Extract-Explain-Abstract: A Rhetorical Role-Driven Domain-Specific Summarisation Framework for Indian Legal Documents

Anonymous EMNLP submission

Abstract

Legal documents are characterized by their length, intricacy, and dense use of jargon, making efficacious summarisation both paramount and challenging. This paper introduces the Rhetorical Role-based Extract-Explain-Abstract (EEA) Framework, a novel three-stage methodology for summarisation of Indian legal documents in low-resource settings. The approach begins by segmenting legal texts using rhetorical roles, such as facts, issues and arguments, through a domain-specific phrase corpus and extraction based on TF-IDF. In the explanation stage, the segmented output is enriched with logical connections, leveraging rhetorical structure theory to ensure coherence and legal fidelity. The final abstraction phase condenses these interlinked segments into cogent, high-level summaries that preserve critical legal reasoning. We focus primarily on small language models (SLMs) because they can be efficiently deployed on local GPUs for cost-effective fine-tuning on specific legal domains or drafting styles. Experiments on Indian legal datasets show that the EEA framework typically outperforms in ROUGE, BERT scores and human evaluations. We also employ InLegalBERT score as a metric to capture domain specific semantics of Indian legal documents.

1 Introduction

Lawyers spend up to 2.3 hours a week searching for documents and another 2 hours recreating missing ones, resulting in an annual loss of $9,071 per lawyer, or 9.8% of their productivity (MetaJure, 2015). Here, summarisation plays a crucial role in condensing information while maintaining meaning. Text summarisation can be extractive; selecting key phrases verbatim, or abstractive; generating new text using NLP methods. Despite advancements in both supervised and unsupervised approaches, accurately aligning summaries with legal standards remains difficult. Large Language Models (LLMs) have advanced legal text summarisation by contextual understanding (Devlin et al., 2019). Although fine-tuning on legal corpora can improve domain performance (Bommarito and Katz, 2022), it introduces jurisdictional biases and the risk of misclassification. We highlight the use of Small Language Models (SLMs) which are preferred in low-resource settings with respect to availability of data, computational power, as fine-tuning can become resource heavy. This directly addresses the challenges of high resource demands and domain-specific accuracy in legal summarisation tasks (Wang et al., 2024). We address these challenges with a three-stage, zero-shot summarisation framework and make our code available[1]. First, rhetorical structure theory (RST) is applied to extract legally significant segments (Mann and Thompson, 1988). Second, a chain-of-thought prompt guides the model in interpreting these segments. Third, an abstractive module generates a cogent summary.

The remainder of this paper is structured as follows: Section 2 reviews related work; Section 3 describes our methodology; Section 4 covers experimental setup and evaluation; Section 5 presents results; Section 6 delineates human evaluation.

2 Related Work

We divide this section into two parts– extractive and abstractive for simplicity.

2.1 Extractive summarisation

This approach seeks to identify and select salient sentences or phrases in a source document to later create a cogent summary. Reduction (Jing, 2000) is an early extractive unsupervised method, while others include graph-based LexRank (Erkan and Radev, 2011) and BERT-based PacSum (Zheng and Lapata, 2019). Supervised methods include SummaRuNNer (Nallapati et al., 2017) and BERT-

[1] https://anonymous.4open.science/r/EEA-89FA/

Sum(Liu and Lapata, 2019). Domain-specific approaches like LetSum (Farzindar and Lapalme, 2004) and KMM (Saravanan et al., 2006) rank sentences using TF-IDF and k-mixture models. CaseSummarizer (Polsley et al., 2016) targets legal documents, while Gist (Liu and Chen, 2019) ranks sentences based on inclusion likelihood.

2.2 Abstractive summarisation

Abstractive methods generate novel sentences to convey key information. LegalSumm (Feijo and Moreira, 2021) uses the RulingBR dataset (Galassi et al., 2018); being the first of it's kind for legal documents and ArgLegalSumm (Elaraby and Litman, 2022) applies to Canadian legal cases. Reliability of abstractive model's is challenged by hallucination, facts' credibility, and an arduous path to transition to adaptation due to their specialized corpus. Chain of Thought (CoT) methodology enhances the quality of summarisation by guiding models through intermediate reasoning steps (Wei et al., 2022) applied in the SumCoT technique (Wang et al., 2022). Zero-shot-CoT introduces a template-based prompting method (Kojima et al., 2022) that enables multi-hop reasoning across tasks, as demonstrated in (Liu et al., 2021).

3 Proposed Model

We propose a three-stage Extract-Explain-Abstract (EEA) framework(Figure 3 in Appendix A.6) for summarisation, extending Extract-then-Abstract methods (Bhandari and Das, 2023; Dey et al., 2021). Extraction selects relevant sentences, Explanation organizes them logically, and Abstraction generates a concise summary. Using Chain-of-Thought prompting, we replace element-aware extraction (Wang et al., 2023) with rhetorical role-based extraction from a domain-specific corpus. We also perform an ablation study by successively removing Explanation and Extraction stages, namely, the Extract-Abstract (EA) and Abstract strategies.

3.1 Extraction

Our extraction builds on LetSum (Farzindar and Lapalme, 2004), a TF-IDF-based legal summarisation method, enhanced with a domain-specific phrase corpus developed and verified by legal experts using generative AI. Sentences are classified by matching phrases to legal categories (e.g., Facts, Issues) (see Figure 1 in Appendix A.1). After preprocessing (abbreviation removal, lowercasing), sentences are scored by TF-IDF plus phrase matches, assigned to the highest-scoring rhetorical role, and filtered for redundancy by thresholding.

Algorithm 1 Categorizing Legal Sentences using TF-IDF and Role-Based Phrases

Require: Cleaned sentence S, role-based phrases P_c, TF-IDF scores $T(p)$, threshold θ
Ensure: Category with the highest valid score

```
Initialize score dictionary Score(c) ← 0, ∀c ∈ C
for all categories c ∈ C do
    for all phrases p ∈ P_c do
        if p appears in S then
            Score(c) ← Score(c) + (1 + T(p))
        end if
    end for
end for
Determine Best Matching Role:
C* ← {c ∈ C | Score(c) > θ}
if C* ≠ ∅ then
    Return category with highest score in C*
else
    Return No Matching Role
end if
```

3.2 Explanation

The extraction stage outputs text segmented by rhetorical roles (Mann and Thompson, 1988; Saravanan et al., 2008). Using RST, we identify logical dependencies: Facts are the foundational units, linked to Issues, which lead to Arguments supported by facts, statutes, and precedents (Figure 2 in Appendix A.5). The lower court's decision and the current court's ruling depend on this information (Moens et al., 2007; Saravanan et al., 2008). Since segmented text lacks explicit links, we build contextual connections across segments to enrich the summary with causal relations (Mann and Thompson, 1988; Saravanan et al., 2008).

3.3 Abstraction

The abstraction phase summarises pre-segmented legal content from the Extraction and Explanation stages, preserving the logical structure established earlier. Due to the length of Indian legal documents exceeding typical model input limits, a chunking strategy splits documents into token-limited segments, each summarised individually (Shukla et al., 2022). These partial summaries are then concate-

IN-Abs

Method	LlaMa 3.2-1B		LlaMa 3.2-3B		Qwen-2.5-7B		Phi-4-14B	
	InLegalBERT	BertScore	InLegalBERT	BertScore	InLegalBERT	BertScore	InLegalBERT	BertScore
Abstract	0.8706	0.8100	0.9347	0.8304	0.9466	**0.8334**	**0.9544**	**0.8409**
EA	0.8919	0.8145	0.9453	**0.8312**	0.9421	0.8311	0.9452	0.8359
EEA	**0.9193**	**0.8272**	**0.9458**	0.8293	**0.9472**	0.8305	0.9413	0.8378

ILC

Method	LlaMa 3.2-1B		LlaMa 3.2-3B		Qwen-2.5-7B		Phi-4-14B	
	InLegalBERT	BERTScore	InLegalBERT	BERTScore	InLegalBERT	BERTScore	InLegalBERT	BERTScore
Abstract	0.8790	0.8122	0.9464	**0.8365**	0.9347	**0.8366**	**0.9433**	**0.8336**
EA	0.8953	0.8163	0.9482	0.8321	0.9340	0.8337	0.9431	0.8290
EEA	**0.9301**	**0.8219**	**0.9506**	0.8275	**0.9356**	0.8303	0.9366	0.8284

Table 1: InLegalBERT and BERTScore scores across summarisation strategies (rows) and models (columns).

nated and reprocessed to produce the final summary. The abstraction module prioritizes legal accuracy, redundancy reduction, and improved readability through rephrasing and synthesis.

