

Summarizing Long Regulatory Documents with a Multi-Step Pipeline

Mika Sie

Utrecht University
mikasia6@gmail.com

Ruby Beek

Power2X
ruby.beek@power2x.com

Michiel Bots

Power2X
michiel.bots@power2x.com

Sjaak Brinkkemper

Utrecht University
s.brinkkemper@uu.nl

Albert Gatt

Utrecht University
a.gatt@uu.nl

Abstract

Due to their length and complexity, long regulatory texts are challenging to summarize. To address this, a multi-step extractive-abstractive architecture is proposed to handle lengthy regulatory documents more effectively. In this paper, we show that the effectiveness of a two-step architecture for summarizing long regulatory texts varies significantly depending on the model used. Specifically, the two-step architecture improves the performance of decoder-only models. For abstractive encoder-decoder models with short context lengths, the effectiveness of an extractive step varies, whereas for long-context encoder-decoder models, the extractive step worsens their performance. This research also highlights the challenges of evaluating generated texts, as evidenced by the differing results from human and automated evaluations. Most notably, human evaluations favoured language models pretrained on legal text, while automated metrics rank general-purpose language models higher. The results underscore the importance of selecting the appropriate summarization strategy based on model architecture and context length.

1 Introduction

Automatic text summarisation (ATS) involves generating a compressed, concise, and fluent version of an input text while preserving its main key points. A summary proves useful because it helps people process and understand texts faster and better. Summarizing regulatory texts is important for making complex legal language more accessible and ensuring compliance by condensing information into a concise, understandable format.¹

Current ATS methods use either extractive or abstractive summarization. An advantage of extractive summarization is that it captures sentences and

information literally, resulting in a factually consistent summary. However, the summary is harder to read and less intuitive as sentences are copied and combined. Abstractive summaries are more coherent and fluent as they summarize texts in a human-like fashion. But it also has disadvantages because an intricate understanding of the original text is required and the summary can be factually inconsistent. In this paper, our aim is to explore the advantages of both strategies, as we leverage them for the summarisation of very lengthy, regulatory documents.

A regulatory text is a formal document issued by a government or regulatory body that outlines rules, guidelines, or standards to govern the conduct, practices, or operations within a specific industry, sector, or jurisdiction. Regulatory documents are difficult to process due to their extensive size, unique structure, numerous citations and references, ambiguity, and domain-specific vocabulary. Current automatic summarization tools face challenges with regulatory texts, either because their length exceeds the context length of LLMs, or because the length and structure of the input document raise the risk of omissions in the summary. Leaving out important information could have major negative effects.

This paper compares two-step and multi-step summarisation methods for regulatory documents, comparing the effectiveness of different neural model architectures and combinations. Our approach consists of the following steps, illustrated in Figure 1. First, the document is segmented into smaller units or ‘chunks’. Each chunk is then processed by an extractive summarization model, and all resulting summaries are concatenated. This extractive step may need to be conducted iteratively. The outcome of extraction is then summarized in an abstractive manner, creating a final summary. Combining these two summarization steps could prove useful in handling the large size of the original text. It uses extracted salient sentences to de-

¹Code and models are available on [GitHub](#) and [HuggingFace](#).

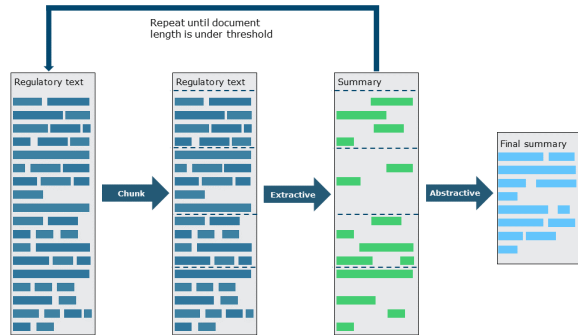


Figure 1: Summarization process proposed by this research. Dotted lines indicate the borders of the chunks.

velop a coherent, fluent summary. Similar architectures have been used on different types of texts and have shown promising results (Pilault et al., 2020; Zhang et al., 2022; Klaus et al., 2022; Bleiweiss, 2023). However, summarizing long regulatory documents using this architecture has been researched less extensively. In particular, our goal is to evaluate various models used for each step to identify the most effective combination of models for the summarization task, paying particular attention to whether preliminary extraction is more beneficial if performed with domain-specific (legal) rather than domain-general models. A second important goal is to compare the effect of context length on the quality of the generated summaries: models allowing longer context lengths need less extraction. Given the growing trend for large language models to allow longer document lengths, it is increasingly important to understand whether such models are able to acquire a comprehensive understanding of a full document, or whether preliminary distillation of information is helpful (Li et al., 2024).

2 Related work

Long document summarisation Pretrained language models (LMs) struggle with long texts due to limitations on input context length. For example, BERT (Devlin et al., 2019) and T5 (Raffel et al., 2020) have a context length of 512 tokens while PEGASUS’s (Zhang et al., 2020a) and BART’s (Lewis et al., 2020) context length is 1024 tokens.

To counter this limitation, some long document architectures incorporate a different self-attention mechanism, calculating attention between specific parts of the sequence instead of calculating the attention for every possible combination of the sequence. This enables them to process long sequences because the computation requirements will

not grow quadratically. Longformer (Beltagy et al., 2020) is an encoder-only architecture based on RoBERTa (Liu et al., 2019), designed to handle long-range dependencies more efficiently than standard transformers, and accepting inputs of up to 4096 tokens. It employs a combination of global attention and sliding window attention instead of full attention, which scales linearly with the input sequence. LED (Beltagy et al., 2020) adds a decoder to the Longformer architecture, turning it into the Longformer Encoder-Decoder model. The decoder does use the full attention mechanism but LED retains its linear computation capability. Similar examples of LMs designed for longer documents include BigBird (which accepts a context length of 4096 tokens; Zaheer et al., 2020), LongT5 (Guo et al., 2022) and PegasusX (Phang et al., 2023), both of which accept contexts of 16,384 tokens.

