

Legal Case Document Summarization: Extractive and Abstractive Methods and their Evaluation

Abhay Shukla¹ Paheli Bhattacharya¹ Soham Poddar¹ Rajdeep Mukherjee¹
Kripabandhu Ghosh² Pawan Goyal¹ Saptarshi Ghosh^{1*}

¹Indian Institute of Technology Kharagpur, India

²Indian Institute of Science Education and Research Kolkata, India

Abstract

Summarization of legal case judgement documents is a challenging problem in Legal NLP. However, not much analyses exist on how different families of summarization models (e.g., extractive vs. abstractive) perform when applied to legal case documents. This question is particularly important since many recent transformer-based abstractive summarization models have restrictions on the number of input tokens, and legal documents are known to be very long. Also, it is an open question on how best to evaluate legal case document summarization systems. In this paper, we carry out extensive experiments with several extractive and abstractive summarization methods (both supervised and unsupervised) over three legal summarization datasets that we have developed. Our analyses, that includes evaluation by law practitioners, lead to several interesting insights on legal summarization in specific and long document summarization in general.

1 Introduction

In Common Law systems (followed in India, UK, USA, etc.) law practitioners have to read through hundreds of case judgements/rulings in order to identify relevant cases that they can cite as precedents in an ongoing case. This is a time-consuming process as case documents are generally very long and complex. Thus, automatic summarization of legal case documents is an important problem (Gelbart and Smith, 1991; Bhattacharya et al., 2019; Zhong et al., 2019; Liu and Chen, 2019). It is additionally challenging due to two primary reasons as demonstrated in Table 1 – (i) legal documents as well as their summaries are much longer than most other types of documents, and (ii) since it is expensive to get *Law Experts* to write summaries, the datasets are usually much smaller, making it difficult to use supervised models.

*Corresponding author: saptarshi@cse.iitkgp.ac.in

Dataset	Language	Domain	#Doc	Avg # Tokens	
				Doc	Summ
CNN/DM (Hermann et al., 2015)	EN	News	312K	781	56
Gigawords (Napoles et al., 2012)	EN	News	4.02M	31	8
arXiv (Cohan et al.)	EN	Academic	216K	6,914	293
PubMed (Cohan et al.)	EN	Academic	133K	3,224	214
TL,DR, TOS,DR (Manor and Li, 2019)	EN	Contracts	506	106	17
BigPatent (Sharma et al.)	EN	Patent	1.34M	3,573	117
RulingBR (Feijó and Moreira, 2018)	Portuguese	Court Rulings	10,623	1,397	100
<i>This work</i>					
IN-Ext (Indian docs, extractive summ)	EN	Court Rulings	50	5,389	1,670
IN-Abs (Indian docs, abstractive summ)	EN	Court Rulings	7,130	4,378	1,051
UK-Abs (UK docs, abstractive summ)	EN	Court Rulings	793	14,296	1,573

Table 1: Comparing some existing summarization datasets with the three legal summarization datasets developed in this work. Last two columns give the average number of tokens per document and per summary.

A plethora of solutions exists for text summarization, for e.g., extractive and abstractive, supervised and unsupervised, etc. (Huang et al., 2020a). Also, several legal domain-specific methods have been designed for case document summarization (Zhong et al., 2019; Liu and Chen, 2019). However, detailed systematic analyses are rare on how the different families of summarization models perform on legal case documents. Our prior work (Bhattacharya et al., 2019) took an early step in this direction, but it mostly considered extractive methods. The state-of-the-art in document summarization has advanced rapidly in the last couple of years, and there has not been much exploration on how recent transformer-based summarization models perform on legal documents (Feijo and Moreira, 2021; Bajaj et al., 2021).

To bridge this gap, we (1) develop three legal case judgement summarization datasets from case documents from the Indian and UK Supreme Courts (see Table 1; details in Section 3), and (2) reproduce/apply representative methods from several families of summarization models on these datasets, and analyse their performances. To our knowledge, this is the first study on how a wide spectrum of summarization methods perform over legal case documents. We list below some interesting insights that come out from our analyses.

- **Domain-specific vs Domain-agnostic methods:** We apply several domain-independent sum-

marization methods, including **unsupervised extractive** (e.g., LexRank (Erkan and Radev, 2004), DSDR (He et al., 2012), and PacSum (Zheng and Lapata, 2019)), **supervised extractive** (e.g., SummaRunner (Nallapati et al., 2017), and BERTSUMM (Liu and Lapata, 2019)), and **supervised abstractive** (e.g., BART (Lewis et al., 2020), and Longformer (Beltagy et al., 2020)) on legal case documents. We then reproduce several legal domain-specific summarization methods, for e.g., MMR (Zhong et al., 2019), CaseSummarizer (Polsley et al., 2016) (unsupervised) and Gist (Liu and Chen, 2019) (supervised). In many cases, we observe general (domain-agnostic) methods to perform better than domain-specific methods.

- **Domain-specific training/fine-tuning:** Using models pretrained on legal corpora, like LegalPegasus (*leg*), consistently improves performance. We also explore and compare multiple ways of generating legal data for training supervised models and further fine-tuning pretrained models.

- **How to deal with long documents:** A key challenge in using existing abstractive summarizers on legal documents is that the input capacity of such models is often much lower than the length of legal documents. Accordingly, we experiment with three different approaches for summarizing long legal case documents – (i) applying *long document summarizers* such as Longformer (Beltagy et al., 2020) that are designed to handle long documents, (ii) applying *short document summarizers* such as BART (Lewis et al., 2020) and LegalPegasus (*leg*) together with approaches for *chunking* the documents, and (iii) reducing the size of the input document by first performing an extractive summarization and then going for abstractive summarization. In general, we find the chunking-based approach to perform better for legal documents, especially with fine-tuning, although Longformer performs the best on the UK-Abs dataset containing the longest documents, according to some of the metrics.

- **Evaluation of summary quality:** As noted in (Bhattacharya et al., 2019), *Law Experts* advise to not only evaluate the full-document summaries, but also check how well a summary is able to represent the different logical rhetorical segments in a legal case document (such as Facts, Final Judgment, etc. – see Appendix, Section A.1). To this end, we perform (i) document-wide automatic evaluations, (ii) segment-wise automatic evaluations, as well as (iii) evaluations by Law practitioners (the

actual end-users of legal summarization systems).

We show that simply computing document-wide metrics gives an incomplete picture of the quality of legal document summarization. In particular, we see some differences between automatic evaluation and evaluation by domain experts. For instance, supervised methods like SummaRunner, and finetuned BART usually achieve higher ROUGE scores, but the law practitioners often prefer the summaries generated by simpler unsupervised methods such as DSDR and CaseSummarizer. Again, the ROUGE scores achieved by the best extractive models are at par with those achieved by the best abstractive models. However, the practitioners often prefer the extractive summaries over the abstractive ones.

Availability of resources: The three legal summarization datasets curated in this work and the implementations of various summarization models are publicly available at <https://github.com/Law-AI/summarization>.

2 Related Work

We give an overview of existing summarization algorithms (Dong, 2018; Huang et al., 2020a).

Extractive domain-independent methods: There exists a wide range of general/domain-agnostic *unsupervised* summarizers such as Reduction (Jing, 2000), and the graph-based LexRank algorithm (Erkan and Radev, 2004). LSA (Gong and Liu, 2001) is a matrix-factorization based method and DSDR (He et al., 2012) relies on data reconstruction. PacSum (Zheng and Lapata, 2019) is a recent BERT-based method. Among *supervised* neural summarizers, SummaRuNNer (Nallapati et al., 2017) and BERTSum (Liu and Lapata, 2019) treat document summarization as a binary classification problem (in-summary vs. out-of-summary).

Extractive domain-specific methods: Several domain-specific approaches have been specifically designed for summarizing legal case documents. Among unsupervised methods, (1) LetSum (Farzindar and Lapalme, 2004) and (2) KMM (Saravanan et al., 2006) rank sentences based on term distribution models (TF-IDF and k-mixture model respectively); (3) CaseSummarizer (Polsley et al., 2016) ranks sentences based on their TF-IDF weights coupled with legal-specific features; (4) MMR (Zhong et al., 2019) generates a template-based summary using a 2-stage classifier and a Maximum Margin

Relevance (Zhong et al., 2019) module.

To our knowledge, Gist (Liu and Chen, 2019) is the only supervised method specifically designed for summarizing legal case documents. Gist first represents a sentence with different handcrafted features. It then uses 3 models – MLP, Gradient Boosted Decision Tree, and LSTM – to rank sentences in order of their likelihood to be included in the summary. We reproduce all these methods (implementation details in Appendix, Section A.2).

Abstractive methods: Most abstractive summarization models have an input token limit which is usually shorter than the length of legal case documents. Approaches from this family include Pointer-Generator (See et al., 2017), BERTSumAbs (Liu and Lapata, 2019), Pegasus (Zhang et al., 2020) and BART (Lewis et al., 2019) (input token limits for these models are at most 1024). Models like Longformer (Beltagy et al., 2020) introduce transformer architectures with more efficient attention mechanisms that enables them to summarize long documents (up to 16×1024 input tokens).