4 Experimentation

4.1 Setup

All tests were run on an NVIDIA Tesla P100 GPU. Inference time varied with document length and complexity. We used QLoRA quantization(Dettmers et al., 2023) via the Unsloth framework (Daniel Han and team, 2023), significantly reducing memory and compute needs without performance loss, enabling scalable experimentation.

4.2 Models

We focus primarily on smaller language models (SLMs) because they can be efficiently deployed on local, on-premises GPUs, enabling cost-effective fine-tuning for specific legal domains or drafting styles. This targeted fine-tuning also helps reduce hallucinations compared to larger LLMs. For these reasons, we focus on open-source models that balance performance and efficiency, primarily using LLaMA 3.2's 1B and 3B variants (Van Der Maaten et al., 2024). We also benchmarked larger open models like Qwen-2.5 7B (Yang et al., 2023) and Phi-4 14B (Behl et al., 2024) to ensure scalability while maintaining transparency and personalization for legal research.

4.3 Datasets

We experimented with two benchmark Indian legal datasets:

i. ILC (Indian Legal Corpus): 3,073 documents from Indian Caselaw, avg. document length 2,380 tokens, avg. summary length 559 tokens. (Trivedi et al., 2023).

ii. IN-Abs: 7,030 Supreme Court judgments, avg. document length 1,051 words, avg. summary length 4,378 words. (Shukla et al., 2022).

5 Results

To evaluate the quality of generated summaries, we report ROUGE scores (Table 2) for lexical overlap, and both BERTScore and InLegalBERTScore (Table 1) for semantic similarity. While BERTScore leverages contextual embeddings from general-purpose language models, it does not explicitly account for legal terminology or structure. In contrast, InLegalBERTScore is a domain-specific semantic similarity metric derived by computing the cosine similarity between the embeddings of generated summaries and reference summaries using InLegalBERT (Paul et al., 2023)—a transformer-based model pretrained on an Indian legal corpus, including statutes, judgments, and legal commentaries. This grounding allows it to capture the nuances, formal phrasing, and domain-specific semantics characteristic of Indian legal language.

Our results demonstrate that the EEA strategy consistently yields higher InLegalBERTScores. These findings highlight the limitations of generic evaluation metrics in specialized domains: while general-purpose metrics like BERTScore provide a coarse-grained estimate of semantic similarity, they may overlook critical legal distinctions. We theorise that Phi-4's EEA semantic scores are inferior to other strategies due to it's shorter pretrained con-

IN-Abs

Method	LlaMa 3.2-1B			LlaMa 3.2-3B			Qwen-2.5-7B			Phi-4-14B		
	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
Abstract	0.2372	0.0941	0.1352	0.3636	0.1578	0.2003	0.4450	**0.1950**	0.2209	0.4326	0.1642	0.2140
EA	0.2385	0.0960	0.1380	0.3836	0.1529	0.2033	0.4388	0.1855	0.2178	0.4353	0.1611	0.2099
EEA	**0.2387**	**0.0962**	**0.1443**	**0.3924**	**0.1642**	**0.2079**	**0.4498**	0.1896	**0.2242**	**0.4648**	**0.1855**	**0.2235**

ILC

Method	LlaMa 3.2-1B			LlaMa 3.2-3B			Qwen-2.5-7B			Phi-4-14B		
	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L	R-1	R-2	R-L
Abstract	0.2275	0.1039	0.1411	0.3867	**0.1714**	0.2176	0.4401	**0.1811**	**0.2286**	0.4725	0.1728	**0.2357**
EA	0.2441	0.1001	0.1436	0.4051	0.1649	0.2165	0.4426	0.1772	0.2243	0.4603	0.1668	0.2229
EEA	**0.3053**	**0.1097**	**0.1732**	**0.4335**	0.1700	**0.2266**	**0.4480**	0.1771	0.2236	**0.4787**	**0.1756**	0.2232

Table 2: ROUGE-1, ROUGE-2, and ROUGE-L scores across summarisation strategies (rows) and models (columns).

text length. This limits the Explanation stage's ability to retain causal connections from the rhetorical roles. This, however, does not inhibit the ROUGE scores as the lexical overlap between the chunks is still retained.

Additionally, we validate the utility of our phrase corpus through ablation experiments using extraction-only summaries (Appendix A.2) and observe a notable reduction in vacuous or non-committal model responses (Appendix A.5).

6 Human Evaluation

Metric	EEA	EA	Abstract
Content Coverage	3.47	3.17	3.12
Readability	3.67	3.58	3.40
Coherence	3.92	3.70	3.38
Legal Relevance	3.62	3.17	3.32
Usefulness	3.55	3.32	3.12

Table 3: Mean scores across five qualitative metrics comparing summarisation paradigms.

We carried out a human evaluation with three human experts, to qualitatively assess the effectiveness of each summarisation paradigm.

Table 4 presents results from three legal experts reviewing 20 randomly assigned summaries (10 from each dataset) generated by the Abstract, EA, and EEA methods with the model chosen at random. Reviewers assessed two metrics: (i) Like (personal endorsement for summaries) and (ii) Replace (most appropriate choice to replace the ground truth summary). EEA received the most Like votes and Replace selections, indicating stronger preference over baselines. Pearson correlation for the five evaluation metrics was **0.7053** and Cohen's kappa for Like/Replace was **0.5925**.

Human Evaluation Scores

	Like			Replace		
Reviewer	EEA	EA	Abstract	EEA	EA	Abstract
Reviewer 1	9	5	6	10	6	4
Reviewer 2	10	5	5	11	7	2
Reviewer 3	9	5	6	11	5	4

Table 4: Expert preferences and replacements across summarisation paradigms. "Like" reflects positive endorsements, while "Replace" indicates gold summary replacements.

7 Conclusion

This paper presents the Extract-Explain-Abstract (EEA) framework, a novel three stage approach to address issues in summarisation of Indian legal documents. The "Explain" stage of our framework enhances contextual precision, preserving necessary nuances. The use of SLMs proves to be particularly effective in limited-resource settings like ours. Moreover, our methodology generates the least number of punts outperforming both EA and Abstractive summarisation methods. This promises accurate and context-aware summarisation in resource-constrained environments.

8 Limitations

Our extraction phase relies on legal sentence fragments validated by experts for rhetorical role tagging, ensuring accuracy for Indian legal documents but posing bottlenecks in adapting the framework to new jurisdictions or changing terminologies (Dettmers et al., 2023). The rhetorical roles (Fact, etc.) are optimized for India's common law system, limiting direct applicability to civil law jurisdictions like Germany or Brazil, where legal reasoning follows more structured statutory frameworks. For evaluation, due to resource limitations, we relied on compressed variants of smaller open-source models like LLaMA 1B/3B using QLoRA quantization rather than against state-of-the-art parameter LLMs(Touvron et al., 2023). While this ensured feasibility on commercial hardware, it potentially underestimates the upper bounds of abstractive summarisation quality achievable with larger foundational models. Despite chain-of-thought prompting and RST-based filtering, the abstraction stage introduces nonfactual entity hallucinations (Bommasani et al., 2021)(e.g., misattributed precedents) at times. We also carried out experimentation on the CivilSum dataset (Malik et al., 2024) which proved challenging due to the terse nature of the summaries (avg. of 104 words). The pre-trained SLM cannot achieve a better score than an LLM; however we can finetune SLMs to further improve overall results including for CivilSum. This provides grounds for future work. Due to the availability of only three human evaluators, the number of case-summary pairs submitted for assessment was limited as legal documents are intense in nature. This constraint impacts the extent of our evaluation.

Ethics Statement

This work uses only public domain legal documents and does not make use of any personal data. Our system is intended solely for informational and research purposes and is not a valid substitute for professional legal advice in any way.

