and Matthias Grabmair. 2024b. Query-driven relevant paragraph extraction from legal judgments. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 13442–13454.
- Santosh T.y.s.s, Isaac Misael Olguín Nolasco, and Matthias Grabmair. 2025b. LeCoPCR: Legal concept-guided prior case retrieval for European court of human rights cases. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 1654–1661, Albuquerque, New Mexico. Association for Computational Linguistics.
- Santosh Tyss, Vatsal Venkatkrishna, Saptarshi Ghosh, and Matthias Grabmair. 2024c. Beyond borders: Investigating cross-jurisdiction transfer in legal case summarization. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4136–4150.
- Linwu Zhong, Ziyi Zhong, Zinian Zhao, Siyuan Wang, Kevin D Ashley, and Matthias Grabmair. 2019. Automatic summarization of legal decisions using iterative masking of predictive sentences. In *Proceedings of the seventeenth international conference on artificial intelligence and law*, pages 163–172.

A LexGenie Prompts

A.1 Keyphrase Generation

Table 3 provides an example of paragraph and generated keyphrases. Prompt 3 provides the prompt used for generating keyphrases.

A.1.1 Prompt

A.2 Structure Generation

Prompt 4 and 5 provide detailed prompts used for generating topic name for each cluster and final LLM call to organize the generated topics and sub-topics into a coherent table of contents respectively.

You are an ECHR lawyer trying to create Legal Case-Law Guides that provide an in-depth overview of Convention ECHR case law on a particular Article or Transversal Theme. You will receive a paragraph extracted from case law; your task is to generate keywords that capture the essence of the paragraph so these keywords reflect the relevant Article or Transversal Theme and can be used to cluster cases, identify important cases, generate the table of contents and content for the Guides.

[Instructions]

1. Identify cross-references between paragraphs and reveal their connections;
2. Make sure keywords reflect the overall context of the paragraph by linking the description of circumstances to the requirements provided as criteria for legal doctrines and norms;
3. Map keywords like 'sometimes', 'exceptionally', 'in the present case' with the view to make sure that there is correspondence between legal standards and circumstances;
4. Focus on keywords detailing the application of substantive or procedural limb/branch explaining the scope of application of the Article;
5. Make sure to map accordingly keywords that detail the application of the Article to a variety of persons such as victims, state agents, witnesses, relatives, and similar;
6. Make sure to map accordingly keywords that detail the application of the Article depending on the jurisdiction, material, or temporal and those which detail the repartition or just satisfaction;
7. Distinguish conditions for the application of the Article in the context of violence/force from conditions that detail other events such as accidents or industrial activities;
8. Carefully identify key phrases that describe risks and operational choices from keywords that describe the creation and application of regulatory framework and conditions for responsibility of and accountability of various actors;
9. Highlight keywords that describe thresholds or conditions concerning intensity, frequency, and ordering in assessing each of the above.

[Paragraph]

{paragraph}

Please return ONLY the keywords for the given paragraph in one line and nothing else. Make sure to keep keywords in arguments together so they make sense.

Prompt 3: Generating keyphrases from paragraphs of case law judgements.

You are given a list of paragraphs extracted from the European Court of Human Rights case law, and your task is to generate a detailed topic label to represent these paragraphs in ECHR case law guidelines. Here is the list of paragraphs:

[DOCUMENTS]

Based on the information above, generate a detailed topic label in the following format and nothing more:
topic: <topic label>

Prompt 4: Generating topic name for each cluster.

A.3 Content Generation

Prompt 6 provides the prompt for iterative content generation approach for each leaf sub-section in the table of contents.

B LexGenie UI

Figure 7 displays the UI interface and functionalities offered through LexGenie.

I have a list of topics related to European Court of Human Rights (ECHR) case law documents. I would like you to organize these topics into a coherent and structured Table of Contents (ToC) similar to a legal document ECHR guidelines. Please group related topics under appropriate sections and subsections, ensuring a logical flow. The ToC should include main headings, subheadings, and possibly further subdivisions where necessary with 4 spaces indentation and without general sections such as introduction and conclusion. The final structure should resemble an outline for comprehensive legal report guidelines that align with the topics from ECHR. Here is the list of topics:

[Topics]
 {Topics}

Please only return a well-structured ToC and nothing else.