Bajaj et al. (2021) developed a two-step extractive-abstractive approach for long document summarization – they use a pre-trained BART model over compressed documents generated by identifying salient sentences. In this work, we reproduce a simplified version of this method.

Gidiotis and Tsoumakas (2020) presented a divide and conquer approach for long document summarization; they split the documents and summaries, using sentence similarity, into an ensemble of smaller summarization problems. In this work, we apply a method inspired by Gidiotis and Tsoumakas (2020) to fine-tune abstractive models.

To our knowledge, the only method for abstractive *legal* document summarization is LegalSumm (Feijo and Moreira, 2021). The method uses the RulingBR dataset (in Portuguese language) which has much shorter documents and summaries than the datasets in this work (see Table 1). A limitation of LegalSumm is that it can generate summaries only up to 200 tokens (which is much smaller than our target summaries); hence we do not apply this method in this work.

3 Datasets for Legal Summarization

There are very few publicly available datasets for legal case document summarization, especially in English (see Table 1). In this work, we develop the following three datasets:

(i) Indian-Abstractive dataset (IN-Abs): We collect Indian Supreme Court judgements from the website of Legal Information Institute of India (<http://www.liiofindia.org/in/cases/cen/INSC/>) which provides free and non-profit access to databases of Indian law. Abstractive summaries (also called “headnotes”) are available for some of these cases; of which we include 7,130 case documents, together with their *headnotes/summaries* as part of the dataset. We reserve 100 randomly-selected document-summary pairs for evaluation and the remaining 7,030 pairs are used for training the supervised models.

(ii) Indian-Extractive dataset (IN-Ext): Different law practitioners may have different preferences about the summary of a legal case document. Per discussion with *Law Experts* (two recent LLB graduates and a Professor from the Rajiv Gandhi School of Intellectual Property Law, a reputed Law school in India), we understand that they are *not* much satisfied with the summaries in the IN-Abs dataset. According to these experts, legal case documents have various *rhetorical segments*, and the summary should contain a representation from each segment. Based on the above preference, the two LLB graduates first rhetorically labelled each sentence from 50 case documents from the Indian Supreme Court (total 9,380 sentences), with one of the following labels – *Facts* (abbreviated as FAC), *Argument* (ARG), *Statute* (STA), *Precedent* (PRE), *Ratio of the decision* (Ratio), and *Ruling by Present Court* (RPC). Descriptions of these rhetorical labels are given in the Appendix (Section A.1). Then they wrote *extractive* summaries for the same 50 documents, each of length approximately one-third of that of the documents. They summarized each rhetorical segment separately; however, they preferred to summarize the segments ‘Ratio’ and ‘Precedent’ together. Each LLB graduate was paid a (mutually agreed) honorarium of INR 800 for labeling and summarizing each document.

Since 50 document-summary pairs are not sufficient for training supervised models, when applying these models on IN-Ext, they were trained over the 7,030 document-summary pairs in the IN-Abs train set. We ensure that there is *no overlap* between this training set and the IN-Ext dataset.

(iii) UK-Abstractive dataset (UK-Abs): The UK Supreme court website (<https://www.supremecourt.uk/decided-cases/>) provides all cases judgements that were ruled since

Dataset	Type of Summary	Compression Ratio	Test Set Size	Training Set Size
IN-Ext	Ext, segmented	0.31	50	7030
IN-Abs	Abs, non-segmented	0.24	100	
UK-Abs	Abs, segmented	0.11	100	693

Table 2: The three datasets developed in this work.

the year 2009. For most of the cases, along with the judgements, they also provide the official *press summaries* of the cases, which we consider as the reference summary. The summaries are abstractive in nature and are divided into three segments – ‘Background to the Appeal’, ‘Judgement’, and ‘Reasons for Judgement’. We gathered a set of 793 case documents (decided during the years 2009–2021) and their summaries. We reserve 100 document-summary pairs for evaluation and use the remaining 693 document-summary pairs for training the supervised models.

Table 2 provides a summary of the datasets, while Table 1 compares the length of the documents in these datasets with those in other datasets. Note that the documents in UK-Abs are approximately double the length of the IN-Abs and IN-Ext documents, and have a very low compression ratio (0.11); hence the UK-Abs dataset is the most challenging one for automatic summarization.

4 Experimental Setup and Evaluation

Target length of summaries: During inference, the trained summarization models need to be provided with the target length of summaries L (in *number of words*). For every document in the IN-Ext dataset, we have two reference summaries (written by two experts). For a particular document, we consider L to be the average of the number of words in the two reference summaries for that document. For IN-Abs and UK-Abs datasets, L is taken as the number of words in the single abstractive reference summary for a given document.

Given a document, every model is made to generate a summary of length at most L words. Some algorithms (e.g. KMM, Gist) return a ranking of sentences according to their summary-worthiness. The final summary is obtained by selecting sentences in descending order of the ranked list till the limit of L words is reached.

Evaluation of summary quality: We report ROUGE-1, ROUGE-2, and ROUGE-L F-scores (computed using <https://pypi.org/project/py-rouge/>, with *max_n* set to 2, parameters *limit_length* and *length_limit* not

used, and other parameters kept as default), and BertScore (Zhang et al., 2019) (computed using <https://pypi.org/project/bert-score/> version 0.3.4) that calculates the semantic similarity scores using the pretrained BERT model. We calculate two kinds of ROUGE and BERTScore as follows:

(a) *Overall document-wide scores:* For a given document, we compute the ROUGE and BERTScore of an algorithmic summary with respect to the reference summary. For IN-Ext, we compute the scores individually with each of the two reference summaries and take the average. The scores are averaged over all documents in the evaluation set.

(b) *Segment-wise scores:* In legal case judgement summarization, a segment-wise evaluation is important to understand how well each rhetorical segment has been summarized (Bhattacharya et al., 2019). We can perform this evaluation only for the IN-Ext and UK-Abs datasets (and *not* for IN-Abs), where the reference summaries are written segment-wise. For each rhetorical segment (e.g., *Fact* or *Background*), we extract the portion of the gold standard summary that belongs to that segment. Then we compute the ROUGE score between the entire algorithmic summary and segment-specific part of the reference summary. We compute the average ROUGE score for a particular segment, averaged over all documents in the evaluation set.¹

In the segment-wise evaluation, we only report ROUGE Recall scores, and *not* F-scores. This is because the summarization algorithms output only a coherent set of sentences as summary, and do *not* specify which part of the summary belongs to which segment; computing ROUGE Precision or F-Score in this case would be misleading.

Expert evaluation: We select a few methods (that achieve the highest ROUGE scores) and get the summaries generated by them for a few documents evaluated by three Law experts (Section 7.3).

Consistency scores: It is important to measure the consistency of an algorithmic summary with the original document, given the possibility of hallucination by abstractive models (Pagnoni et al., 2021). To this end, we experimented with the SummaC_{CONV} summary consistency checker (Laban et al., 2022). However, we find that it gives very

¹In this paper, we report segment-wise ROUGE scores only since both segment-wise ROUGE scores as well as segment-wise BERTScores give similar insights.

low consistency scores to the expert-written reference abstractive summaries – the average scores for the expert summaries in IN-Abs and UK-Abs are 0.485 and 0.367 respectively. A probable reason for these counter-intuitive scores could be that the SummaC_{CONV} model could not be fine-tuned on a legal domain-specific dataset, owing to its unavailability. Curating such a dataset to check for factual consistency of summaries of legal documents, together with developing a suitable consistency measure for summaries in the legal domain are envisioned as immediate future works. The present SummaC_{CONV} consistency scores are therefore concluded to be unreliable for legal document summarization, and hence are not reported.

5 Extractive Summarization Methods

We consider some representative methods from four classes of extractive summarizers: (1) Legal domain-specific unsupervised methods: LetSum, KMM, CaseSummarizer, and MMR. (2) Legal domain-specific supervised methods: Gist. (3) Domain-independent unsupervised methods: LexRank, LSA, DSDR, Luhn, Reduction and PacSum. (4) Domain-independent supervised methods: SummaRuNNer and BERTSum.

Short descriptions of all the above methods are given in Section 2. The implementation details for the domain-specific methods we implemented, and publicly available code repositories are stated in the Appendix (Section A.2 and Section A.3).

Training supervised extractive models: The supervised methods (Gist, SummaRuNNer and BERTSUM) require labelled training data, where every sentence must be labeled as 1 if the sentence is suitable for inclusion in the summary, and 0 otherwise. As stated in Section 3, we use parts of the IN-Abs and UK-Abs datasets for training the supervised methods. However, since both these datasets have *abstractive* summaries, they cannot be directly used to train the extractive summarizers.