Extending context length is often a goal in recent releases of decoder-only LLMs, such as the GPT family of models. Other examples include LLaMA-2-7B-32k (Tog), which is an LLM based on LLaMA-2 (Touvron et al., 2023) with a context length of 32768.

Multi-step summarisation The idea of multi-step methods is to leverage both extractive and abstractive techniques to alleviate the burden of summarising very long documents. Pilault et al. (Pilault et al., 2020) add one extractive step before generating the abstractive summary. The extractive parts are then used beside the original text as input for the transformer. A related approach is taken in CreativeSumm (Kim et al., 2022) for the summarisation of lengthy movie scripts. Liu et al. (2018) summarise Wikipedia articles by first performing an extractive step, using the extracted sentences as additional input to the summariser. Bleiweiss (2023) propose a two-step method for long biographical novels. Klaus et al. (2022) make use of a two-step method to summarize legal regulatory documents. Klaus et al. use TextRank (Mihalcea and Tarau, 2004), a graph-based extractive summarization approach, for the first extractive step and BERT (Devlin et al., 2019) or RoBERTa (Liu et al., 2019) for a second extractive step.

A generalisation of the two-step strategy was proposed in the form of Summ^N (Zhang et al., 2022). Summ^N splits the data samples and generates coarse summaries, possibly over multiple stages (N), before producing a final fine-grained abstractive summary. This method outperformed previous

state-of-the-art methods on different datasets. Different from our work, Summ^N makes use of abstractive summarisation for both the coarse-grained and the final, more fine-grained summarisation steps. Instead, we use extractive summarisation for the first stage.

Inspired by multi-step methods, we experiment in this paper with various combinations of extractive and abstractive steps, in an effort to identify the best architecture for summarisation of long, regulatory documents.

Divide-and-conquer (chunking) strategies An interesting class of approaches to long document summarisation involves a ‘divide-and-conquer’ strategy. Briefly, the idea is to chunk the document into sub-parts before summarisation, where sub-part identification may also exploit the document structure. Examples of this are the context-aware chunking strategy for academic articles used in DANCER (Gidiotis and Tsoumakas, 2020) and the work of Shen and Lam (2022), whose model directly learns the correspondence between document sections and summary parts. In our work, we also explore the role of chunking strategies and their effectiveness in producing coherent summaries.

Domain-specific Legal Language Models An important question in the processing of texts in specialised domains is whether in-domain pretraining is beneficial, given that specialised domains have stylistic and other peculiarities. Relevant to the present paper is the case of legal text (of which regulatory texts are a subset), which has well-studied distinctive stylistic characteristics (Turtle, 1995; Kanapala et al., 2019; Jain et al., 2021). Studies have shown that in-domain pretraining can be beneficial in downstream NLP tasks (Gururangan et al., 2020) and domain-specific LMs have been developed for healthcare (Huang et al., 2020; Lee et al., 2020), science (Beltagy et al., 2019) and finance (Yang et al., 2020; Wu et al., 2023), among many others. Pre-trained LMs for law include Lawyer LLaMA (Huang et al., 2023), Lawformer (Xiao et al., 2021), LegalLongformer (Mamakos et al., 2022), PEGASUS-Billsum (Zhang et al., 2020a), LegalBERT (Chalkidis et al., 2020b), CaseLawBERT (Zheng et al., 2021), PoL-BERT (Henderson et al., 2022) and LexLM (Chalkidis et al., 2023). In an early study, Chalkidis et al. (2020b) showed that LegalBERT consistently outperformed BERT-based models on a variety of NLP tasks, including EURLEX57K (Chalkidis et al.,

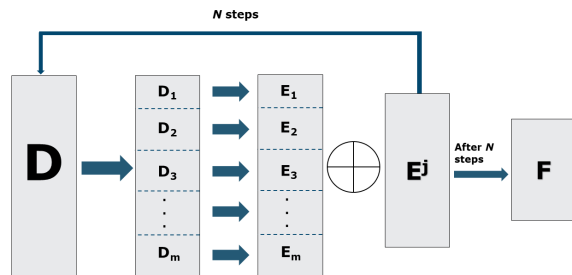


Figure 2: Visualisation of the summarisation process. N represents the amount of extractive steps and the \oplus symbol represents the concatenation process.

2020c), ECHR-CASES (Chalkidis et al., 2020a), and CONTRACTS-NER (Chalkidis et al., 2017). Building on this work, Mamakas et al. (2022) introduced LegalLongformer, initialised with LegalBERT’s parameters, to handle long legal texts. Chalkidis et al. (2023) introduced LexLM, a model pre-trained on a multinational English legal data. Additionally, they introduced a version of LexLM utilizing the Longformer (Beltagy et al., 2020) attention mechanism, enhancing the capability to handle long legal documents. In comparative evaluations, LexLM models outperformed other legal LMs, such as CaseLawBERT and PoL-BERT, particularly in prior knowledge assessment and downstream task performance. Notably, RoBERTa (Liu et al., 2019) also showed strong performance, occasionally surpassing some specialized legal models.

Building on these observations, in our experiments we also compare general-purpose models with a representative subset of legal LMs, particularly for the extractive summarisation step.

3 Method

Our approach to the summarisation of long regulatory documents is a multi-step process consisting of extraction followed by abstraction, where extraction is intended to alleviate the problem of limited context length accepted by a model. In particular, if the length of a source document $|D|$ exceeds the context length K of an abstractive model, creating an intermediate extractive summary could help identify essential information across the document span, a more informed strategy than truncating the document to fit within K .