Acknowledgements

We would like to express our sincere appreciation to Unsloth for providing an efficient and accessible framework that enabled low-resource fine-tuning and quantization of large language models. Their contributions were instrumental in scaling our experiments across different model sizes while maintaining computational feasibility. We also extend our gratitude to the legal domain experts who played a critical role in constructing and validating the domain-specific phrase corpus. Their expertise ensured the contextual accuracy and legal relevance of the summaries produced.

References

Harkirat Behl et al. 2024. Phi-4 technical report. arXiv preprint arXiv:2412.08905.

S. Bhandari and S. Das. 2023. An extract-then-abstract based method to generate disaster-news headlines using a dnn extractor followed by a transformer abstractor. Information Processing & Management, 60(3).

M. Bommarito and D. Katz. 2022. Legal artificial intelligence: Law and the transformation of legal services. Harvard Journal of Law & Technology.

Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. In Stanford Center for Research on Foundation Models.

Michael Han Daniel Han and Unsloth team. 2023. Unsloth.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Quantization-aware low-rank adapter for efficient finetuning of llms. arXiv preprint arXiv:2305.14314.

J. Devlin, M. Chang, K. Lee, and K. Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In Proc. NAACL-HLT.

M. Dey, R. Das, and S. Ghosh. 2021. Ease: Extractive-abstractive summarization end-to-end using the information bottleneck principle. In New Frontiers in Summarization (NewSum), EMNLP.

M. Elaraby and D. Litman. 2022. Arglegalsumm: Improving abstractive summarization of legal documents with argument mining. arXiv preprint.

G. Erkan and D. R. Radev. 2011. Lexrank: Graph-based lexical centrality as salience in text summarization. arXiv preprint.

A. Farzindar and G. Lapalme. 2004. Letsum, an automatic legal text summarizing system. In Proc. JURIX.

A. Feijo and C. Moreira. 2021. Legalsumm: A portuguese dataset for legal document summarization. Artificial Intelligence and Law.

Andrea Galassi, Marco Lippi, and Paolo Torroni. 2018. Attention-based argument mining for legal documents. In Stefano Ferilli, editor, AI*IA 2018 – Advances in Artificial Intelligence, volume 11298 of Lecture Notes in Computer Science, pages 225–237. Springer, Cham.

H. Jing. 2000. Sentence reduction for automatic text summarization. In Proc. ANLP-NAACL.

T. Kojima et al. 2022. Large language models are zero-shot reasoners. arXiv preprint.

Y. Liu and H. Chen. 2019. Gist: A generative summarization framework for legal text. In Proc. CIKM.

Y. Liu and M. Lapata. 2019. Text summarization with pretrained encoders. arXiv preprint.

Y. Liu, M. Ott, N. Goyal, et al. 2021. Is multi-hop reasoning really explainable? towards benchmarking reasoning interpretability. arXiv preprint.

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. ACM.

William C. Mann and Sandra A. Thompson. 1988. Rhetorical Structure Theory: Toward a functional theory of text organization. Text, 8(3), 243-281.

MetaJure. 2015. Lawyers waste as much as six hours a week on document management issues. Accessed: 2025-05-17.

Marie-Francine Moens, Erik Boiy, Raquel Mochales Palau, and Chris Reed. 2007. Automatic detection of argumentation structure in legal documents. In Proceedings of the 11th international conference on Artificial Intelligence and Law, pages 225–230.

R. Nallapati, B. Zhou, C. Gulcehre, and B. Xiang. 2017. Summarunner: A recurrent neural network based sequence model for extractive summarization. arXiv preprint.

Shounak Paul, Arpan Mandal, Pawan Goyal, and Saptarshi Ghosh. 2023. Pre-trained language models for the legal domain: A case study on indian law. In Proceedings of 19th International Conference on Artificial Intelligence and Law - ICAIL 2023.

L. Polsley, A. Jhunjhunwala, and M. A. Gonzalez. 2016. Case summarizer: A system for automated summarization of legal texts. In Proc. COLING.

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics.

M. Saravanan, B. Ravindran, and S. Raman. 2006. Improving legal document summarization using graphical models. Knowledge and Information Systems.

Muthiah Saravanan, Balaraman Ravindran, Sanjay Raman, and K. Ghosh. 2008. Automatic identification of rhetorical roles using conditional random fields for legal document summarization. AI and Law, 16(2):109–129.

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 (Volume 1: Long Papers), pages 1048–1064, Online only. Association for Computational Linguistics.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Baptiste Rozière, Naman Goyal, Suchin Fu, Yacine Magloire, Sylvain Lam de Scaillat, et al. 2023. Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971.

Pawan Trivedi, Digha Jain, Shilpa Gite, and Ketan Kotecha. 2023. Indian legal corpus (ilc): A dataset for summarizing indian legal proceedings using natural language. Engineered Science, 21:100–110.

Laurens Van Der Maaten et al. 2024. The llama 3 herd of models. arXiv preprint arXiv:2407.21783.

Fali Wang, Zhiwei Zhang, Xianren Zhang, Zongyu Wu, TzuHao Mo, Qiuhao Lu, Wanjing Wang, Rui Li, Junjie Xu, Xianfeng Tang, Qi He, Yao Ma, Ming Huang, and Suhang Wang. 2024. A comprehensive survey of small language models in the era of large language models: Techniques, enhancements, applications, collaboration with llms, and trustworthiness. arXiv preprint arXiv:2411.03350.

H. Wang, Z. Yang, and H. Yin. 2023. Element-aware summarization with large language models: Expert-aligned evaluation and chain-of-thought method. arXiv preprint.

X. Wang et al. 2022. Large language models are reasoning teachers. arXiv preprint.

J. Wei et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. arXiv preprint.

An Yang et al. 2023. Qwen technical report. arXiv preprint arXiv:2309.16609.

W. Zheng and M. Lapata. 2019. Sentence centrality revisited for unsupervised summarization. In Proc. ACL.

A Appendix

A.1 Rhetorical Roles

Following prior work and consultations with legal domain experts, we incorporate the following rhetorical roles to structure and analyse legal summaries. These roles capture the functional components of legal reasoning commonly found in judicial documents:

- *Fact*: This role refers to the factual background of the case, including events, entities, and circumstances that led to the legal dispute.
- *Issue*: The issue identifies the central legal question or controversy under judicial consideration.
- *Argument*: This role captures the claims and counterclaims presented by the parties involved.
- *Statute*: This denotes references to specific legislative provisions, including constitutional articles and regulations that are applicable to the case.
- *Precedent*: This includes citations to prior judicial decisions that bear relevance to the present case.
- *Ruling by Lower Court*: This role outlines the findings and decisions made by subordinate courts in earlier stages of litigation.
- *Rationale behind Present Court's Ruling*: This captures the reasoning and justification provided by the current court in arriving at its decision.

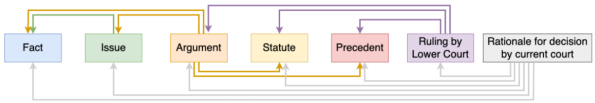

Figure 1: Dependency flow across rhetorical roles: facts lead to issues, which frame arguments, supported by statutes and precedents, culminating in the court's rationale.

A.2 Extraction

To test the accuracy of our phrase corpus, we generate purely extractive summaries of the datasets. Each sentence is scored by checking for specific linguistic cues and phrase patterns that correspond to summary categories. These categories are derived by mapping codes to groups: codes `Facts` and `Issues` indicate *Introduction*; `Argument` and `Ruling by Lower Court` indicate *Context*; `Statute` and `Precedent` indicates *Analysis*; and `Rationale behind by the Present Court's Ruling` indicates *Conclusion*. Sentences containing cue phrases associated with these codes are scored accordingly, reflecting their relevance to each category. To capture the informativeness of sentences, a TF-IDF score is computed for each word in the document. The final score for each sentence within a category is the sum of its category relevance (from cue phrase matches mapped via codes) and its TF-IDF importance. This combination balances thematic relevance and lexical content quality. Top-ranked sentences from each category are selected to build a coherent summary covering all key sections of the document.