We explore three methods – *Maximal*, *Avr*, and *TF-IDF* – for converting the abstractive summaries to their extractive counterparts. Best performances for the supervised methods are observed when the training data is generated through the *Avr* method; hence we describe *Avr* here and report results of the supervised methods trained on data generated through *Avr*. Descriptions of *Maximal* and *TF-IDF* are stated in the Appendix (Section A.4).

Avr: We adopt the technique given by Narayan et al. (2018). For each sentence in the abstractive gold-standard summary, we select 3 sentences from the source document (full text) that have the maximum average of ROUGE-1, ROUGE-2 and ROUGE-L scores w.r.t. the sentence in the abstractive summary. Then we take the union of all the sentences thus selected, and label them 1 (to be included in the summary). All other sentences in the source document are assigned a label of 0.

6 Abstractive Summarization Methods

We apply several abstractive methods for legal document summarization, including both pretrained models and models finetuned for legal document summarization. A key challenge in applying such methods is that legal documents are usually very long, and most abstractive summarization models have restrictions on the number of input tokens.

6.1 Pretrained Abstractive Models

6.1.1 Models meant for short documents

We consider Legal-Pegasus (*leg*) which is already pretrained on legal documents, and BART (Lewis et al., 2020) (max input length of 1024 tokens). We use their pre-trained versions from the HuggingFace library; details in the Appendix (Section A.5).

The input token limit in these models (1024) is much smaller than the number of words in a typical legal case document. Hence, to apply these models on legal case documents, we apply a chunking-based approach as described below:

Chunking-based approach: We first divide a document into small chunks, the size of each chunk being the maximum number of tokens (say, n) that a model is designed/pre-trained to accept without truncating (e.g., $n = 1024$ for BART). Specifically, the first n tokens (without breaking sentences) go to the first chunk, the next n tokens go to the second chunk, and so on. Then we use a model to summarize every chunk. For a given document, we equally divide the target summary length among all the chunks. Finally, we append the generated summaries for each chunk in sequence.

6.1.2 Models meant for long documents

Models like Longformer (LED) (Beltagy et al., 2020) have been especially designed to handle long documents (input capacity = 16,384 tokens), by including an attention mechanism that scales linearly

with sequence length. We use Legal-LED specifically finetuned on legal data (details in Appendix, Section A.5). The model could accommodate most case documents fully. A few documents in UK-Abs are however longer (see Table 2), those documents were truncated after 16,384 tokens.

6.1.3 Hybrid extractive-abstractive approach

To focus only on important parts of the document in the chunking-based approach, we use a hybrid of an extractive approach and an abstractive approach, similar to Bajaj et al. (2021). First, the document length is reduced by selecting salient sentences using a BERT-based extractive summarization model. Then a BART model is used to generate the final summary (Bajaj et al., 2021). Since, in our case, we often require a summary length greater than 1024 (see Table 1), we use a chunking-based BART (rather than pre-trained BART) in the second step. We call this model **BERT_BART**.

6.2 Finetuning Abstractive Models

Fine-Tuning transformer models has shown significant improvement in most downstream tasks. Hence, we finetune BART, Longformer, and Legal-Pegasus on our proposed datasets. We also use finetuned BART as part of our BERT_BART model.

Generating finetuning data: Finetuning supervised models needs a large set of doc-summary pairs. However, our considered models (apart from Longformer) have a restricted input limit which is lesser than the length of documents in our datasets. Hence, we use the following method, inspired from Gidiotis and Tsoumakas (2020), to generate finetuning data for chunking based summarization.

Consider (d, s) to be a (training document, reference summary) pair. When d is segmented into n chunks d_1, d_2, \dots, d_n , it is not logical for the same s to be the reference summary for each chunk d_i . In order to generate a suitable reference summary s_i for each chunk d_i , first we map every sentence in s to the most similar sentence in d . Here, we use a variety of sentence-similarity measures, as detailed below. Then for every chunk d_i , we combine all sentences in s which are mapped to any of the sentences in d_i , and consider those sentences as the summary s_i (of d_i). Following this procedure, from each document, we get a large number of (d_i, s_i) pairs which are then used for finetuning.

Sentence similarity measures for generating finetuning data: We experiment with several tech-

niques for measuring sentence similarity between two sentences – (i) Mean Cosine Similarity (**MCS**), (ii) Smooth Inverse Frequency (**SIF**), (iii) Cosine similarity between BERT [CLS] token embeddings (**CLS**), and (iv) **MCS_RR** which incorporates rhetorical role information. Out of these, we find MCS to perform the best. Hence we describe MCS in detail here. Descriptions of the other methods can be found in the Appendix (Section A.6).

In Mean Cosine Similarity (**MCS**) (Ranasinghe et al., 2019), we calculate the mean of token-level embeddings (obtained using SBERT (Reimers and Gurevych, 2019)) to obtain the representation for a given sentence. We then compute the cosine similarity between two such sentence embeddings.

We used all the methods stated above to generate fine-tuning datasets for IN-Abs and UK-Abs. We finetune three different versions of the BART model, BART_CLS, BART_MCS, and BART_SIF, using the three sentence similarity measures described above. Out of these, BART_MCS performs the best (as we will see in Section 7). Therefore, we use MCS for generating finetuning data for the other models, to obtain Legal-Pegasus-MCS and BART_MCS_RR (where the finetuning data is generated based on rhetorical labels). We also use the finetuned BART_MCS model with BERT_BART method to get BERT_BART_MCS.

The hyper-parameters used to finetune the different abstractive models are stated in Table 9 in the Appendix (Section A.5).

7 Results and Analyses

This section analyzes the performance of different summarization models. For IN-Ext, In-Abs and UK-Abs datasets, Table 3, Table 4 and Table 5 report the overall evaluation of a few of the best-performing methods, respectively. Table 6 and Table 7 show the segment-wise evaluation of a few best-performing methods on the IN-Ext and UK-Abs datasets respectively. Detailed results are given in Tables 10–14 in the Appendix (Section A.7).

7.1 Evaluation of Extractive methods

Overall Evaluation (Tables 3–5): Among the unsupervised general methods, Luhn (on IN-Ext) and DSDR (on IN-Abs and UK-Abs) show the best performances. Among the unsupervised legal-specific methods, CaseSummarizer performs the best on both In-Abs and UK-Abs datasets, while LetSum performs the best on IN-Ext. Among supervised

Algorithm	ROUGE Scores			BERTScore
	R-1	R-2	R-L	
<i>Extractive Methods</i>				
Unsupervised, Domain Independent				
Luhn	0.568	0.373	0.422	0.882
Pacsum_bert	0.59	0.41	0.335	0.879
Unsupervised, Legal Domain Specific				
MMR	0.563	0.318	0.262	0.833
KMM	0.532	0.302	0.28	0.836
LetSum	0.591	0.401	0.391	0.875
Supervised, Domain Independent				
SummaRunner	0.532	0.334	0.269	0.829
BERT-Ext	0.589	0.398	0.292	0.85
Supervised, Legal Domain Specific				
Gist	0.555	0.335	0.391	0.864
<i>Abstractive Methods</i>				
Pretrained				
BART	0.475	0.221	0.271	0.833
BERT-BART	0.488	0.236	0.279	0.836
Legal-Pegasus	0.465	0.211	0.279	0.842
Legal-LED	0.175	0.036	0.12	0.799
Finetuned				
BART_MCS	0.557	0.322	0.404	0.868
BART_MCS_RR	0.574	0.345	0.402	0.864
BERT_BART_MCS	0.553	0.316	0.403	0.869
Legal-Pegasus_MCS	0.575	0.351	0.419	0.864
Legal-LED	0.471	0.26	0.341	0.863

Table 3: Document-wide ROUGE-L and BERTScores (Fscore) on the IN-Ext dataset. All values averaged over the 50 documents in the dataset. The best value in a particular class of methods is highlighted in **bold**.

extractive methods, SummaRuNNer performs the best across both domain-independent and domain-specific categories, on the IN-Abs and UK-Abs datasets. BERT-Ext is the best performing model on the IN-Ext dataset.

Segment-wise Evaluation: Table 6 and Table 7 show the segment-wise ROUGE-L Recall scores of some of the best performing methods on the IN-Ext and UK-Abs datasets respectively. Section 4 details the process of obtaining these scores. According to overall ROUGE scores, it may seem that a particular method performs very well (e.g., LetSum on In-Ext), but that method *may not perform the best across all the segments* (e.g. among the extractive methods, LetSum performs the best in only 1 out of the 5 segments in In-Ext). This observation shows the importance of segment-wise evaluation. It is an open challenge to develop an algorithm that shows a balanced segment-wise performance. Some more interesting observations on segment-wise evaluations are given in the Appendix (Section A.8).