The overall process is visualised in Figure 2. We view a source document $D = \{D_1, D_2, D_3, \dots, D_m\}$ as a sequence of chunks D_i .

A chunk is summarised by an extractive summarisation model, which produces an intermediary summary $E^j = E_1^j \oplus E_2^j \oplus E_3^j \oplus \dots \oplus E_M^j$, where E_i^j represents an intermediate summary of chunk D_i at extractive step j . Thus, the intermediate summary E^j comprises the summaries of all the chunks concatenated in the same order as in the original text. The extractive summarisation model has a compression ratio $R \in [0, 1]$. One way to define R is in terms of the ration of the length of an article and that of its summary (Grusky et al., 2018); below, we also explore other possible definitions for R . Before the summarisation is performed, the number of extractive steps N taken is determined, such that the extractive summary produced at step $n \leq N$ is the input to the extractive step $n + 1 \leq N$. The extractive summary after N steps is the input to the abstractive summarisation model, which yields the final summary F .

3.1 Dataset

The dataset used to fine-tune the abstractive model is EUR-Lex-Sum (Aumiller et al., 2022). This dataset consists of documents from the European Union law platform with corresponding manually curated summaries. Only the English part of the dataset, composed of 1504 document -summary pairs, was used for this task. It has been divided into training, validation, and test sets, containing 1129 pairs, 187 and 188 pairs, respectively. The dataset is characterised by a small number of documents whose length far exceeds that of the others. To ensure consistency in our evaluation, we define any document whose word count is more than two standard deviations above the mean as an outlier and remove it from the training, validation and test subsets originally provided by Aumiller et al. (2022). In total, 62 instances were removed by this criterion. The final dataset consists of 1091 training, 172 validation and 179 test samples.

3.2 Architecture

Extractive step(s) As described above, documents are first summarised over N extractive steps. Note that the extractive step is only performed if the length of the document exceeds the context length K of the abstractive model. The number of extractive steps needed ultimately depends on the compression ratio R that we require for the summarisation, corresponding to two-step (one extraction step followed by abstraction) and multi-step approaches. We experiment with three different strategies for

computing R for an abstractive model with context length K and a document of length $|D|$. Note that K and $|D|$ are fixed in advance for a given model and document.

Our first strategy is to use a **fixed compression ratio**, empirically setting $R = 0.4$. In this case, $N \geq 1$ and is estimated as follows (see Appendix A.1 for details of how this is derived):

$$N = \left\lceil \frac{\log\left(\frac{K}{|D|}\right)}{\log(R)} \right\rceil \quad (1)$$

The second strategy is to use a **dependent compression ratio**, which depends on the document’s size and the abstractive model’s context length, resulting in $N = 1$:

$$R = \frac{K}{|D|} \quad (2)$$

The final strategy is a **hybrid ratio**, where we perform $N - 1$ extractive steps with a fixed ratio, with a final extractive step N using a dependent ratio. The hybrid ratio could be more effective than the fixed ratio because it is focused on ensuring that the final intermediate summary optimally fits the context length of the abstractive model. We define the hybrid ratio as follows:

$$R = \begin{cases} 0.4 & \text{for steps } 1, 2, \dots, N - 1 \\ \frac{K}{|D|} & \text{for step } N \end{cases} \quad (3)$$

Extractive models One of our goals is to compare the impact of domain-specific LMs and general-purpose LMs. In what follows, non-domain-specific LMs will be referred to as ‘general’ LMs, and domain-specific legal LMs will be referred to as ‘legal’ LMs. The top panel of Table 1 lists all the extractive summarisation models used. Based on this comparison, we aim to identify the optimal extractive model.

We compare all the extractive models with the three ratio types described above, with a view to determining the optimal extractive strategy to support abstractive summarisation. To identify the optimal extractive model, we compare the impact of different extractive models and compression ratios on downstream abstractive summarisation with BART (Lewis et al., 2020). Specifically, we compare the output of a BART summariser, finetuned on using input from different extractive models. We compare this to a baseline BART model with

Model	Context length	Legal LM	Type	Architecture
RoBERTa (Liu et al., 2019)	512	✗	Extractive	Encoder
Longformer (Beltagy et al., 2020)	4096	✗	Extractive	Encoder
LegalBERT-SC (Chalkidis et al., 2020b)	512	✓	Extractive	Encoder
LexLM (Chalkidis et al., 2023)	512	✓	Extractive	Encoder
LexLM - Longformer (Chalkidis et al., 2023)	4096	✓	Extractive	Encoder
BART (Lewis et al., 2020)	1024	✗	Abstractive	Encoder-Decoder
T5 (Raffel et al., 2020)	512	✗	Abstractive	Encoder-Decoder
LongT5 (Guo et al., 2022)	16384	✗	Abstractive	Encoder-Decoder
Pegasus (Zhang et al., 2020b)	1024	✗	Abstractive	Encoder-Decoder
PegasusX (Phang et al., 2022)	16384	✗	Abstractive	Encoder-Decoder
Llama3 (AI, 2024)	8192	✗	Abstractive	Decoder

Table 1: Summarisation models used. The context length is expressed in number of tokens. Top: models used for extractive summarisation; bottom: models used for abstractive summarisation.

no extractive steps. In total, we compare sixteen model configurations. The optimal extractive strategy under this experimental setting was then used to fine-tune subsequent abstractive models.

Abstractive step The abstractive step was only performed once the length of the intermediate summary $|E^j|$ is within the context length K of an abstractive summarisation model. The abstractive step involves creating the final summary F by an abstractive summarisation model fine-tuned on the intermediate summary E^j .