To create the final summary, we divide the total allowed length according to set proportions for each category: **10%** for the Introduction, **24%** for the Context, **60%** for the Analysis, and **6%** for the Conclusion (Shukla et al., 2022). These proportions are applied to a target summary length of about one-third (34%) of the original document's total word count. For each category, we then pick the highest-scoring sentences until we reach the allocated word limit. Thus, the summary fairly and evenly covers all important parts of the document. To avoid repetition, once a sentence is included, it will not be added again under another category. This method helps keep the summary clear, well-structured, and informative.

Dataset	ROUGE-1	ROUGE-2	ROUGE-L
ILC	0.5024	0.2539	0.2530
IN-Abs	0.5620	0.3060	0.2610

Table 5: Performance comparison of extraction-only methods on ROUGE metrics across ILC and IN-Abs datasets.

A.3 Implementation Details

We employ chunking of texts into chunks of size equal to the maximum input length (`chunk size + max_new_tokens`) that the model was trained on. The first chunk would contain `chunk size` words without breaking sentences. If the last sentence of the chunks exceeds the limit, it is put in the next chunk. Each chunk is summarised and the summaries are concatenated. The combined summary undergoes abstraction to generate the final summary. Chunking size set for models are as

follows:

- LLaMa 3.2 1B: 128k
- LLaMa 3.2 3B: 128k
- Qwen 2.5 7B: 32k
- Phi-4 14B: 16k

For the **ILC** and **IN-Abs** datasets, we query the model using the following prompt template:

```
You are a proficient legal assistant
specializing in summarising legal texts.
Summarise the following judgement in
{0.5 * len(summary_chunk)} words or less,
include relevant information and only
output the summary nothing else.
```

For both configurations, we use a temperature of **0.6** and a top-p value of **0.9** for decoding. Sampling is disabled with `do_sample = False` to ensure deterministic outputs. We also set `max_new_tokens = 5000` to accommodate the length of legal summaries.

We employ the following prompt template:

```
 You are a proficient legal assistant.
Your goal is to generate a clear,
concise, and accurate summary that
logically flows from the facts through
to the ruling.  Follow the steps
below, making sure to use the tagged
information to guide your summary:
1. **Key Facts**: Highlight critical
facts and their impact.
2. **Issues**: Explain legal issues and
their factual basis.
3.  **Arguments**:  Summarise key
arguments from all sides.
4. **Lower Court Rulings**: Explain
prior rulings with relevant connections.
5. **Statutes**: Clarify statutes and
their application.
6. **Precedents**: Discuss relevant
precedents.
7.  **Present Court's Ruling**:
Summarise final ruling with integrated
reasoning.
```

This prompt explicitly leverages rhetorical role annotations to structure the generated summary, ensuring that all critical legal components are adequately represented.

A.4 Human Evaluation

The human evaluation was conducted on the basis of following parameters:
- **Content Coverage** ensures the summary has all key legal points and remains factually accurate.
- **Readability** measures clarity, simplicity, and grammatical quality.
- **Coherence** assesses logical flow and smooth transitions between ideas.
- **Legal Relevance** ensures only pertinent legal information is included, excluding irrelevant details.
- **Usefulness** reflects the summary's practicality and trustworthiness for legal professionals.

The evaluators, comprising three trained legal experts, with LL.B.s from Pravin Gandhi College of Law, provided detailed qualitative feedback. Summaries generated by Abstract method slightly lacked in legal relevance while EA method often had redundant information in its summary which hindered readability. In contrast, the EEA model produced summaries that demonstrated stronger alignment with legal facts and usefulness, while preserving overall readability. These observations highlight the inherent complexity of legal summarisation. As per reviewers, the IN-Abs and ILC dataset summaries had enough room to fit legally relevant information, adding to their usefulness. Overall the reviewers preferred the EEA rhetorical structure of the summaries for its legal relevance. They stated that segmented text made the summaries more readable, but they often contained repetitive information.

A.5 Punts and Text Degeneration

Dataset	IN-Abs				ILC			
	M-1	M-2	M-3	M-4	M-1	M-2	M-3	M-4
Abstract	10	12	0	0	368	4	0	0
EA	24	4	0	0	275	3	0	0
EEA	2	4	0	0	36	2	0	0

Table 6: Punts counts across summarisation strategies (rows) and models (columns).
M-1: LlaMa3.2-1B; M-2: LlaMa 3.2-3B; M-3: Qwen-2.5-7B; M-4: Phi-4-14B

In the context of legal summarisation, it is essential to evaluate not only the informativeness and accuracy of the generated summaries but also their appropriateness. Two undesirable phenomena that often arise in language model outputs are punts and text degeneration, both of which can undermine the usefulness of a summary in legal settings.
Punts refer to generic, non-committal responses generated by a model when it is uncertain or lacks specific knowledge.