7.2 Evaluation of Abstractive methods

Overall Evaluation (Tables 3–5): Among the pre-trained models, Legal-Pegasus generates the best

Algorithm	ROUGE Scores			BERTScore
	R-1	R-2	R-L	
<i>Extractive Methods (U: Unsupervised, S: Supervised)</i>				
DSDR (U)	0.485	0.222	0.270	0.848
CaseSummarizer (U)	0.454	0.229	0.279	0.843
SummaRunner (S)	0.493	0.255	0.274	0.849
Gist (S)	0.471	0.238	0.308	0.842
<i>Finetuned Abstractive Methods</i>				
BART_MCS	0.495	0.249	0.330	0.851
BERT_BART_MCS	0.487	0.243	0.329	0.853
Legal-Pegasus_MCS	0.488	0.252	0.341	0.851
Legal-LED	0.471	0.235	0.332	0.856

Table 4: Document-wide ROUGE-L and BERTScores (Fscore) on the IN-Abs dataset, averaged over the 100 test documents. Results of some of the top-performing methods are shown here (all results in Table 11).

Algorithm	ROUGE Scores			BERTScore
	R-1	R-2	R-L	
<i>Extractive Methods (U: Unsupervised, S: Supervised)</i>				
DSDR (U)	0.484	0.174	0.221	0.832
CaseSummarizer (U)	0.445	0.166	0.227	0.835
SummaRunner (S)	0.502	0.205	0.237	0.846
Gist	0.427	0.132	0.215	0.819
<i>Finetuned Abstractive Methods</i>				
BART_MCS	0.496	0.188	0.271	0.848
BERT_BART_MCS	0.476	0.172	0.259	0.847
Legal-Pegasus_MCS	0.476	0.171	0.261	0.838
Legal-LED	0.482	0.186	0.264	0.851

Table 5: Document-wide ROUGE-L and BERTScores (Fscore) on UK-Abs dataset, averaged over the 100 test documents. Results of some of the top-performing methods are shown here (all results in Table 12).

summaries (Table 3), followed by BART-based methods. This is expected, since Legal-Pegasus is pre-trained on legal documents. This short document summarizer, when used with chunking to handle long documents, notably outperforms Legal-LED, which is meant for long documents. For IN-Ext dataset, BERT_BART performs the best maybe due to extractive nature of the summaries.

All models show notable improvement through fine-tuning. Overall, the best performances are noted by Legal-Pegasus (IN-Ext and IN-Abs) and BART_MCS (UK-Abs).

Segment-wise Evaluation (Tables 6, 7): Again, none of the methods performs well across all segments, and fine-tuning generally improves performance. Interestingly, though Legal-LED performs poorly with respect to document-wide ROUGE scores, it shows better performance in segment-wise evaluation – it gives the best performance in the FAC and ARG segments of IN-Ext and in 2 out of the 3 segments of UK-Abs. Since the UK-Abs dataset contains the longest documents, possibly

Algorithms	Rouge L Recall				
	RPC (6.42%)	FAC (34.85%)	STA (13.42%)	Ratio+Pre (28.83%)	ARG (16.45%)
<i>Extractive Methods (U: Unsupervised, S: Supervised)</i>					
LexRank (U)	0.039	0.204	0.104	0.208	0.127
Luhn (U)	0.037	0.272	0.097	0.175	0.117
LetSum (U)	0.036	0.237	0.115	0.189	0.1
SummaRunner (S)	0.059	0.158	0.08	0.209	0.096
Gist (S)	0.041	0.191	0.102	0.223	0.093
<i>Finetuned Abstractive Methods</i>					
BART_MCS_RR	0.061	0.192	0.082	0.237	0.086
Legal-Pegasus_MCS	0.037	0.192	0.09	0.257	0.101
Legal-LED	0.053	0.245	0.086	0.187	0.124

Table 6: Segment-wise ROUGE-L Recall scores of the best methods in Table 3 on the IN-Ext dataset. All values are averaged over the 50 documents in the dataset. The best scores for each segment in a particular class of methods are in **bold**. Results of all methods in Table 13.

Algorithms	Rouge-L Recall		
	Background (39%)	Final Judgement (5%)	Reasons (56%)
<i>Extractive Methods (U: Unsupervised, S: Supervised)</i>			
SummaRunner (S)	0.172	0.044	0.165
BERT-Ext (S)	0.203	0.034	0.135
Gist (S)	0.123	0.041	0.195
<i>Finetuned Abstractive Methods</i>			
Legal-Pegasus_MCS	0.166	0.039	0.202
Legal-LED	0.187	0.058	0.172

Table 7: Segment-wise ROUGE-L Recall scores of the best methods in Table 5 on the UK-Abs dataset. All values averaged over the 100 documents in the evaluation set. Best scores for each segment in a particular class of methods are in **bold**. Results of all methods in Table 14.

Legal-LED has an advantage over chunking-based methods when evaluated segment-wise.

Overall performance on long legal case documents: We experimented with three approaches for summarizing long documents – (i) models with modified attention mechanism such as Legal-LED, (ii) methods based on chunking the documents, and (iii) reducing the size of the input by initial extractive summarization and then going for abstractive summarization (BERT_BART). When we see the overall (document-wide) ROUGE scores, **Legal-Pegasus** and **BART** (when used along with chunking), are seen to perform the best, followed by BERT_BART. However for segment-wise performances **Legal-LED** shows greater potential.

7.3 Expert evaluation

Finally, we evaluate some of the model-generated summaries via three domain experts. Since it is expensive to obtain evaluations from Law experts, we chose to conduct this evaluation for a few documents/summaries from the IN-Abs dataset.

Recruiting the 3 experts: We recruited the two recent LLB graduates (who wrote the reference sum-

maries in IN-Ext) from the Rajiv Gandhi School of Intellectual Property Law (RGSOIPL), India, who were mentored by a Professor of the same Law school (as mentioned in Section 3) while carrying out the annotations. Additionally, we recruited a senior Faculty of Law from the West Bengal National University of Juridical Sciences (WBNUJS), India. Note that both RGSOIPL and WBNUJS are among the most reputed Law schools in India.

Each annotator was paid a (mutually agreed) honorarium of INR 200 for evaluation of each summary. The annotators were clearly informed of the purpose of the survey. Also we discussed their experiences after the survey about. Through all these steps, we tried our best to ensure that the annotations were done rigorously.

Survey setup: We select the summaries generated by 7 algorithms which give relatively high ROUGE-L F-Score on IN-Abs – see Table 8. Then, we show the annotators 5 selected documents and their summaries generated by the 7 algorithms (35 summaries evaluated in total). An annotator was asked to evaluate a summary on the basis of the following parameters – (1) how well a summary represents each rhetorical segment, i.e., the final judgement (**RPC**), facts (**FAC**), relevant statutes/laws cited (**STA**), relevant precedents cited (**PRE**), the reasoning/rationale behind the judgement (**Ratio**), and the arguments presented in the case (**ARG**). (2) how well important information has been covered in the summary (**Imp Inf**). (3) Readability and grammatical coherence (**Read**). (4) An overall score for the summary (**Overall**).

Each summary was rated on a Likert scale of 0 – 5 on each parameter, independently by the 3 annotators. Thus, a particular method got 15 scores for each parameter – for 5 documents and by 3 annotators. Table 8 reports (i) the mean/average, and (ii) the median of all these 15 scores for each method and for each parameter.

Inter-Annotator Agreement: We calculate pairwise Pearson Correlation between the ‘Overall’ scores given by the three annotators over the 35 summaries, and then take the average correlation value as the IAA. Refer to the Appendix (Section A.9) for why we chose this IAA measure. The average IAA is 0.525 which shows moderate agreement between the annotators².

²<https://www.andrews.edu/~calkins/math/edrm611/edrm05.htm>

Algorithms	RPC		FAC		STA		PRE		Ratio		ARG		Imp.Inf.		Read.		Overall	
	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.
DSDR	4.2	5	3.8	4	3.7	4	3.1	3.7	3.7	4	1.9	3.7	3.7	4	4.3	4	3.9	4
CaseSummarizer	2.1	2	3.8	4	3.6	4	3	3.6	3.5	3	2.4	3	3.2	3	4.3	4	3.6	3
SummaRuNNer	2.1	3	4.2	4	2.4	3	3.3	3	2.9	3	2.1	2.9	3.2	3	4.1	4	3.2	4
Gist	3.3	4	1.8	3	2.6	3	3.5	3	3.2	4	2.1	3.2	3	3	3.9	4	3.2	3
Legal-Pegasus	1.4	1	3.9	4	3.2	4	2.4	3.2	2.9	3	2	2.9	3	3	3.5	4	3	3
BART-MCS	0.9	1	2.8	3	2.9	3	3.3	3	2.5	3	1.8	2.5	2.8	3	2.7	3	2.8	3
BART-MCS-RR	0.8	1	2.7	3	3.1	3	2.6	3	2.6	3	1.3	2.6	2.6	3	2.9	3	2.6	3

Table 8: Evaluation of some summaries from the IN-Abs dataset, by three domain experts (two recent LLB graduates and a Senior faculty of Law). The evaluation parameters are explained in the text. Scores are given by each expert in the range [0-5], 5 being the best. The Mean and Median (Med.) scores for each summarization algorithm and for each parameter are computed over 15 scores (across 5 documents; each judged by 3 experts).