We compare a variety of abstractive models, listed in the bottom panel of Table 1. The context length of the abstractive summarisation model is an important consideration as it affects the number of extractive steps. A longer context length implies that fewer extractive steps need to be taken. By hypothesis, the quality of the final summary should be higher the fewer the extractive steps, since there is less potential in this case for information loss. To quantify this, we chose models that permit a direct comparison of context length effects, while keeping architecture largely constant. We compare T5 (Raffel et al., 2020) against LongT5 (Guo et al., 2022), and Pegasus (Zhang et al., 2020b) against PegasusX (Phang et al., 2022) to determine the effect of a long context length in the abstractive summarization model. Finally, we include Llama3 (AI, 2024), as an example of a SOTA large language model based on a decoder architecture (T5 and Pegasus are encoder-decoder models).

Full parameter fine-tuning was performed for all abstractive models except Llama3, which was fine-tuned using QLoRA (Detmeters et al., 2023) as full parameter fine-tuning was not feasible due to its size. Data had to be prepared in a different way for Llama3 as it is the sole decoder-only model used in our experiments. A single combined se-

quence is used instead of separate input and output sequences. To accommodate a summary of 1500 tokens, 1500 tokens are subtracted from the model’s context length, resulting in an effective context length of 6692 tokens for Llama3. The extractive summarisation process was adjusted to summarise the reference text to fit within this 6692-token limit, ensuring minimal truncation. See Appendix A.2 and A.3 for more details on model finetuning, including hyperparameters.

3.3 Evaluation

Evaluation metrics Multiple evaluation metrics were used to assess the proposed architecture from different aspects. This research employed ROUGE-1, ROUGE-2, ROUGE-L (Lin, 2004), BERTScore (Zhang et al., 2020b), BARTScore (Yuan et al., 2021), and BLANC (Vasilyev et al., 2020a). Details of the implementations used for the evaluation metrics are in Appendix A.4.

Expert evaluation Besides automated metrics, we also performed a small-scale qualitative human evaluation involving expert readers. The human evaluation provides insights into the quality of the summaries, complementing the quantitative data from automated metrics with qualitative feedback. After selecting the optimal extractive model and training the abstractive models, we generate summaries of a new text which is not in the training dataset.² Summaries generated with the different abstractive models were compared by the expert readers. This document was chosen specifically because the expert readers were already familiar with the contents and, hence, were able to judge summary quality more reliably.

The evaluators were two experts from the com-

²The text in question is the Carbon Border Adjust Mechanism document (European Union, 2023).

Criterion	Description
Factual Correctness	Evaluation of how factually correct the summary is relative to the source document.
Usability	Assessment of how practical and user-friendly the summary is.
Accuracy	Assessment of the precision and correctness of the information in the summary.
Fluency	Assessment of the summary’s smoothness and ease of reading in terms of form, content, and grammar.
Coherence	Measure of how logical the summary is to it is linguistic context.

Table 2: Criteria for human evaluation.

pany ANON, a collaborator on this project whose personnel have extensive experience with regulatory documents issued by the European Union. The experts were asked to read summaries generated by different summarization architectures and evaluate them based on a set of criteria. The criteria included *Factual Correctness*, *Usability*, *Accuracy*, *Fluency*, and *Coherence*. Each criterion was rated on a scale from 1 to 5. Detailed descriptions of these criteria can be found in Table 2 and are based on the findings of Howcroft et al. (2020)’s meta-review of constructs used in human evaluation of Natural Language Generation systems. In addition to scoring the summaries, experts were also asked to comment on the quality of summaries.

Due to resource and time constraints, we selected specific architectures to be included in the qualitative evaluation. To analyse the impact of different extractive models, we compare different versions of BART, using (1) the best extractive model; (2) no extractive step; (3) the best legal LM for extraction; and (4) the best long-context extractive model. To analyse the impact of different abstractive strategies, we also include (5) the best long-context abstractive model; and (6) the best decoder-only model.

4 Results

4.1 Comparison of extractive models

Table 3 contains the results on different metrics for abstractive summarisation using BART, in combination with different extractive strategies. It can be seen that RoBERTa with a dependent ratio scores the highest on ROUGE-1, ROUGE-2, ROUGE-L, and BERTScore. RoBERTa with a hybrid ratio achieves the highest score on BLANC. On the other hand, the best BARTScore is obtained when we do not combine any extractive summarisation to compress the input to BART.

In the rest of this section, we discuss these results in light of the different experimental conditions.

4.1.1 Effect of number of extractive stages

The results indicate that models using the dependent ratio type generally achieve higher performance across most metrics. Notably, the RoBERTa model with the dependent ratio type attains the highest scores in ROUGE-1 (0.4873), ROUGE-2 (0.1974), ROUGE-L (0.2247), and BERTScore (0.8721), suggesting superior performance in these areas. However, the BART model without any extractive steps achieves the best scores in BARTScore (-3.4154) and BLANC (0.1700), indicating a stronger performance in these specific metrics despite not utilizing extraction.

Using a multi-step architecture, that is, one that performs multiple extractive iterations (up to N ; see Section 3), sentences from different document chunks get combined during the summarization process. This could introduce noise and consequently fail to capture the most relevant and coherent information, resulting in lower performance. It seems that using a single extractive step is more effective at capturing the most important sentences out of a chunk relative to the context of the global document. We hypothesise that this explains the superiority of the dependent ratio (where $N = 1$) on most metrics.

Effect of Legal Language Models General-purpose LMs such as RoBERTa achieve slightly higher scores across all metrics except BARTScore, compared to legal LMs. For this comparison, RoBERTa was compared against LegalBERT and LexLM, and Longformer was compared against LexLM-Longformer to accommodate for the context lengths.