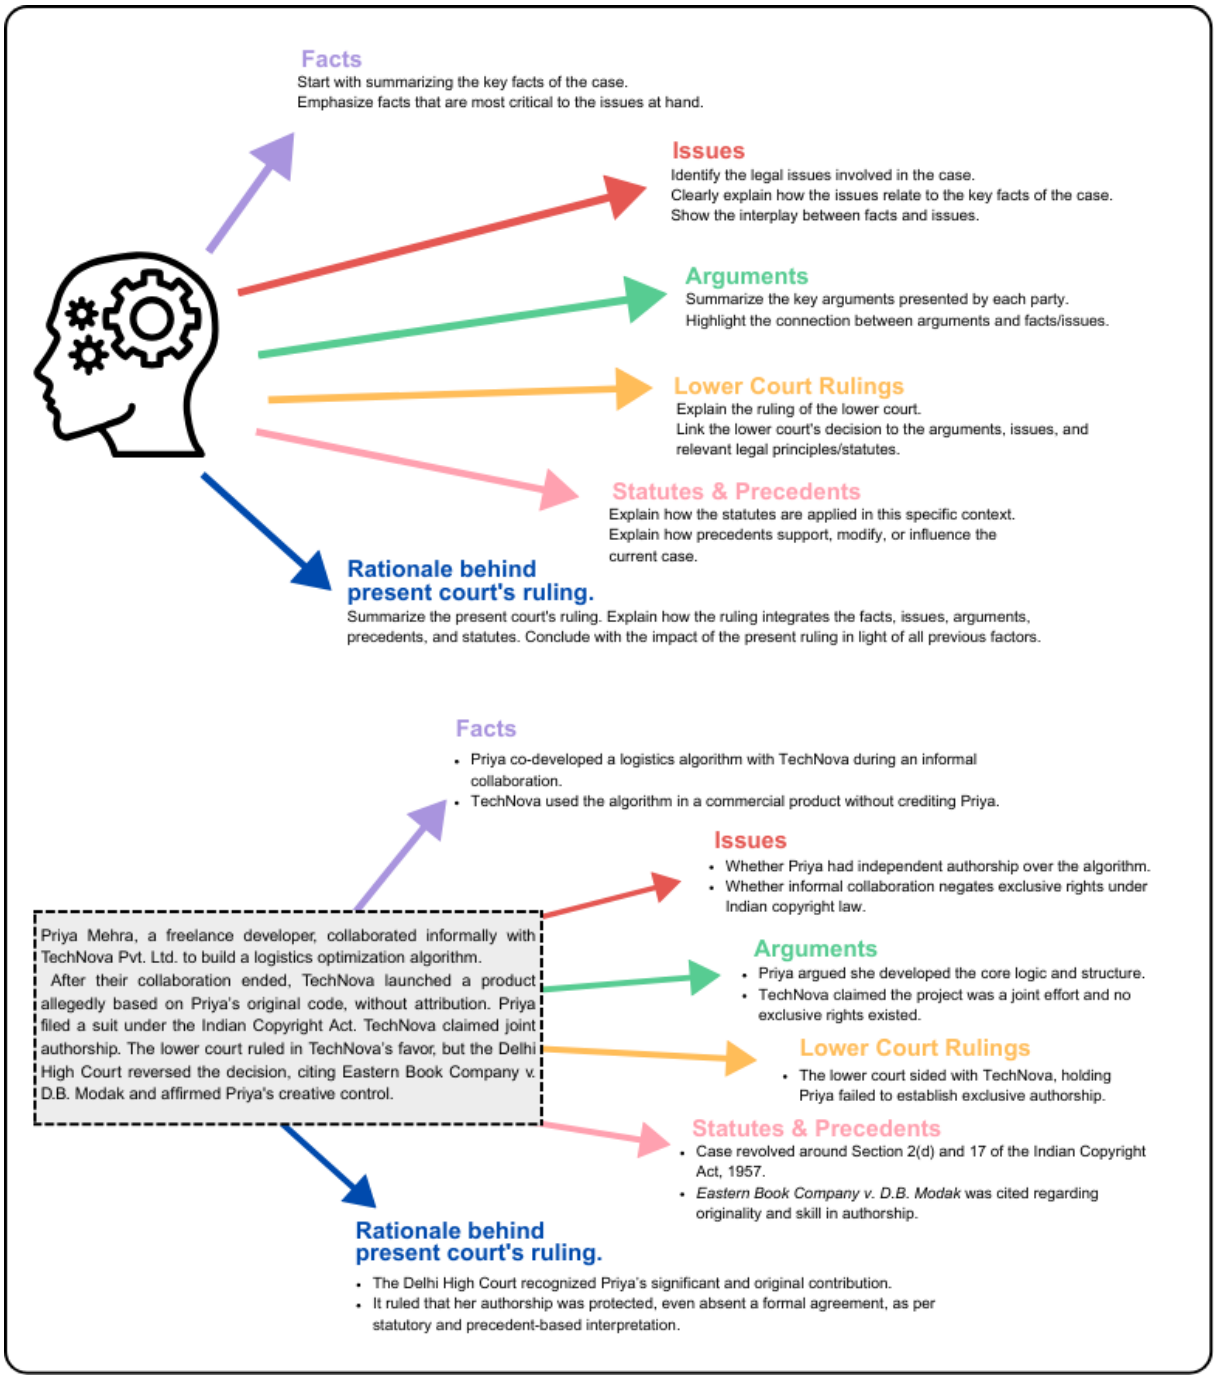

Figure 2: CoT reasoning with rhetorical roles.

Text degeneration refers to pathological patterns in generation, such as repetitive or looping phrases, which can occur due to poor decoding strategies or exposure bias in training. For detecting punts, we employed transformer-based sentence embeddings from https://huggingface.co/sentence-transformers/paraphrase-mpnet-base-v2 (Reimers and Gurevych, 2019) to measure semantic similarity. We calculated cosine similarity of responses with respect to generic punts prevalent, such as `I cannot provide a summary of the judgment` or `Please consult a lawyer`. Threshold was set to 0.65. For detecting text degeneration, we utilised combinations of Jaccard similarity, n-gram repetition and intra-summary cosine similarity using the sentence embeddings. However, we noticed that there were more false positives than observed in the responses, even with high confidence thresholds. Hence, we do not report these scores.

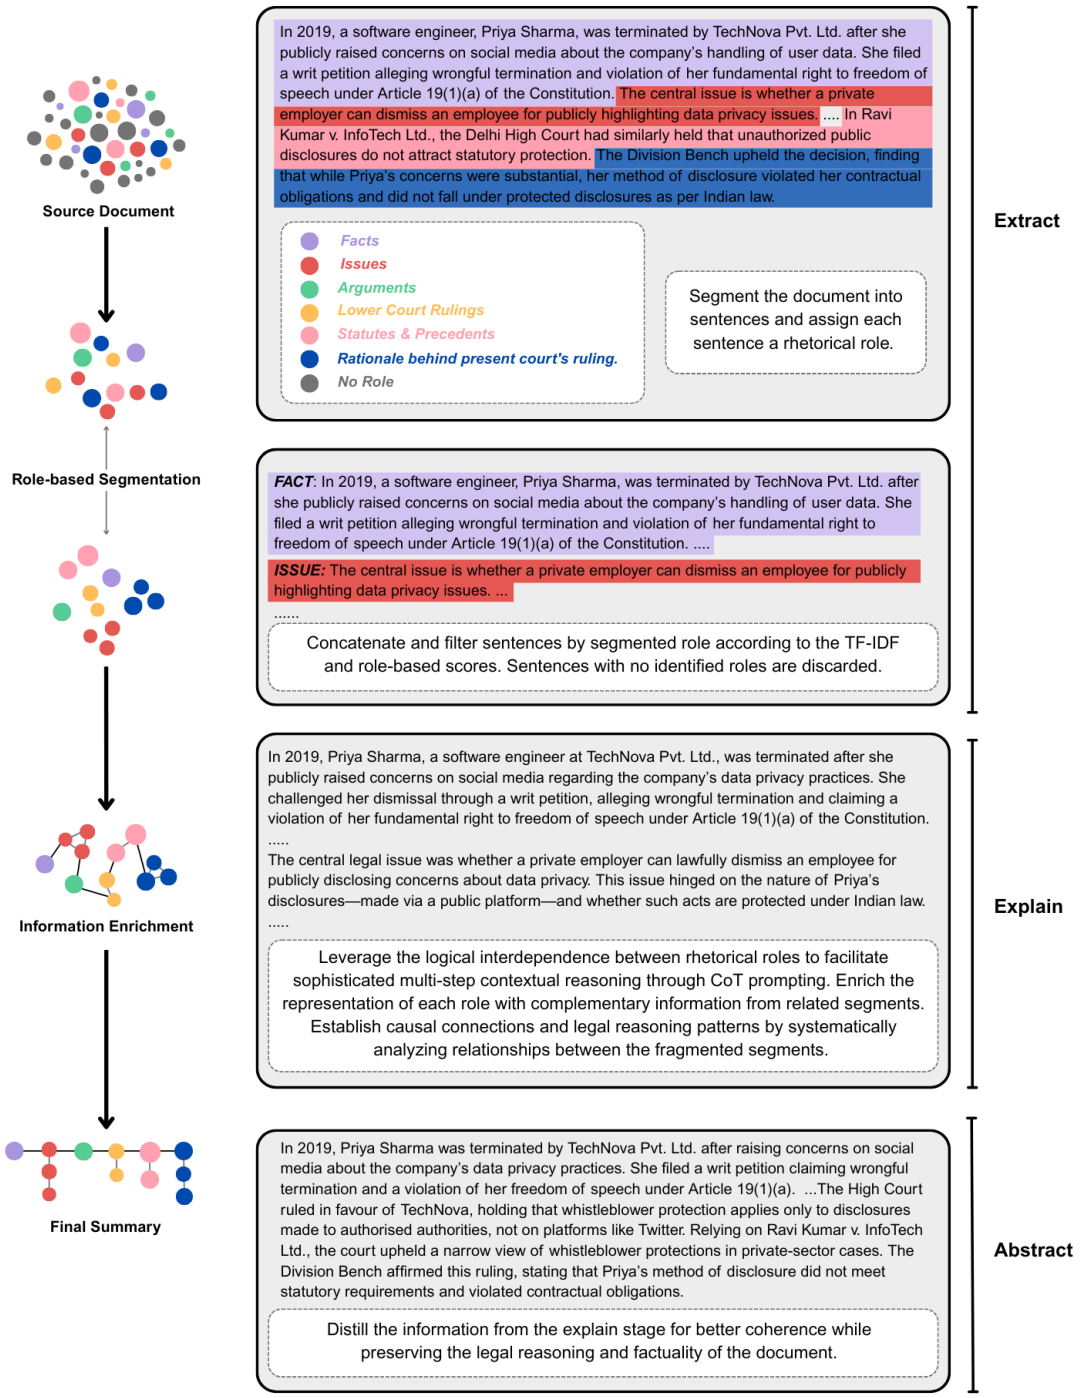

Figure 3: Extract-Explain-Abstract Framework.

A.6 Summary Outputs

We provide sample summary outputs for the two datasets: IN-Abs and ILC generated using three models—LlaMa 3.2-3B, Qwen-2.5-7B, and Phi-4-14B. We have included a reference summary and three summaries generated respectively: tagged (EEA), tagless (EA) and original (abstractive) summary. Table 7 shows the summary sample of IN-Abs dataset using Qwen-2.5-7B and table 8 shows the summary sample of ILC dataset using Phi-4-14B.