Results (Table 8): According to the Law experts, important information (Imp. Inf.) could be covered best by DSDR, followed by CaseSummarizer and SummaRuNNer. In terms of readability (Read.) as well, DSDR, CaseSummarizer and SummaRuNNer have higher mean scores than others. Finally, through the Overall ratings, we understand that DSDR is of higher satisfaction to the Law practitioners than the other algorithms, with CaseSummarizer coming second. These observations show a discrepancy with the automatic evaluation in Section 7 where supervised methods got better ROUGE scores than unsupervised ones.

Importantly, we again see that none of the summaries could achieve a balanced representation of all the rhetorical segments (RPC – Arg). For instance, DSDR (which gets the best overall scores) represents the final judgement (RPC) and statutes (STA) well, but misses important precedents (PRE) and arguments (ARG).

In general, the experts opined that the summaries generated by several algorithms are good in the initial parts, but their quality degrades gradually from the middle. Also, the experts felt the abstractive summaries to be less organized, often having incomplete sentences; they felt that the abstractive summaries have potential but need improvement.

Correlation between expert judgments and the automatic metrics: As stated above, there seems to be some discrepancy between expert judgements and the automatic metrics for summarization. To explore this issue further, we compute the correlation between the expert judgments (average of the ‘Overall’ scores of the three annotators) and the automatic metrics (ROUGE-1,2, L Fscores and BERT-Scores). The human evaluation was conducted over 5 documents and 7 algorithms. So, for each metric, correlation was calculated between the 5 human-assigned overall scores and the 5 metric scores, and then an average was taken across all the

7 algorithms (details in Appendix Section A.9).

Following this procedure, the correlation of the mean ‘Overall’ score (assigned by experts) with ROUGE-1 F-Score is 0.212, that with ROUGE-2 F-Score is 0.208, that with ROUGE-L F-Score is 0.132 and the correlation with BERTScore is 0.067. These low correlation scores again suggest that *automatic summarization metrics may be insufficient* to judge the quality of summaries in specialized domains such as Law.

8 Concluding discussion

We develop datasets and benchmark results for legal case judgement summarization. Our study provides several guidelines for long and legal document summarization: (1) For extractive summarization of legal documents, DSDR (unsupervised) and SummaRuNNer (supervised) are promising methods. (2) For abstractive summarization, Legal-Pegasus (pretrained and finetuned) is a good choice. (3) For long documents, fine-tuning models through chunking seems a promising way. (4) Document-wide evaluation does not give the complete picture; domain-specific evaluation methods, including domain experts, should also be used.

Acknowledgements

The authors acknowledge the anonymous reviewers for their suggestions. The authors thank the Law domain experts from the Rajiv Gandhi School of Intellectual Property Law, India (Amritha Shaji, Ankita Mohanty, and Prof. Uday Shankar) and from the West Bengal National University of Juridical Sciences, India (Prof. Shouvik Guha) who helped in developing the gold standard summaries (IN-Ext dataset) and evaluating the summaries. The research is partially supported by the TCG Centres for Research and Education in Science and Technology (CREST) through a project titled “Smart Legal Consultant: AI-based Legal Analytics”.

References

- Legal pegasus. <https://huggingface.co/nsi319/legal-pegasus>. [Online].
- Ahsaas Bajaj, Pavitra Dangati, Kalpesh Krishna, Pradhiksha Ashok Kumar, Rheeya Uppaal, Bradford Windsor, Eliot Brenner, Dominic Dotterrer, Rajarshi Das, and Andrew McCallum. 2021. [Long document summarization in a low resource setting using pre-trained language models](#). *CoRR*, abs/2103.00751.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *CoRR*, abs/2004.05150.
- Paheli Bhattacharya, Kaustubh Hiware, Subham Rajgaria, Nilay Pochhi, Kripabandhu Ghosh, and Saptarshi Ghosh. 2019. A comparative study of summarization algorithms applied to legal case judgments. In *Proc. European Conference on Information Retrieval*.
- Paheli Bhattacharya, Shounak Paul, Kripabandhu Ghosh, and Saptarshi Ghosh. 2021. [DeepRhole: deep learning for rhetorical role labeling of sentences in legal case documents](#). *Artificial Intelligence and Law*.
- Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. A discourse-aware attention model for abstractive summarization of long documents. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 615–621.
- Yue Dong. 2018. [A survey on neural network-based summarization methods](#). *CoRR*, abs/1804.04589.
- Günes Erkan and Dragomir R. Radev. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. *J. Artif. Int. Res.*, 22(1).
- Atefeh Farzindar and Guy Lapalme. 2004. Letsum, an automatic legal text summarizing system. *Proc. Legal knowledge and information systems (JURIX)*.
- Diego Feijo and Viviane P. Moreira. 2021. Improving abstractive summarization of legal rulings through textual entailment. *Artificial Intelligence and Law*.
- Diego Feijó and Viviane Moreira. 2018. Rulingbr: A summarization dataset for legal texts. In *Proc. International Conference on Computational Processing of the Portuguese Language*, pages 255–264.
- Dephne Gelbart and JC Smith. 1991. Beyond boolean search: Flexicon, a legal tex-based intelligent system. In *Proceedings of the 3rd international conference on Artificial intelligence and law*, pages 225–234. ACM.
- Alexios Gidiotis and Grigorios Tsoumakas. 2020. [A divide-and-conquer approach to the summarization of academic articles](#). *CoRR*, abs/2004.06190.
- Yihong Gong and Xin Liu. 2001. Generic text summarization using relevance measure and latent semantic analysis. In *Proc. International conference on Research and development in information retrieval (SIGIR)*.
- Zhanying He, Chun Chen, Jiajun Bu, Can Wang, Lijun Zhang, Deng Cai, and Xiaofei He. 2012. Document summarization based on data reconstruction. In *Proc. AAAI Conference on Artificial Intelligence*, pages 620–626.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Dandan Huang, Leyang Cui, Sen Yang, Guangsheng Bao, Kun Wang, Jun Xie, and Yue Zhang. 2020a. [What have we achieved on text summarization?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 446–469, Online. Association for Computational Linguistics.
- Dandan Huang, Leyang Cui, Sen Yang, Guangsheng Bao, Kun Wang, Jun Xie, and Yue Zhang. 2020b. What have we achieved on text summarization? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 446–469.
- Hongyan Jing. 2000. Sentence reduction for automatic text summarization. In *Proc. Applied Natural Language Processing Conference*.
- Chris Kedzie, Kathleen Mckeown, and Hal Daumé III. 2018. Content selection in deep learning models of summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1818–1828.
- Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. 2022. [SummaC: Re-visiting NLI-based models for inconsistency detection in summarization](#). *Transactions of the Association for Computational Linguistics*, 10:163–177.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2019. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). *CoRR*, abs/1910.13461.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proc. ACL*, pages 7871–7880.

- Chao-Lin Liu and Kuan-Chun Chen. 2019. Extracting the gist of chinese judgments of the supreme court. In *Proc. International Conference on Artificial Intelligence and Law (ICAIL)*.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. In *Proc. EMNLP-IJCNLP*.
- Laura Manor and Junyi Jessy Li. 2019. Plain English summarization of contracts. In *Proceedings of the Natural Legal Language Processing Workshop 2019*, pages 1–11.
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. Summarunner: A recurrent neural network based sequence model for extractive summarization of documents. In *Proc. AAAI Conference on Artificial Intelligence*.
- Courtney Napoles, Matthew Gormley, and Benjamin Van Durme. 2012. Annotated Gigaword. In *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction (AKBC-WEKEX)*, pages 95–100.
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Ranking sentences for extractive summarization with reinforcement learning. In *Proc. NAACL-HLT*, pages 1747–1759.
- Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov. 2021. Understanding Factuality in Abstractive Summarization with FRANK: A Benchmark for Factuality Metrics. In *Proc. NAACL-HLT*, pages 4812–4829.
- Seth Polsley, Pooja Jhunjhunwala, and Ruihong Huang. 2016. Casesummarizer: A system for automated summarization of legal texts. In *Proc. International conference on Computational Linguistics (COLING) System Demonstrations*.
- Tharindu Ranasinghe, Constantin Orasan, and Ruslan Mitkov. 2019. Enhancing unsupervised sentence similarity methods with deep contextualised word representations. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 994–1003, Varna, Bulgaria. INCOMA Ltd.
- 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. In *Legal knowledge and information systems, JURIX*.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083.
- Eva Sharma, Chen Li, and Lu Wang. BIGPATENT: A large-scale dataset for abstractive and coherent summarization. In *Proc. ACL*, pages 2204–2213.
- Sajad Sotudeh, Arman Cohan, and Nazli Goharian. 2021. On generating extended summaries of long documents. In *The AAAI-21 Workshop on Scientific Document Understanding (SDU 2021)*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J Liu. 2020. Pegasus: pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning*, pages 11328–11339.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Hao Zheng and Mirella Lapata. 2019. Sentence centrality revisited for unsupervised summarization. In *Proc. ACL*, pages 6236–6247.
- 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 *Proc. International Conference on Artificial Intelligence and Law (ICAIL)*.