These results indicate that, when used as extractors for preliminary document compression, the broad range of training data types that general-purpose LMs are exposed to gives them an advantage in locating important information in the document. In contrast, legal LMs can suffer from a ‘narrow’ focus, resulting in less coherent and com-

Extractive model	Ratio type	R1	R2	RL	BERTScore	BARTScore	BLANC
N/A	No extraction	0.4590	0.1954	0.2174	0.8702	-3.4154	0.1029
RoBERTa	Fixed	0.4670	0.1798	0.2171	0.8692	-3.5654	0.1040
RoBERTa	Dependent	0.4873	0.1974	0.2247	0.8721	-3.5590	0.1272
RoBERTa	Hybrid	0.4809	0.1889	0.2193	0.8700	-3.5781	0.1296
LegalBERT	Fixed	0.4390	0.1766	0.2158	0.8700	-3.4893	0.1099
LegalBERT	Dependent	0.4619	0.1854	0.2174	0.8713	-3.5143	0.1117
LegalBERT	Hybrid	0.4469	0.1774	0.2137	0.8665	-3.5714	0.1098
LexLM	Fixed	0.4571	0.1745	0.2123	0.8692	-3.6130	0.1154
LexLM	Dependent	0.4859	0.1954	0.2227	0.8713	-3.5441	0.1277
LexLM	Hybrid	0.4582	0.1792	0.2135	0.8665	-3.5639	0.1102
Longformer	Fixed	0.4436	0.1686	0.2103	0.8684	-3.5901	0.1029
Longformer	Dependent	0.4613	0.1874	0.2194	0.8712	-3.5835	0.1238
Longformer	Hybrid	0.4778	0.1862	0.2181	0.8703	-3.5697	0.1256
LexLM-Longformer	Fixed	0.4250	0.1584	0.2041	0.8659	-3.6141	0.0959
LexLM-Longformer	Dependent	0.4751	0.1852	0.2164	0.8689	-3.5344	0.1272
LexLM-Longformer	Hybrid	0.4619	0.1819	0.2189	0.8692	-3.5833	0.1199

Table 3: Results for all extractive summarization models in combination with BART.

Abstractive model	Ratio type	R1	R2	RL	BERTScore	BARTScore	BLANC
BART	No extraction	0.4590	0.1954	0.2174	0.8702	-3.4154	0.1029
BART	Dependent	0.4873	0.1974	0.2247	0.8721	-3.5590	0.1272
T5	No extraction	0.3033	0.1241	0.1994	0.8443	-2.1585	0.0760
T5	Dependent	0.2934	0.0926	0.1857	0.8404	-2.2234	0.0812
LongT5	No extraction	0.3261	0.1309	0.2192	0.8497	-2.2195	0.1128
LongT5	Dependent	0.2854	0.0969	0.0969	0.8444	-2.0423	0.1051
Pegasus	No extraction	0.3305	0.1293	0.2260	0.8499	-1.8067	0.0923
Pegasus	Dependent	0.3067	0.0911	0.2021	0.8435	-1.8940	0.0952
PegasusX	No extraction	0.3673	0.1622	0.2304	0.8523	-2.4528	0.1086
PegasusX	Dependent	0.3052	0.1162	0.1960	0.8413	-2.4305	0.0999
Llama3	No extraction	0.4088	0.1816	0.2107	0.7854	-3.3424	0.1177
Llama3	Dependent	0.4474	0.1885	0.2284	0.8687	-3.1268	0.1231

Table 4: Evaluation results of all abstractive models with and without an extractive step.

	Extr. model	Ratio	Abstr. model	FC	U	Acc	Fl	Coh
1	RoBERTa	Dep.	BART	2.0	2.0	1.5	1.5	2.0
2	-	NE	BART	3.5	1.0	2.0	3.0	1.5
3	LexLM	Dep.	BART	4.0	3.5	3.0	3.0	3.0
4	Longformer	Dep.	BART	2.0	2.0	2.5	1.5	2.0
5	-	NE	PegasusX	3.5	1.0	2.5	3.0	1.0
6	RoBERTa	Dep.	Llama3	3.0	2.5	2.5	2.5	2.0

Table 5: Average human evaluation results. Dep: Dependent ratio; NE: No extraction; FC: Factual Correctness; U: Usability; Acc: Accuracy; Fl: Fluency; Coh: Coherence

prehensive extractive summaries. This insight suggests that general LMs can be effective for domain-specific tasks, at least for preparatory steps such as the one considered here.

Effect of context length for the extractive step

Models with shorter context lengths for the extractive step achieve higher scores across all metrics. RoBERTa was compared against Longformer for general LMs and LegalBERT and LexLM against LexLM-Longformer for legal LMs. This approach ensures a fair comparison by accommodating general and legal language model differences.

This finding is surprising, since one would assume that longer-context models would perform

better by capturing more global context. However, when sequences are excessively long, the models might struggle to maintain and encode all relevant information, leading to reduced sensitivity to portions of the input, in line with findings such as those reported by [Fu et al. \(2023\)](#), among others.

This could explain why shorter context models, which deal with more manageable chunks of information, consistently perform better in the extraction task.

Optimal extractive model Based on Table 3, RoBERTa with a dependent ratio will be chosen as the optimal extractive model and is used in the remainder of the experiments reported below.

4.2 Comparison of abstractive models

For every abstractive model, two versions are compared: one leveraging RoBERTa with a dependent ratio and one without using any extractive step at all. The results for all abstractive models and their variants can be seen in Table 4. For clarity, models that incorporate an extractive step will be referred to by the name of the abstractive model. Models that do not use an extractive step will be denoted by appending “-NE” to the name of the abstractive model, where “NE” signifies “No Extraction”.

Effect of extractive step The performance of encoder-decoder abstractive summarization models generally worsens when using one extractive step, though this differs per model. This is evident in the results for T5, LongT5, Pegasus, and PegasusX, where the versions without extraction tend to outperform their counterparts with an extractive step. BART presents a more varied picture as it differs per metric in which variant scores higher. Since encoder-decoder models generate a condensed representation of the text, one explanation for these results is that by introducing an intermediate extractive summary we compromise the performance of the encoder. This could happen because the intermediate summary is less coherent than the input document as a whole.