Reference Summary

Proceedings were commenced under Chapter III B of the Rajasthan Tenancy Act 1955 for determining the ceiling area for Raghubir Singh, the father of the appellant. In a revision application in the first round of litigation. The Board of Revenue directed that the ceiling area for Raghubir Singh may be determined according to the old law, i.e. Act of 1955, and not according to the Rajasthan Imposition of Ceiling on Agricultural Holdings Act, 1973. Thereafter, by his order dt. 5.5.76, the SDO held, inter alia, that tile father and the son, appellant, constituted two separate units and each of them was entitled to get 62 bighas and 8 biswas. Aggrieved by the SDO 's order, Raghubir Singh filed an appeal before Revenue Appellate Authority, which was dismissed. He then filed a revision application under S.230 of the Act of 1955 before the Board of Revenue for Rajasthan. He urged that he was in possession of 112 bighas only and sought permission to adduce additional evidence, which was disallowed. However, the Member of the Board held inter alia that the provisions of the old law applied to the case, but the SDO had committed an error of law in determining the ceiling area under the new Act of 1973. The Board further held that there is no provision for separate units in Chapter III B of the Act of 1955, and remanded the case to the SDO, Hanumangarh, for fresh determination of the Ceiling area for Raghubir Singh. Before the High Court, it was urged that the State having not appealed against the order of the SDO dated 5.5.76, it became final, and the Board of 723 Revenue had no jurisdiction to set aside that part of the SDO 's order which had gone against the State and in his favour. It was submitted that the power of general superintendence and control over all revenue courts which vested in the Board could not be exercised to the detriment of the writ petitioners much less by way of suo motu exercise of powers. The High Court dismissed the petitions. Before this court it was submitted that the State not having filed an appeal, or a revision, the Board of Revenue could not, while hearing the revision petition of Raghubir Singh, set aside the orders of the SDO and Revenue Appellate Authority under S.221 of the Act of 1955. Dismissing the appeal, this court, HELD:. .(1) S.221 is not subject to the other provisions of the Act. It is clear from the language of Section 221 of the Act 1955 that the Board of Revenue has general powers of superintendence and control over all revenue courts. It is both administrative as well as judicial powers. It is open to the Board to exercise its powers of superintendence on all its subordinate courts in order to regulate the functioning of the subordinate courts so as to keep them within their respective spheres of jurisdiction. If the subordinate court disregards any specific provision of law and does something illegal it is open to the Board of Revenue to interfere and set the matter right. (727 E) Karan Singh vs Board of revenue, Rajasthan and Permessar Singh vs Kailaspati AIR 1916 Pat. 292 (FB), distinguished. Kana vs Board of Revenue ILR , approved. (2) There is no restriction on the powers of the Board to set aside the order of the SDO provided it comes to the conclusion that interest of justice requires exercise of such powers. (728 D) (3) In not determining the ceiling area according to the Act of 1955, the SDO committed a grave illegality in not merely ignoring the law but also ignoring the directions of the Board of Revenue itself. (728 F) 724

EEA Summary

ROUGE-1: 0.6433; ROUGE-2: 0.3424; ROUGE-L: 0.3356; InLegalBERT: 0.9737; BertScore: 0.8726

Summary

Key Facts

The case concerns the determination of the ceiling area for Raghubir Singh under Chapter III B of the Rajasthan Tenancy Act (Act No. 3 of 1958). Initially, the Sub Divisional Officer, Hanumanggarh, determined the ceiling area. Raghubir Singh appealed this decision to the Revenue Appellate Authority, Bikaner, which set aside the Sub Divisional Officerâ€™s order and remanded the case. The Board of Revenue for Rajasthan then directed that the ceiling area should be determined according to the old law, i.e., the Act of 1958, rather than the new law, the Rajasthan Imposition of Ceiling on Agricultural Holdings Act, 1973. Subsequently, the Sub Divisional Officer held that Surendrapal Singh, a minor son of Raghubir Singh, was in possession of ancestral land. Raghubir Singh appealed this order to the Revenue Appellate Authority, which upheld it. However, the Board of Revenue, upon reconsideration, found that Surendrapal Singh was only a minor and remanded the case for a fresh determination of the ceiling area under the old law.

Issues

The primary issue in this case is whether the Board of Revenue had the jurisdiction to set aside the order of the Sub Divisional Officer and the Revenue Appellate Authority, given that the State did not file an appeal or a revision. Additionally, the case raises questions about the application of Section 9 of the Rajasthan Tenancy Act, 1958, which vests general superintendence and control over all revenue courts and officers in the Board of Revenue.

Arguments

- **Raghubir Singh**: Argued that since the State did not file an appeal or a revision, the order of the Sub Divisional Officer became final, and thus the Board of Revenue lacked jurisdiction to set aside that part of the order favorable to Raghubir Singh.
- **Appellant's Counsel**: Contended that the Board of Revenue had the executive powers of superintendence and control over all revenue courts and officers, and could therefore set aside the orders of the Sub Divisional Officer and the Revenue Appellate Authority if it deemed it necessary for the interest of justice.

Lower Court Rulings

- The Revenue Appellate Authority upheld the order of the Sub Divisional Officer.
- The Board of Revenue, upon reconsideration, found that Surendrapal Singh was a minor and remanded the case for a fresh determination of the ceiling area under the old law. - The Board also held that the Sub Divisional Officer had committed an error of law by applying the new law instead of the old law.

Statutes

- **Section 9 of the Rajasthan Tenancy Act, 1958**: Provides that the general superintendence and control over all revenue courts and officers shall be vested in the Board of Revenue. This section is analogous to Section 9 of the Government of India Act, which grants superintendence over all courts subject to its appellate jurisdiction.
- **Board of Revenue's Argument**: Claimed that this provision allowed it to set aside the order of the Sub Divisional Officer if it deemed it necessary for the interest of justice.

Precedents

- **Karan Singh vs Board of Revenue, Rajasthan**: The Division Bench of the Rajasthan High Court held that the Board of Revenue could not exercise its powers of superintendence and control on a judicial side, as it had appellate jurisdiction. However, the High Court held that the Board could still exercise its powers of superintendence and control in appropriate cases.

Present Court's Ruling

- The present court dismissed the appeal, holding that the Board of Revenue had the jurisdiction to set aside the order of the Sub Divisional Officer and the Revenue Appellate Authority based on the provisions of Section 9 of the Rajasthan Tenancy Act, 1958, and the analogy with Section 9 of the Government of India Act.
- The court found that the Board's action was justified in the interest of justice, as there was no provision for separate units in Chapter III B of the Act, and the case should be remanded for a fresh determination of the ceiling area.

EA Summary

ROUGE-1: 0.6352; ROUGE-2: 0.3237; ROUGE-L: 0.3520; InLegalBERT: 0.9681; BertScore: 0.8733

Summary Case Background

The case concerns the determination of the ceiling area for Raghubir Singh under Chapter III B of the Rajasthan Tenancy Act (Act No. 3 of 1950). Initially, the Sub Divisional Officer, Hanumangarh, determined the ceiling area. However, an appeal by Raghubir Singh led the Revenue Appellate Authority to set aside the order and remand the case.

Revision Application

Raghubir Singh subsequently filed a revision application before the Board of Revenue for Rajasthan. The Board directed that the ceiling area should be determined according to the old law (Act of 1950) rather than the new law (Rajasthan Imposition of Ceiling on Agricultural Holdings Act, 1973).

Sub Divisional Officer's Order

In response to the Boardâ€™s direction, the Sub Divisional Officer, Hanumanggarh, issued an order on May 5, 2, determining that Surendrapal Singh (the appellant herein) was a minor son of Raghubir Singh and that the land in possession of Raghubir Singh was ancestral. Raghubir Singh appealed this order to the Revenue Appellate Authority, Bikaner, which upheld the Sub Divisional Officer's order.

Board of Revenue's Decision

The Board of Revenue, however, concluded that Surendrapal Singh was actually 19 years old at the time the ceiling proceedings were finalized, and thus the provisions of the old ceiling law applied. The Board found that the Sub Divisional Officer had committed an error of law by applying the new law. Consequently, the Board set aside the Sub Divisional Officer's and the Revenue Appellate Authority's orders and remanded the case for a fresh determination of the ceiling area.

High Court Writ Petition

Raghubir Singh challenged the Board's decision in a writ petition before the Division Bench of the High Court. He argued that since the State did not appeal the Sub Divisional Officer's order, it became final, and the Board lacked jurisdiction to set aside the order in favor of the State.