A Appendix

A.1 Rhetorical Role Labels in a Legal Case Document

According to our legal experts, rhetorical role labels/segments define a semantic function of the sentences in a legal case documents. A good summary should contain a concise representation of each segment. These rhetorical segments are defined as follows:

- (i) *Facts (abbreviated as FAC)*: refers to the chronology of events that led to filing the case;
- (ii) *Argument (ARG)*: arguments of the contending parties;
- (iii) *Statute (STA)*: Established laws referred to by the present court;
- (iv) *Precedent (PRE)*: Precedents/prior cases that were referred to;
- (v) *Ratio of the decision (Ratio)*: reasoning/rationale for the final judgement given by the present court;
- (vi) *Ruling by Present Court (RPC)*: the final judgement given by the present court.

A.2 Implementations details of Domain-Specific Extractive Summarization Methods

We state here the reproducibility details of the legal domain-specific summarization methods, which could not be stated in the main paper due to lack of space.

- **Legal Dictionary**: Some domain-specific summarization methods like CaseSummarizer and Gist use a set of legal keywords for identifying importance of sentences in the input document. We identify these keywords using a glossary from the legal repository <https://www.advocatekhoj.com/library>. This website provides several legal resources for Indian legal documents, including a comprehensive glossary of legal terms.
- **MMR**: The original paper experiments on BVA decision of the US jurisdiction. The MMR method creates a template-based summary considering various semantic parts of a legal case document, and selecting a certain number of sentences from each semantic part. Specifically, the summary is assumed to contain (i) one sentence from the procedural history, (ii) one sentence from issue, (iii) one sentence from the service history of the veteran, (iv) a

variable number of Reasoning & Evidential Support sentences selected using Maximum Margin Relevance, (v) one sentence from the conclusion. Pattern-based regex extractors are used to identify the sentences (i)-(iii) and (v).

Reasoning & Evidential Support sentences are identified using a 2-step supervised classification method – in the first step, sentences predictive of a case’s outcome are detected using Convolutional Neural Networks. In the second step, a Random Forest Classifier is used to specifically extract the “Reasoning & Evidential Support” sentences from the predictive sentences. In the absence of such annotated training datasets to build a 2-stage classification framework for India and UK, we adopt only the Maximum Margin Relevance module of their work as a baseline.

This method decides the inclusion of a sentence S_i to the summary based on $\lambda \times Sim(S_i, Case) + (1 - \lambda) \times Sim(S_i, Summary)$, where $Case$ indicates the set of sentences in the original case document and $Summary$ represents the current set of sentences in the summary. λ acts as the weight that balances the relevance and diversity; we consider $\lambda = 0.5$.

- **Gist**: Gist uses the following handcrafted features to represent every sentence in the input case document (which is to be summarized) –
 - (i) Quantitative features: number of words, number of characters, number of unique words, and position of the sentence
 - (ii) Case category information: The original paper produced summaries of Chinese documents which contain information like whether a document is recorded as a judgment or as a ruling (which is a category of judicial judgments) and specific words that are used by the courts to indicate subcategories of the judgments. These information are absent in Indian and UK Supreme Court Case documents. So we do not consider this category of features.
 - (iii) Specific Legal Terms: We use a legal dictionary for the purpose (from <https://www.advocatekhoj.com/library/glossary/a.php>, as stated in the main paper).
 - (iv) Word Embeddings: To construct the embedding of a sentence, we take the average of the embeddings of the words in the sentence. To this end, we train a word2vec model on the training corpus (7030 documents of the IN-Abs and 693 documents of the UK-Abs dataset). During

evaluation, the trained word2vec model is used to derive the embeddings.

(v) One-hot vectors of first k POS tags in the sequence, where $k = 10$ as mentioned in the paper

(vi) Word Embeddings of the opening words: we take the average of the embeddings of the first 5 words in the sentence, since the paper did not clearly mention how to obtain them.

Based on the above features, Gist uses 3 models – MLP, Gradient Boosted Decision Tree, LSTM and a combination of LSTM and MLP classifiers – to rank sentences in order of their likelihood to be included in the summary. We observe the best performance by using Gradient Boosted Decision Tree as the ML classifier, which we report.

• **CaseSummarizer:** The original method of CaseSummarizer was developed for Australian documents. All sentences in the input document are ranked using the following score: $w_{new} = w_{old} + \sigma(0.2d + 0.3e + 1.5s)$, where w_{old} is the sum of the TF-IDF values of its constituent words, normalized over the sentence length, d is the number of ‘dates’ present in the sentence, e is the number of named entity mentions in the sentence, s is a boolean variable indicating whether the sentence is at the start of any section, and σ is the standard deviation among the sentence scores.

The Indian case documents used in our study (IN-Ext and IN-Abs) are less structured than Australian case documents, and they do not contain ‘section headings’. So, in place of that feature we used a count of the number of legal terms (identified by a legal dictionary) present in the sentence. We could find section numbers of Acts in our gold standard summaries, for example, “section 302 of the Indian Penal Code”. Hence, for the parameter “d” in the formulation, we included both dates and section numbers. The authors did not clearly mention how they have identified the “entities” in the texts. So, we have used the Stanford NER Tagger for identifying entities within the sentence. For ensuring a fair comparison, we have used the same setting on UK-Abs too.

• **LetSum and KMM:** Both the LetSum and KMM methods initially assign rhetorical labels to sentences (using certain cue-phrases and Conditional Random Fields respectively). The sentences are then ranked, for which LetSum uses TF-IDF scores and KMM uses a K-Mixture Model based score. However, the rhetorical role information is *not* used for generating the summary. Rather, the rhetorical

labels are used as a *post-summarization step* that is mainly used for displaying the summary in a structured way. We therefore implement only the sentence ranking modules for these methods – i.e., TF-IDF based summarization for LetSum and K-mixture model based summarization for KMM.

A.3 Implementation Details of Domain-Independent Extractive Summarization Methods

We use the publicly available implementations of the domain-independent extractive methods from the following sources:

- LexRank, LSA, Luhn and Reduction: <https://pypi.org/project/sumy/>
- PacSum: <https://github.com/mswellhao/PacSum>
- SummaRuNNer: <https://github.com/hpzha0/SummaRuNNer>
- BERTSUM: <https://github.com/nlpyang/PreSumm>. The original BERTSUM model uses a post-processing step called *Trigram Blocking* that excludes a candidate sentence if it has a significant amount of trigram overlap with the already generated summary (to minimize redundancy in the summary). However, we observed that this step leads to summaries that are too short, as also observed in (Sotudeh et al., 2021). Hence we ignore this step.

A.4 Methods for obtaining Training Data for Extractive Supervised Methods

As stated in Section 5, we tried three methods for generating training data for extractive supervised methods from abstractive reference summaries. The best-performing **Avr** method (which we finally used in our experiments) was described in Section 5. Here we describe the other two methods that we tried.

(i) **Maximal:** In this approach proposed in (Nalapati et al., 2017) the basic premise was to maximize the ROUGE score between the extractive and the abstractive gold-standard summaries. However global optimization is computationally expensive; a faster greedy strategy is – keep adding sentences to the extractive summary one by one, each time selecting the sentence that when added to the already extracted summary has the maximum

Model	Fine-tuning parameters
BART	Learning rate - 2e-5, Epochs - 3, Batch size - 1 Max input length - 1024, Max output length - 512
Legal-Pegasus	Learning rate - 5e-5, Epochs - 2, Batch size - 1 Max input length - 512, Max output length - 256
Legal-LED	Learning rate - 1e-3, Epochs - 3, Batch size - 4 Max input length - 16384, Max output length - 1024

Table 9: Hyper-parameters used in finetuning BART, Legal-Pegasus and Legal-LED.

ROUGE score with respect to the abstractive gold-standard summary. This process is repeated till the ROUGE score does not increase anymore. Finally, all the sentences in this extractive summary are labelled as 1, the rest as 0.

(ii) **TF-IDF**: We calculated the TF-IDF vectors for all the sentences in the source document and those in the summary. For each sentence in the summary, we find three sentences in the full text that are most similar to it. The similarity is measured as the cosine-similarity between the TF-IDF vectors of a sentence in the summary and a sentence in the source document, and similarity should be greater than 0.4. We label the sentences in the source document that are similar to some summary-sentence as 1, rest as 0.