LLama3, the decoder-only model seems to benefit from an additional extractive step, obtaining better results on all metrics when compared to the version with no extraction. The beneficial effect of extraction here is likely due to the limited context of Llama3 and the risk of loss of sensitivity to longer inputs, as decoding proceeds (Fu et al., 2023). These shortcomings could be mitigated by performing some preliminary input compression and identification of core information.

Effect of context length for the abstractive step

Long context models generally outperform their short context counterparts, with some exceptions. Long context models without an extractive step outperform short context models without an extractive step on all metrics, except BARTScore. When an extractive step is used, results vary as short context models show advantages on specific metrics. In other words, models with shorter input contexts benefit from input compression, as expected. Long context models without an extractive step generally outperform short context models with an extractive step across all metrics.

4.3 Human evaluation

Human evaluation scores are in Table 5. Experts’ individual scores and comments are in Appendix B. Recall that the human evaluation was performed after selecting the optimal extractive model and fine-tuning all abstractive models. Overall, the expert evaluators preferred architectures that relied on a legal LM or a long context model in the extractive step. Indeed, the model that was preferred across all criteria was BART coupled with a LexLM extractor with a dependent compression ratio. The experts’ comments suggested that this architecture did have shortcomings, but these were counterbalanced by other factors. For example, one expert noted that the summary correctly grasps the key points of the regulation, making it quite useful, despite the fact that it is incomplete and has shortcomings on fluency and coherence.

Common criticisms of the summaries by the experts included excessive repetition in the case of some architectures, which severely decrease the quality of the produced summary. Furthermore, while some summaries may appear well-structured and readable, they fail to capture the essential points of the regulation or contain factual errors.

A somewhat surprising outcome is that LLaMA-3 scores relatively poorly on coherence and fluency, compared to the best-performing model. It should be noted that the two evaluators diverged significantly in their scores for this model on these criteria (compare Tables 8 and 9 in Appendix B). Furthermore, as noted above, LLaMA was treated somewhat differently since it is the only decoder-only model. In particular, we subtracted 1500 tokens from the model’s context length to accommodate the extractive summary; this too could have impacted results, though we adjusted the extractive summarisation process to ensure minimal truncation.

Despite the fact that this is a small-scale evaluation (a point we return to in Section 5), there are interesting divergences between expert judgments and the conclusions drawn based on the automatic metrics, an observation which is quite common in the NLG and summarisation literature (cf. Belz and Reiter, 2006; Reiter, 2018; Celikyilmaz et al., 2021).

In particular, experts suggest that legal LMs help achieve more satisfactory summaries if used in the extractive step. On the other hand, both automatic and human evaluation suggest that BART is a com-

petitive model for summarisation, especially if preceded by an extractive step.

5 Conclusion

In this paper, we focused on summarisation of long regulatory documents. Our findings indicate that while models with a longer context length do not benefit from extraction, an extractive step renders BART, an encoder-decoder architecture, highly competitive. A small-scaled evaluation with human experts confirms this finding. However, experts also indicate a preference for summaries relying on extraction with a domain-specific, legal language model.

Future work should consider whether these findings are generalisable to other domains. Furthermore, a more extensive human evaluation is required to ensure that our findings are reliable. This is particularly crucial given that human expert judgments are not perfectly aligned with the outcomes of our metric-based evaluation, which echoes findings from other studies. A further possible research direction is to use a state-of-the-art LLM as an evaluator or ‘judge’ for generated texts, a strategy which recent research suggests is increasingly viable (Liu et al., 2023; Zheng et al., 2023), though also one that requires some caution in view of results suggesting self-bias on the part of LLMs (Panickssery et al., 2024), as well as lower reliability in comparison with expert judgment (Bavaresco et al., 2024).

References

- Llama-2-7B-32K-Instruct — and fine-tuning for Llama-2 models with Together API.
- Meta AI. 2024. [Introducing meta llama 3: The most capable openly available llm to date](#). Accessed: 2024-06-04.
- Dennis Aumiller, Ashish Chouhan, and Michael Gertz. 2022. [Eur-lex-sum: A multi- and cross-lingual dataset for long-form summarization in the legal domain](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022*.
- Anna Bavaresco, Raffaella Bernardi, Leonardo Bertolazzi, Desmond Elliott, Raquel Fernández, Albert Gatt, Esam Ghaleb, Mario Giulianelli, Michael Hanna, Alexander Koller, André F. T. Martins, Philipp Mondorf, Vera Neplenbroek, Sandro Pezzelle, Barbara Plank, David Schlangen, Alessandro Suglia, Aditya K. Surikuchi, Ece Takmaz, and Alberto Testoni. 2024. [LLMs instead of Human Judges? A Large Scale Empirical Study across 20 NLP Evaluation Tasks](#). *arXiv preprint*.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [Scibert: A pretrained language model for scientific text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing*.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *Preprint*, arXiv:2004.05150.
- Anja Belz and E Reiter. 2006. [Comparing Automatic and Human Evaluation of NLG Systems](#). In *Proceedings of the 11th Conference of the European Chapter of the Association for Computational Linguistics (EACL’06)*, pages 313–320.
- Avi Bleiweiss. 2023. [Two-step text summarization for long-form biographical narrative genre](#). In *Proceedings of the 4th Workshop on Computational Approaches to Discourse (CODI 2023)*, pages 145–155, Toronto, Canada. Association for Computational Linguistics.
- Asli Celikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2021. [Evaluation of Text Generation: A Survey](#). *arXiv preprint*. ArXiv:2006.14799 [cs].
- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2020a. [Neural legal judgment prediction in english](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Ilias Chalkidis, Ion Androutsopoulos, and Achilleas Michos. 2017. [Extracting contract elements](#). In *Proceedings of the International Conference on Artificial Intelligence and Law*.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020b. [Legal-bert: The muppets straight out of law school](#). In *Findings of the Association for Computational Linguistics Findings of ACL: EMNLP 2020*.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, and Ion Androutsopoulos. 2020c. [Large-scale multi-label text classification on eu legislation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Ilias Chalkidis, Nicolas Garneau, Anders Søgaard, Cătălină Goantă, and Daniel Martin Katz. 2023. [Lex-files and legallama: Facilitating english multinational legal language model development](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#). *Preprint*, arXiv:2305.14314.