Prompt 5: Organize topics into a hierarchical structure.

Paragraph	The applicant submitted that the manner in which he had been forced to undergo the medical intervention had amounted to torture. The taking of the urine sample had been coercive, and he had never given his consent to the procedure.
Keyphrases	forced medical intervention, coercive, lack of consent, urine sample, torture, manner of procedure.

Table 3: An example of a paragraph along with its generated keyword representations.

You are a legal expert tasked with generating content for a Case Law Guidelines section based on the given section heading, current section content, and a set of paragraphs extracted from case law documents. Your goal is to synthesize the information from these paragraphs to extend and create clear and accurate content without sections like introductions or subsections. The content should be strictly related to the heading and logically coherent, and the relevant paragraphs from the case law documents should be cited by their IDs. Provide thorough explanations, elaborate on key points, and include examples where relevant. Follow the instructions below carefully to ensure the guidelines are precise and informative.

[Instructions]

1. Review the provided set of paragraphs extracted from case law documents;
2. Consider only those paragraphs that are strictly related to the keywords in the heading;
3. Develop content based on the information principles contained in the paragraphs and ensure the content is clear and concise;
5. Citations: whenever a guideline is influenced by or derived from a specific paragraph, cite that paragraph by its id and number in parentheses as (id#paragraph_number);
6. Maintain a professional and formal tone throughout;
7. Only generate the content in relation to the keywords in the heading and focus on the specific standards implied by those keywords;
8. Return a coherent answer comprising general observations and standards from the Convention and specific observations and standards implied by the keywords in the heading;
9. Extend the previously generated content with the new content, revising and integrating it smoothly to form a coherent narrative;

[Heading]

{Heading}

[Previous Content]

{Previous_Content}

[Paragraphs]

{Paragraphs}

Return the generated content and nothing else. Make sure to use only the related paragraphs to the heading.

[Your response]

Prompt 6: Content Generation.

LexGenie

RETRIEVED CONTEXT TABLE OF CONTENTS GENERATOR CONTENT GENERATOR DOWNLOAD REPORT

GENERATE TABLE OF CONTENTS + ADD PARAGRAPH

Showing 100 of 788 paragraphs

- [A. B AND C v. IRELAND](#), 2010 § 188
- [A.D.T. v. THE UNITED KINGDOM](#), 2000 § 27
- [A.D.T. v. THE UNITED KINGDOM](#), 2000 § 32
- [A.D.T. v. THE UNITED KINGDOM](#), 2000 § 38
- [A.D.T. v. THE UNITED KINGDOM](#), 2000 § 40

Threshold Documents

Enter Comma-separated Keywords
Rights of LGBTI persons SUBMIT

LexGenie

RETRIEVED CONTEXT TABLE OF CONTENTS GENERATOR CONTENT GENERATOR DOWNLOAD REPORT

Table of contents

- I. Introduction to ECHR Case-Law on LGBT Rights
 - A. Overview of the European Convention on Human Rights
 - B. Importance of LGBT Rights in ECHR Jurisprudence
- II. Discrimination and Equality Rights
 - A. General Principles of Non-Discrimination
 - Article 14 of the ECHR
 - Intersectionality in Discrimination Cases
 - B. Discrimination Against LGBT Individuals
 - Discrimination Against LGBT Individuals in the Context of Assembly and Adoption Rights
 - Discrimination on Grounds of Sexual Orientation
 - Discrimination Based on Sexual Orientation in Adoption
 - Discrimination Based on Sexual Orientation in Tenancy Succession Rights
 - Discrimination Against Same-Sex Couples in Marriage and Pension Rights
 - Discrimination and Enjoyment of Rights under Article 14
- III. Rights of Transgender Persons
 - A. Legal Recognition and Rights
 - Legal Recognition and Rights of Transgender Persons

Threshold Documents

Enter Comma-separated Keywords
Rights of LGBTI persons SUBMIT

Figure 7: LexGenie interface. Given a legal topic as query, it automatically retrieves relevant documents and generates a table of content structure for the report. Finally, content for each sub-section in report is populated and the whole report is available for download.