A.5 Implementation details of Abstractive Summarization Methods

We use the publicly available implementations of the abstractive methods from the following sources:

- BART: https://huggingface.co/facebook/BART_large
- Legal-Pegasus (trained on legal documents): <https://huggingface.co/nsi319/legal-pegasus>
- Legal-LED (trained on legal documents): <https://huggingface.co/nsi319/legal-led-base-16384>

The hyper-parameters for finetuning are given in Table 9.

A.6 Methods for obtaining finetuning data for abstractive summarization models

As stated in Section 6.2, we experimented with several sentence similarity measures for generating finetuning data for abstractive models. The best performing sentence similarity measure, MCS, was described in Section 6.2. Here we describe the other sentence similarity measures that we tried.

(i) *Smooth Inverse frequency with cosine similarity (SIF)* (Ranasinghe et al., 2019): This approach

is similar to the MCS approach; only here instead of mean, we consider a weighted mean, and we use a pre-trained BERT model. The weight of every token w is given by $\frac{a}{a+p(w)}$. Where $p(w)$ is the estimated frequency of a word in the whole dataset. In other word, the weight for a word would be inversely proportional to the number of word occurrences.

(ii) *Cosine similarity with BERT [CLS] token (CLS-CS)*: Here we consider the cosine similarity of the encodings of the CLS tokens of the two sentences (as given by the pre-trained BERT model).

(iii) *MCS_RR*: Here, we using Rhetorical Roles (RR) for generating finetuning data that incorporates legal domain knowledge. As described earlier in Section 3, a legal case document consists of 7 rhetorical segments such as Facts, Statutes, etc. We incorporate this knowledge into our abstractive summarization process by combining it with the divide and conquer approach presented in (Gidiotis and Tsoumakas, 2020) (which is originally designed for summarizing research articles that are already segmented into logical segments).

We first use a state-of-the-art classifier for rhetorical labeling of sentences in a legal document (Bhattacharya et al., 2021) to assign one of the labels – RPC, FAC, STA, RLC, Ratio, PRE, ARG – to each sentence of a document. We collate sentences of a particular role as one segment. Thus, effectively, we partition a document into 7 segments, each segment corresponding to a rhetorical role. Then we apply the same approach as stated above to generate the summary of each segment; for this, we use the MCS sentence similarity measure (which performs the best, as we shall see later in Section 7). Note that, some of these rhetorical segments themselves may be longer than the input token limit of BART and Pegasus; in such cases, we further divide the rhetorical segments into smaller chunks, and then generate the summary of each chunk.

A.7 Detailed Summarization Results

Table 10, Table 11 and Table 12 contain the *document-wide* ROUGE and BERTScores for the IN-Ext, IN-Abs and UK-Abs datasets respectively. These tables give the results for all summarization methods that we have applied (while the tables in the main text report results of only some of the best-performing methods).

Table 13 and Table 14 contain the *segment-wise* ROUGE scores over the IN-Ext and UK-Abs

Algorithm	ROUGE Scores			BERTScore
	R-1	R-2	R-L	
<i>Extractive Methods</i>				
Unsupervised, Domain Independent				
LexRank	0.564	0.344	0.388	0.862
Lsa	0.553	0.348	0.397	0.875
DSDR	0.566	0.317	0.264	0.834
Luhn	0.568	0.373	0.422	0.882
Reduction	0.561	0.358	0.405	0.869
Pacsum_bert	0.590	0.410	0.335	0.879
Pacsum_tfidf	0.566	0.357	0.301	0.839
Unsupervised, Legal Domain Specific				
MMR	0.563	0.318	0.262	0.833
KMM	0.532	0.302	0.28	0.836
LetSum	0.591	0.401	0.391	0.875
CaseSummarizer	0.52	0.321	0.279	0.835
Supervised, Domain Independent				
SummaRunner	0.532	0.334	0.269	0.829
BERT-Ext	0.589	0.398	0.292	0.85
Supervised, Legal Domain Specific				
Gist	0.555	0.335	0.391	0.864
<i>Abstractive Methods</i>				
Pretrained				
BART	0.475	0.221	0.271	0.833
BERT-BART	0.488	0.236	0.279	0.836
Legal-Pegasus	0.465	0.211	0.279	0.842
Legal-LED	0.175	0.036	0.12	0.799
Finetuned				
BART_CLS	0.534	0.29	0.349	0.853
BART_MCS	0.557	0.322	0.404	0.868
BART_SIF	0.540	0.304	0.369	0.857
BERT_BART_MCS	0.553	0.316	0.403	0.869
Legal-Pegasus_MCS	0.575	0.351	0.419	0.864
Legal-LED	0.471	0.26	0.341	0.863
BART_MCS_RR	0.574	0.345	0.402	0.864

Table 10: Document-wide ROUGE-L and BERTScores (Fscore) on the IN-Ext dataset. All values averaged over the 50 documents in the dataset. The best value in a particular class of methods is in **bold**.

datasets, for all methods that we have applied.

A.8 More Insights from Segment-wise Evaluation

Table 13 shows the segment-wise ROUGE-L Recall scores of all methods on the IN-Ext dataset, considering the 5 rhetorical segments RPC, FAC, STA, ARG, and Ratio+PRE. Similarly, Table 14 shows the segment-wise ROUGE-L Recall scores of all methods on the UK-Abs dataset, considering the 3 segments Background, Reasons, and Final Judgement. In this section, we present some more observations from these segment-wise evaluations, which could not be reported in the main paper due to lack of space.

An interesting observation is that the performances of several methods on a particular segment depend on the *size* and *location* of the said segment in the documents. The FAC (Facts) segment in the

Algorithm	ROUGE Scores			BERTScore
	R-1	R-2	R-L	
<i>Extractive Methods</i>				
Unsupervised, Domain Independent				
LexRank	0.436	0.195	0.284	0.843
Lsa	0.401	0.172	0.259	0.834
DSDR	0.485	0.222	0.27	0.848
Luhn	0.405	0.181	0.268	0.837
Reduction	0.431	0.195	0.284	0.844
Pacsum_bert	0.401	0.175	0.242	0.839
Pacsum_tfidf	0.428	0.194	0.262	0.834
Unsupervised, Legal Domain Specific				
MMR	0.452	0.21	0.253	0.844
KMM	0.455	0.2	0.259	0.843
LetSum	0.395	0.167	0.251	0.833
CaseSummarizer	0.454	0.229	0.279	0.843
Supervised, Domain Independent				
SummaRunner	0.493	0.255	0.274	0.849
BERT-Ext	0.427	0.199	0.239	0.821
Supervised, Legal Domain Specific				
Gist	0.471	0.238	0.308	0.842
<i>Abstractive Methods</i>				
Pretrained				
BART	0.39	0.156	0.246	0.829
BERT-BART	0.337	0.112	0.212	0.809
Legal-Pegasus	0.441	0.19	0.278	0.845
Legal-LED	0.223	0.053	0.159	0.813
Finetuned				
BART_CLS	0.484	0.231	0.311	0.85
BART_MCS	0.495	0.249	0.33	0.851
BART_SIF	0.49	0.246	0.326	0.851
BERT_BART_MCS	0.487	0.243	0.329	0.853
Legal-Pegasus_MCS	0.488	0.252	0.341	0.851
Legal-LED	0.471	0.235	0.332	0.856
BART_MCS_RR	0.49	0.234	0.311	0.849

Table 11: Document-wide ROUGE-L and BERTScores (Fscore) on the IN-Abs dataset, averaged over the 100 test documents. The best value in a particular class of methods is in **bold**.

In-Ext dataset and the Background segment in the UK-Abs dataset are large segments that appear at the beginning of the case documents. On the other hand, the RPC (Ruling by Present Court) segment in In-Ext and the ‘Final judgement’ segment in UK-Abs are short segments appearing at the end of the documents. Most domain-independent models, like Luhn and BERT-Ext, perform much better for the FAC and Background segments, than for the RPC and ‘Final judgement’ segments. Such models may be suffering from the lead-bias problem (Kedzie et al., 2018) whereby a method has a tendency to pick initial sentences from the document for inclusion in the summary.