- Jacob Devlin, Ming Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 1.
- European Union. 2023. [Regulation \(eu\) 2023/0956 of the european parliament and of the council of 10 may 2023 on machinery](#). Accessed: 2024-06-28.
- Zihao Fu, Wai Lam, Qian Yu, Anthony Man-Cho So, Shengding Hu, Zhiyuan Liu, and Nigel Collier. 2023. [Decoder-Only or Encoder-Decoder? Interpreting Language Model as a Regularized Encoder-Decoder](#). *arXiv preprint*. ArXiv:2304.04052 [cs].
- Alexios Gidiotis and Grigorios Tsoumakas. 2020. [A divide-and-conquer approach to the summarization of long documents](#). *IEEE/ACM Transactions on Audio Speech and Language Processing*, 28.
- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. [Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 1.
- Mandy Guo, Joshua Ainslie, David Uthus, Santiago Ontañón, Jianmo Ni, Yun Hsuan Sung, and Yinfei Yang. 2022. [Longt5: Efficient text-to-text transformer for long sequences](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*.
- Suchin Gururangan, Ana Marasovic, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don't stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Peter Henderson, Mark S. Krass, Lucia Zheng, Neel Guha Christopher D. Manning, Dan Jurafsky, and Daniel E. Ho. 2022. [Pile of law: Learning responsible data filtering from the law and a 256gb open-source legal dataset](#). In *Advances in Neural Information Processing Systems*, volume 35.
- David M. Howcroft, Anya Belz, Miruna Clinciu, Dimitra Gkatzia, Sadid A. Hasan, Saad Mahamood, Simon Mille, Emiel van Miltenburg, Sashank Santhanam, and Verena Rieser. 2020. [Twenty years of confusion in human evaluation: Nlg needs evaluation sheets and standardised definitions](#). In *Proceedings of the 13th International Conference on Natural Language Generation*.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685.
- Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. 2020. [Clinicalbert: Modeling clinical notes and predicting hospital readmission](#). *Preprint*, arXiv:1904.05342.
- Quzhe Huang, Mingxu Tao, Chen Zhang, Zhenwei An, Cong Jiang, Zhibin Chen, Zirui Wu, and Yansong Feng. 2023. [Lawyer llama technical report](#). *Preprint*, arXiv:2305.15062.
- Deepali Jain, Malaya Dutta Borah, and Anupam Biswas. 2021. [Summarization of legal documents: Where are we now and the way forward](#). *Computer Science Review*, 40.
- Ambedkar Kanapala, Sukomal Pal, and Rajendra Pambula. 2019. [Text summarization from legal documents: a survey](#). *Artificial Intelligence Review*, 51.
- Eunchong Kim, Taewoo Yoo, Gunhee Cho, Suyoung Bae, and Yun-Gyung Cheong. 2022. [The creativesumm 2022 shared task: A two-stage summarization model using scene attributesutterances](#). In *Proceedings of The Workshop on Automatic Summarization for Creative Writing*.
- Svea Klaus, Ria Van Hecke, Kaweh Djafari Naini, Ismail Sengor Altingovde, Juan Bernabé-Moreno, and Enrique Herrera-Viedma. 2022. [Summarizing legal regulatory documents using transformers](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. [Biobert: A pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2020. [Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Jiaqi Li, Mengmeng Wang, Zilong Zheng, and Muhan Zhang. 2024. [LoogLE: Can Long-Context Language Models Understand Long Contexts?](#) *arXiv preprint*. ArXiv:2311.04939 [cs].
- C Y Lin. 2004. [Rouge: A package for automatic evaluation of summaries](#). *Proceedings of the workshop on text summarization branches out (WAS 2004)*.
- Peter J. Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Łukasz Kaiser, and Noam Shazeer. 2018. [Generating wikipedia by summarizing long sequences](#). In *Proceedings of the 6th International Conference on Learning Representations, ICLR 2018*.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-Eval: NLG Evaluation using Gpt-4 with Better Human Alignment](#). In *Proceedings of the 2023 Conference*