Arguments

- **Appellant's Counsel:** The learned senior counsel for the appellant argued that the Board's powers of superintendence and control under Section 9 of the Act of 1950 do not extend to judicial review. He cited the precedent of Karan Singh vs Board of Revenue, Rajasthan, where the High Court held that the Board cannot exercise its powers of superintendence and control on a judicial side.
- **Statutory Interpretation:** The learned counsel contended that Section 9 of the Act of 1950 vests the Board with general superintendence and control over all revenue courts and officers, and there is no restriction on the Board setting aside the order of the Sub Divisional Officer if it deems it necessary for the interest of justice.

Statutory Provisions

- **Section 9 of the Act of 1950:** Provides for the general superintendence and control over all revenue courts and officers.
- **Section 14 of the Government of India Act, 1935:** Grants superintendence over to High Courts over all inferior courts. - **Section 15 of the Government of India Act, 1935:** Clarifies that superintendence does not extend to questioning judgments of inferior courts unless they are subject to appeal or revision.

Decision

The Division Bench of the High Court dismissed the appeal, holding that the Board of Revenue had the authority to set aside the Sub Divisional Officer's order based on the interest of justice, as there was no restriction in Section 15 of the Act of 1951 that would prevent the Board from doing so.

Conclusion

The case highlights the interplay between the Board's powers of superintendence and control and the limitations on judicial review. The Board of Revenue was found to have the authority to set aside the Sub Divisional Officer's order despite the absence of an appeal from the State, as the interest of justice required such action.

Abstractive Summary
ROUGE-1: 0.5911; ROUGE-2: 0.2976; ROUGE-L: 0.3234; InLegalBERT: 0.9372; BertScore: 0.8333
Summary of the Judgment **Case Details:** - **Appeal No.:** 1367 of 1980 - **High Court Decision:** Judgment and Order dated 19th July, 1979 in D.B. Civil Writ Petition Nos. 681 of 1978 and 769 of 1979 - **Appellants:** Subhash Mallick and S.K. Jain for the appellant (Surendrapal Singh) - **Respondents:** B.D. Sharma for the respondents (Board of Revenue for Rajasthan) **Background:** - **Writ Petitionsions Filed:** By Raghubir Singh (father) and Surendrapal Singh (his minor son) through his next friend, Shri Jagjit Singh, maternal grandfather. - **Relief Claimed:** Setting aside the order dated 6th August, 1977 passed by the Board of Revenue for Rajasthan, Ajmer. **Facts:** - **Ceiling Proceedings:** Initiated under Chapter III B of the Rajasthan Tenancy Act (Act No. 3 of 1955). - **Sub Divisional Officer's Order (10th Aug., 1972):** Determined the ceiling area for Raghubir Singh. - **Revenue Appellate Authority's Order (6th Mar., 1973):** Set aside the Sub Divisional Officer's order and remanded the case. - **Board of Revenue's Order (14th Apr., 1975):** Directed that the ceiling area be determined according to the old law (Act of 1955) rather than the new law (Act of 1973). - **Sub Divisional Officer's Final Order (5th May, 1976):** Held that Surendrapal Singh (appellant) was a minor and used to study in Punjab, and that the land in possession of Raghubir Singh was ancestral. Concluded that Raghubir Singh and his son constitute two separate units, each entitled to 62 Bighas and 8 Biswas, resulting in a total of 124 Bighas and 16 Biswas retained, with 4 Bighas and 6 Biswas to be resumed. - **Appeal to Revenue Appellate Authority (6th Sep., 1976):** Dismissed the appeal and upheld the Sub Divisional Officer's order. - **Revision Application to Board of Revenue (14th Apr., 1975):** Found that Surendrapal Singh was 13 years old when the ceiling proceedings were finalized, and that the Sub Divisional Officer erred in applying the new law. Held that there is no provision for separate units in Chapter III B of the Act of 1955, and set aside the Sub Divisional Officer's and Revenue Appellate Authority's orders, remanding the case for a fresh determination of the ceiling area. **High Court Decision:** - **Contention of Appellant:** The State did not file an appeal under Section 225 of the Act of 1955 or a revision under Section 230 of the Act of 1955, so the Board of Revenue had no jurisdiction to set aside the Sub Divisional Officer's order. - **Section 221 of the Act of 1155:** Provides for the general superintendence and control over all revenue courts, including the Sub Divisional Officer. - **Interpretation of Section s221:** The Board of Revenue has both administrative and judicial powers to oversee the functioning of subordinate courts and to correct illegal actions. - **Application of Section s221:** The Board of Revenue correctly exercised its powers by setting aside the Sub Divisional Officer's order due to the illegal application of the new law and the absence of a provision for separate units. - **Conclusion:** The High Court upheld the Board of Revenue's decision and dismissed the writ petitions. **Judgment:** - **Board of Revenue's Order:** Correctly exercised its powers under Section s221 of the Act of 1155 to set aside the Sub Divisional Officer's order. - **Interest of Justice:** The Board of Revenue acted in the interest of justice by correcting the illegal application of the new law and remanding the case for a fresh determination. - **Costs:** The appeal is dismissed with costs. **Outcome:** - **Appeal Dismissed:** The High Court's judgment is affirmed, and the appeal is dismissed with costs.

Table 7: Sample reference and generated summaries from the IN-Abs test set (Document ID = 7109, Model = Qwen-2.5-7B).

Reference Summary

NGOs shouldnâ€™t use the terms â€œCentralâ€, â€œStateâ€, and â€œNationalâ€ with their names so as to avoid the confusion that itâ€™s a statutory body run by either Central or State Government. Kerala High Court gave the judgment stating the above-cited reasons in the case of State Environment Protection Council vs. State of Kerala [WP(C).No.2366 of 2015(S)] headed by the bench of Honâ€™ble Chief Justice S. Manikumar and Justice Shaji P. Chaly. In the above-cited case, the State Environment Protection Council (Petitioner) of Kannur had filed a writ petition of certiorari to quash the building permit for constructing additional six floors in the building. Writ of Mandamus was also filed against the respondent for illegal construction on the building. According to the facts, Respondent had given the permit to KP Mohammad Ashraf, MD (Global Village, Kannur) in violation of Rules 34(2) and 31(2) of the Kerala Municipal Building Rules, 1999. As a result, an inquiry had been conducted to revoke the building permit and partial occupancy certificate (which was already given). In this matter, instead of taking necessary actions, excuses were made by the respondents. In addition to the facts, the petitioner also contended that the illegal construction of 6 storey building on the side of the National Highway would cause a severe threat to the life of the people. As an effect, an affidavit was filed by the town planner of Kannur stating that building construction had violated the rules of Kerala Municipal Building Rules, 1999 including: In response to the above affidavit, an application had been made by KP Mohammad Ashraf (Respondent) for regularizing the construction of the building, for which the learned counsel of the petitioner had argued that since the rules were violated, further construction of the building should not commence. The learned counsel of the respondents had argued that whether such private NGO based in Kannur has the right to file such writ petition challenging the construction of the building. Arguments were held regarding the name of the petitioner where, it was named as State Environment Protection Council, therefore, giving an impression that itâ€™s been run by the Government or itâ€™s a statutory body. On the contrary, itâ€™s just a private NGO. Court held that, while registering, no private body should give an impression that itâ€™s a statutory body. HC stated that â€œNGOs or associations or societies, should not give any impression, to the public at large that, it is a statutory body, under any enactment, State / Central, as the case may be. Needless to say that NGOs or associations, or societies, registered under the Kerala Societies Registration Act, 1860, as far as possible, should avoid, using the name â€~Centralâ€™ or â€~Stateâ€™ or â€~Nationalâ€. HC stated that regarding the alleged construction of the building and violations, the public shouldnâ€™t be misled by the institution in the exercise of the powers under Article 226 of the Constitution of India. HC directed that the NGOs shouldnâ€™t use the word â€œCentralâ€, â€œStateâ€, and â€œNationalâ€ with their names under the provisions of the Kerala Societies Registration Act, 1860 because it might mislead the public at large resulting in confusions and unauthorized use of the powers.