However, the RPC and ‘Final judgement’ segments are important from a legal point of view, and should be represented well in the summary according to domain experts (Bhattacharya et al., 2019). In fact, the performances of all methods are rela-

Algorithm	ROUGE Scores			BERTScore
	R-1	R-2	R-L	
<i>Extractive Methods</i>				
Unsupervised, Domain Independent				
LexRank	0.481	0.187	0.265	0.848
Lsa	0.426	0.149	0.236	0.843
DSDR	0.484	0.174	0.221	0.832
Luhn	0.444	0.171	0.25	0.844
Reduction	0.447	0.169	0.253	0.844
Pacsum_bert	0.448	0.175	0.228	0.843
Pacsum_tfidf	0.414	0.146	0.213	0.825
Unsupervised, Legal Domain Specific				
MMR	0.440	0.151	0.205	0.83
KMM	0.430	0.138	0.201	0.827
LetSum	0.437	0.158	0.233	0.842
CaseSummarizer	0.445	0.166	0.227	0.835
Supervised, Domain Independent				
SummaRunner	0.502	0.205	0.237	0.846
BERT-Ext	0.431	0.184	0.24	0.821
Supervised, Legal Domain Specific				
Gist	0.427	0.132	0.215	0.819
<i>Abstractive Methods</i>				
Pretrained				
Pointer_Generator	0.420	0.133	0.193	0.812
BERT-Abs	0.362	0.087	0.208	0.803
BART	0.436	0.142	0.236	0.837
BERT-BART	0.369	0.099	0.198	0.805
Legal-Pegasus	0.452	0.155	0.248	0.843
Legal-LED	0.197	0.038	0.138	0.814
Finetuned				
BART_CLS	0.481	0.172	0.255	0.844
BART_MCS	0.496	0.188	0.271	0.848
BART_SIF	0.485	0.18	0.262	0.845
BERT_BART_MCS	0.476	0.172	0.259	0.847
Legal-Pegasus_MCS	0.476	0.171	0.261	0.838
Legal-LED	0.482	0.186	0.264	0.851
BART_MCS_RR	0.492	0.184	0.26	0.839

Table 12: Document-wide ROUGE-L and BERTScores (Fscore) on UK-Abs dataset, averaged over the 100 test documents. The best value for each category of methods is in **bold**.

tively poor for for these segments (see Table 13 and Table 14). Hence, another open challenge in domain-specific long document summarization is to develop algorithms that perform well on short segments that have domain-specific importance.

A.9 Expert Evaluation Details

We mention below some more details of the expert evaluation, which could not be accommodated in the main paper due to lack of space.

Choice of documents for the survey: We selected 5 documents from the IN-Abs test set, specifically, those five documents that gave the best average ROUGE-L F-scores over the 7 summarization methods chosen for the human evaluation.

Ideally, some summaries that obtained lower ROUGE scores should also have been included

Algorithms	Rouge L Recall				
	RPC (6.42%)	FAC (34.85%)	STA (13.42%)	Ratio+Pre (28.83%)	ARG (16.45%)
<i>Extractive Methods</i>					
LexRank	0.039	0.204	0.104	0.208	0.127
Lsa	0.037	0.241	0.091	0.188	0.114
DSDR	0.053	0.144	0.099	0.21	0.104
Luhn	0.037	0.272	0.097	0.175	0.117
Reduction	0.038	0.236	0.101	0.196	0.119
Pacsum_bert	0.038	0.238	0.087	0.154	0.113
Pacsum_tfidf	0.039	0.189	0.111	0.18	0.111
MMR	0.049	0.143	0.092	0.198	0.096
KMM	0.049	0.143	0.1	0.198	0.103
LetSum	0.036	0.237	0.115	0.189	0.1
CaseSummarizer	0.044	0.148	0.084	0.212	0.104
SummaRunner	0.059	0.158	0.08	0.209	0.096
BERT-Ext	0.038	0.199	0.082	0.162	0.093
Gist	0.041	0.191	0.102	0.223	0.093
<i>Pretrained Abstractive Methods</i>					
BART	0.037	0.148	0.076	0.187	0.087
BERT-BART	0.038	0.154	0.078	0.187	0.084
Legal-Pegasus	0.043	0.139	0.076	0.186	0.092
Legal-LED	0.049	0.131	0.078	0.228	0.091
<i>Finetuned Abstractive Methods</i>					
BART_MCS	0.036	0.206	0.082	0.228	0.092
BERT_BART_MCS	0.037	0.205	0.085	0.237	0.094
Legal-Pegasus_MCS	0.037	0.192	0.09	0.257	0.101
Legal-LED	0.053	0.245	0.086	0.187	0.124
BART_MCS_RR	0.061	0.192	0.082	0.237	0.086

Table 13: Segment-wise ROUGE-L Recall scores of all methods on the IN-Ext dataset. All values averaged over the 50 documents in the dataset. The best value for each segment in a particular class of methods is in **bold**.

in the evaluation by the domain experts. But the number of summaries that we could get evaluated was limited by the availability of the experts.

Framing the questions asked in the survey: We framed the set of questions (described in Section 7.3) based on the parameters stated in (Bhat-tacharya et al., 2019; Huang et al., 2020b) about how a legal document summary should be evaluated.

Pearson Correlation as IAA : The human annotators were asked to rate the summaries on a scale of 0-5, for different parameters. Here we discuss the IAA in the ‘Overall’ parameter. For a particular summary of a document, consider that Annotator 1 and Annotator have given scores of 2 and 3 respectively. Now, there are two choices for calculating the IAA – (i) in a regression setup, these scores denote a fairly high agreement between the annotators, (ii) in a classification setup, if we consider each score to be a ‘class’, then Annotator 1 has assigned a ‘class 2’ and Annotator 2 has assigned a ‘class 3’; this implies a total disagreement between the two experts. In our setting, we find the regression setup for calculating IAA more suitable than the Classification setup. Therefore we use Pearson Correlation between the expert scores as the inter-annotator agreement (IAA) measure. For each algorithmic summary, we calculate the corre-

Algorithms	Rouge-L Recall		
	Background (39%)	Final Judgement (5%)	Reasons (56%)
<i>Extractive Methods</i>			
LexRank	0.197	0.037	0.161
Lsa	0.175	0.036	0.141
DSDR	0.151	0.041	0.178
Luhn	0.193	0.034	0.146
Reduction	0.188	0.035	0.158
Pacsum_bert	0.176	0.036	0.148
Pacsum_tfidf	0.154	0.035	0.157
MMR	0.152	0.04	0.17
KMM	0.133	0.037	0.157
LetSum	0.133	0.037	0.147
CaseSummarizer	0.153	0.036	0.17
SummaRunner	0.172	0.044	0.165
BERT-Ext	0.203	0.034	0.135
Gist	0.123	0.041	0.195
<i>Pretrained Abstractive Methods</i>			
BART	0.161	0.04	0.175
BERT-BART	0.143	0.04	0.158
Legal-Pegasus	0.169	0.042	0.177
Legal-LED	0.177	0.066	0.219
<i>Finetuned Abstractive Methods</i>			
BART_MCS	0.168	0.041	0.184
BERT_BART_MCS	0.174	0.047	0.183
Legal-Pegasus_MCS	0.166	0.039	0.202
Legal-LED	0.187	0.058	0.172
BART_MCS_RR	0.165	0.042	0.18

Table 14: Segment-wise ROUGE-L Recall scores of all methods on the UK-Abs dataset. All values averaged over 100 documents in the evaluation set. Best value for each segment in a particular class of methods is in **bold**.

lation between the two sets of ‘Overall’ scores. We then take the average across all the seven ‘Overall’ correlation scores for the seven algorithmic summaries.

Computing the correlation between human judgements and the automatic metrics: Recall that we have 5 documents for the human evaluation. For a particular algorithm, e.g. DSDR, suppose the average ‘Overall score given by human annotators to the summaries of the 5 documents generated by DSDR are $[h_1, h_2, h_3, h_4, h_5]$, where h_i denotes the average ‘Overall’ score given by humans for the i^{th} document’s summary (range [0-1]).

Suppose, the ROUGE-1 FScore of the DSDR summaries (computed with respect to the reference summaries) are $[d_1, d_2, d_3, d_4, d_5]$, where d_i denotes the ROUGE-1 FScore for the i^{th} document’s DSDR-generated summary (range [0-1]).

We then compute the Pearson Correlation c_{DSDR} between the list of human scores and the list of Rouge-1 Fscores for DSDR. We repeat the above procedure for all the 7 algorithms for a particular metric (e.g. ROUGE-1 Fscore) to get 7 c values (e.g., c_{DSDR} , c_{Gist} , etc.) and then take the average of the 7 values. This gives the final corre-

lation between ROUGE-1 Fscore and the overall scores assigned by the human evaluators.

Likewise, we compute the correlation between other automatic metrics (e.g., ROUGE-2 Fscore, BertScore) and the human-assigned overall scores.

A.10 Ethics and limitations statement

All the legal documents and summaries used in the paper are publicly available data on the Web, except the reference summaries for the In-Ext dataset which were written by the Law experts whom we consulted. The law experts were informed of the purpose for which the annotations/surveys were being carried out, and they were provided with a mutually agreed honorarium for conducting the annotations/surveys as well as for writing the reference summaries in the IN-Ext dataset.

The study was performed over legal documents from two countries (India and UK). While the methods presented in the paper should be applicable to legal documents of other countries as well, it is *not* certain whether the reported trends in the results (e.g., relative performances of the various summarization algorithms) will generalize to legal documents of other countries.

The evaluation study by experts was conducted over a relatively small number of summaries (35) which was limited by the availability of the experts. Also, different Law practitioners have different preferences about summaries of case judgements. The observations presented are according to the Law practitioners we consulted, and can vary in case of other Law practitioners.