EEA Summary

ROUGE-1: 0.5614; ROUGE-2: 0.1875; ROUGE-L: 0.2436; InLegalBERT: 0.9446 ; BertScore: 0.8450

Summary of the Case: State Environment Protection Council v. Kannur Municipality
Key Facts:
The State Environment Protection Council, Chelad Kannur, filed a writ petition challenging the issuance of a revised building permit by the Kannur Municipality. The permit allowed for additional construction on a building owned by Mr. K. P. Mohammed Ashraf, Managing Director of Emad Building Global Village, Kannur. The petitioner argued that the permit was issued in violation of the Kerala Municipal Building Rules and posed a threat to public safety due to its proximity to a National Highway.
Issues:
The primary legal issue was whether the Kannur Municipality acted within its authority and in compliance with the Kerala Municipal Building Rules when it issued the revised building permit. Additionally, the case raised concerns about the potential threat to public safety due to unauthorized construction near a National Highway.
Arguments:
- **Petitioner's Argument:** The petitioner contended that the building permit was issued in violation of specific rules, particularly Rules 34 and 31, and that the construction posed a severe threat to public safety. They sought a writ of certiorari to quash the permit and a mandamus to demolish the unauthorized construction.
- **Respondent's Argument:** The Kannur Municipality argued that the permit was issued based on an oversight and that steps were being taken to regularize the construction. They also noted that violations had been rectified and that the construction was being regularized under the Kerala Municipality Building (Regularization of Unauthorized Construction) Rules.
Lower Court Rulings:
The lower court had directed the Kannur Municipality to revoke the building permit and partial occupancy certificate. However, the Municipality issued a subsequent notice stating that the permit was issued due to an oversight and that regularization was underway.
Statutes:
The case involved the Kerala Municipal Building Rules, particularly Rules 34 and 31, which govern building permits and construction standards. The Kerala Municipality Building (Regularization of Unauthorized Construction) Rules were also relevant, as they provided a framework for regularizing unauthorized constructions.
Precedents:
The court considered precedents related to the issuance of building permits and the authority of municipal bodies to regularize unauthorized constructions. The court also addressed the role of NGOs in representing public interest without misleading the public about their statutory status.
Present Courtâ€™s Ruling:

The High Court of Kerala, presided over by Chief Justice S. Manikumar and Justice Shaji P. Chaly, dismissed the writ petition. The court noted that the petitioner, an NGO, should not mislead the public about its statutory status. The court also observed that the application for regularization of the construction was pending and directed the Inspector General of Registration to decide on the registration of NGOs using terms like "Central" or "State." The court disposed of the petition, emphasizing the need for NGOs to avoid giving an impression of being statutory bodies.
This ruling underscores the importance of compliance with municipal regulations and the role of NGOs in public interest litigation.

EA Summary

ROUGE-1: 0.5496; ROUGE-2: 0.1674; ROUGE-L: 0.2414; InLegalBERT: 0.9386; BertScore: 0.8372

Summary of the Judgment
Case Overview:
- **Court:** High Court of Kerala at Ernakulam
- **Date:** January 19, 2016
- **Judges:** Hon'ble Chief Justice S. Manikumar and Hon'ble Mr. Justice Shaji P. Chaly
- **Parties Involved:**
- **Petitioner:** State Environment Protection Council, Chelad Kannur, represented by its President, Satheesh Kumar Pamban S/O Late T. Raghavan.
- **Respondents:**
- State of Kerala, represented by its Principal Secretary, Secretariat, Thiruvananthapuram.
- Secretary, Kannur Municipality.
- Town Planner, Department of Local Self Government Institutions, Kannur.
- Mr. K. P. Mohammed Ashraf, Managing Director, Emad Building Global Village, Kannur.
Facts:
The petitioner, State Environment Protection Council, sought judicial intervention to address unauthorized construction activities by Mr. K. P. Mohammed Ashraf. The petitioner requested the court to issue a writ of certiorari to quash a revised building permit (Exhibit P2) that allowed the addition of six floors and a report (Exhibit P4) directing the regularization of the construction. Additionally, the petitioner sought a mandamus to compel the Kannur Municipality to demolish the unauthorized constructions.
The case arose from the issuance of a building permit by the Secretary of Kannur Municipality, which was allegedly in violation of certain rules. Despite a directive to revoke the permit, it was issued, leading to an inquiry by the Town Planner Department. The inquiry revealed several violations, including insufficient open space and parking spaces. The construction exceeded the permitted limits, although the owner had increased the plot area and obtained necessary certificates. An application for regularization of the unauthorized construction was pending.
Arguments:
The petitioner argued that the construction posed a significant threat to public safety, particularly due to the proximity of a proposed multiplex near a national highway. The Town Planner's counter-affidavit acknowledged the violations and the pending regularization application. The learned counsel for the petitioner suggested that the writ petition be disposed of due to the pending regularization application.
Judgment:
The court addressed two primary issues in its judgment. Firstly, it noted that non-governmental organizations (NGOs) or associations should not mislead the public by using terms like "Central" or "State" in their names, as this could create the false impression of being statutory bodies. The court directed the Inspector General of Registration to decide on the registration of such bodies within two months.
Secondly, the court disposed of the writ petition, considering the pending regularization application for the unauthorized construction. The court emphasized that the pending application for regularization should be resolved before any further action.
Directives:
The court issued specific directives to address the broader issue of NGO nomenclature and the specific case of unauthorized construction: 1. The Inspector General of Registration, Kerala, was directed to make a decision on the registration of NGOs or associations using terms like "Central" or "State" within two months.
2. The court disposed of the writ petition, acknowledging the pending regularization application for the unauthorized construction.
3. The registering authority was instructed to ensure that no private body gives the impression of being a statutory body.
This judgment not only addressed the specific case of unauthorized construction but also highlighted the importance of accurate representation by NGOs and associations to prevent public deception. The court's directives aimed to ensure compliance with regulatory standards and prevent misleading nomenclature by private bodies.

Abstractive Summary

ROUGE-1: 0.4680; ROUGE-2: 0.1647; ROUGE-L: 0.2442; InLegalBERT: 0.9382; BertScore: 0.8451

Summary of Judgment: WP(C).No.2366 OF 2015(S)
Court: High Court of Kerala at Ernakulam
Date: 19th January 2021
Judges: Hon'ble Chief Justice S. Manikumar and Hon'ble Mr. Justice Shaji P. Chaly
Background:

The State Environment Protection Council, Chelad Kannur, filed a writ petition challenging a building permit issued by the Kannur Municipality. The petitioner argued that the permit violated the Kerala Municipal Building Rules 1999 and sought its quashing, along with a directive for the demolition of unauthorized constructions.
Key Issues:
1. **Building Permit Violation:** The petitioner contended that the permit issued to Mr. K. P. Mohammed Ashraf for constructing additional floors was in violation of municipal rules.
2. **Unauthorized Construction:** The construction was deemed a public safety threat due to its proximity to a national highway.
3. **Regulatory Oversight:** The Town Planner Department had directed the revocation of the permit, but the Municipality did not comply.
Counter Affidavit:
The Town Planner of Kannur provided a counter-affidavit detailing: - The original permit allowed a 7-story building, later revised to permit two additional floors.
- Noted violations included insufficient open space and parking. - Efforts were made to rectify some violations, and a partial occupancy certificate was issued with conditions.
- An application for regularization of the unauthorized construction was pending under the Kerala Municipality Building (Regularization of Unauthorized Construction) Rules 2014.
Court's Observations:
1. **NGO Nomenclature:** The court advised NGOs to avoid using terms like "State" or "Central" in their names to prevent public confusion about their statutory status.
2. **Regulatory Compliance:** The court directed the Inspector General of Registration to ensure NGOs do not mislead the public regarding their statutory nature.
Decision:
The court disposed of the writ petition, noting the pending application for regularization. It emphasized the need for regulatory bodies to ensure compliance with municipal rules and directed the Inspector General of Registration to address NGO nomenclature issues.
This judgment highlights the importance of adhering to municipal regulations and ensuring clarity in the representation of non-governmental organizations.

Table 8: Sample reference and generated summaries from the ILC test set (Document title = 'The words Central and State should not be used with the names of the NGOs: Kerala High Court', Model = Phi-4-14B.).