- on *Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *Preprint*, arXiv:1907.11692.
- Dimitris Mamakas, Petros Tsotsi, Ion Androutsopoulos, and Ilias Chalkidis. 2022. [Processing long legal documents with pre-trained transformers: Modding legalbert and longformer](#). In *Proceedings of the Natural Legal Language Processing Workshop 2022*.
- Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into texts. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence rnns and beyond](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*.
- Arjun Panickssery, Samuel R. Bowman, and Shi Feng. 2024. [LLM Evaluators Recognize and Favor Their Own Generations](#). *arXiv preprint*. ArXiv:2404.13076 [cs].
- Jason Phang, Yao Zhao, and Peter Liu. 2023. [Investigating efficiently extending transformers for long input summarization](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3946–3961, Singapore. Association for Computational Linguistics.
- Jason Phang, Yao Zhao, and Peter J. Liu. 2022. [Investigating efficiently extending transformers for long input summarization](#). *Preprint*, arXiv:2208.04347.
- Jonathan Pilault, Raymond Li, Sandeep Subramanian, and Christopher Pal. 2020. [On extractive and abstractive neural document summarization with transformer language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21.
- Sebastian Raschka. 2023. [Practical tips for finetuning llms using lora \(low-rank adaptation\)](#). Accessed: 2024-06-10.
- Ehud Reiter. 2018. [A Structured Review of the Validity of BLEU](#). *Computational Linguistics*, 44(3).
- Xin Shen and Wai Lam. 2022. [Improved divide-and-conquer approach to abstractive summarization of scientific papers](#). In *Proceedings of the 4th International Conference on Natural Language Processing, ICNLP 2022*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and finetuned chat models](#). *Preprint*, arXiv:2307.09288.
- Howard Turtle. 1995. [Text retrieval in the legal world](#). *Artificial Intelligence and Law*, 3.
- Oleg Vasilyev, Vedant Dharnidharka, and John Bohannon. 2020a. [Fill in the BLANC: Human-free quality estimation of document summaries](#). In *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems*, pages 11–20, Online. Association for Computational Linguistics.
- Oleg V. Vasilyev, Vedant Dharnidharka, Nicholas Egan, Charlene Chambliss, and John Bohannon. 2020b. [Sensitivity of BLANC to human-scored qualities of text summaries](#). *CoRR*, abs/2010.06716.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David Rosenberg, and Gideon Mann. 2023. [Bloomberggpt: A large language model for finance](#). *Preprint*, arXiv:2303.17564.
- Chaojun Xiao, Xueyu Hu, Zhiyuan Liu, Cunchao Tu, and Maosong Sun. 2021. [Lawformer: A pre-trained language model for chinese legal long documents](#). *AI Open*, 2.
- Yi Yang, Mark Christopher Siy UY, and Allen Huang. 2020. [Finbert: A pretrained language model for financial communications](#). *Preprint*, arXiv:2006.08097.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [Bartscore: Evaluating generated text as text generation](#). In *Advances in Neural Information Processing Systems*, volume 33.

- Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. Big bird: Transformers for longer sequences. In *Advances in Neural Information Processing Systems*, volume 2020-December.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2020a. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020*, volume PartF168147-15.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020b. Bertscore: Evaluating text generation with bert. In *Proceedings of the 8th International Conference on Learning Representations, ICLR 2020*.
- Yusen Zhang, Ansong Ni, Ziming Mao, Chen Henry Wu, Chenguang Zhu, Budhaditya Deb, Ahmed H. Awadallah, Dragomir Radev, and Rui Zhang. 2022. Summn: A multi-stage summarization framework for long input dialogues and documents. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, volume 1.
- Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, Alban Desmaison, Can Balioglu, Pritam Damania, Bernard Nguyen, Geeta Chauhan, Yuchen Hao, Ajit Mathews, and Shen Li. 2023. [Pytorch fsdp: Experiences on scaling fully sharded data parallel](#). *Preprint*, arXiv:2304.11277.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena](#). In *Proceedings of the NeurIPS 2023 Datasets and Benchmarks Track*.
- Lucia Zheng, Neel Guha, Brandon R. Anderson, Peter Henderson, and Daniel E. Ho. 2021. [When does pretraining help?: Assessing self-supervised learning for law and the casehold dataset of 53,000+ legal holdings](#). In *Proceedings of the 18th International Conference on Artificial Intelligence and Law, ICAIL 2021*.

A Further details on the method

A.1 Derivation of N

The following is the derivation of Equation 1:

1. The length of the intermediary summary $|E^j|$ after the first step is $R \cdot |D|$. After the second step, it is $R^2 \cdot |D|$ and so on. This implies that the length of the intermediary summary after N steps is:

$$|E^N| = R^N \cdot |D|$$

2. Extractive steps are performed until the length of the intermediary summary is within the context length of the abstractive summarisation model, K :

$$R^N \cdot |D| \leq K$$

3. To estimate N , take the logarithm on both sides:

$$N \cdot \log(R) \leq \log\left(\frac{K}{|D|}\right)$$

4. Then, solve for N :

$$N \leq \frac{\log\left(\frac{K}{|D|}\right)}{\log(R)}$$

5. N is then rounded up to the highest integer. So, the formula for estimating the number of extractive steps N needed before the final abstractive step can be taken is:

$$N = \left\lceil \frac{\log\left(\frac{K}{|D|}\right)}{\log(R)} \right\rceil \quad (4)$$

A.2 Hyperparameter settings

Table 6 summarises the hyperparameters used to finetune BART, T5, LongT5, Pegasus and PegasusX.

Hyperparameter	Setting
Learning rate	$5e^{-05}$
Epochs	40
Effective batch size	16
Warmup ratio	0.1
Weight decay	0.01
Early stopping patience	5
Metric for best model	Validation loss
Maximum generation length	1500

Table 6: Hyperparameter settings for BART, T5, LongT5, Pegasus and PegasusX.

A.3 Llama3 hyperparameter settings and training procedure

Table 7 shows the hyperparameters used to finetune Llama3 on the abstractive evaluation task.

Hyperparameter	Setting
Learning rate	$5e^{-05}$
Epochs	10
Effective batch size	16
Warmup ratio	0.1
Weight decay	0.01
Early stopping patience	-
Metric for best model	-
LoRA rank (r)	8
LoRA alpha	16
LoRA dropout	0.1
Precision for frozen model weights	4-bit NF
Precision for low-rank matrices	bfloat16
Precision for calculations	bfloat16
Double Quantization	True

Table 7: Llama3 settings.

Fully Sharded Data Parallel (FSDP) (Zhao et al., 2023) was used to fine-tune Llama3 with the [Hugging Face implementation](#). Due to issues when combining FSDP and QLoRA, the best-performing model could not be loaded, and early stopping patience and best model metric were not set. To mitigate overfitting, we used 10 epochs instead of 40, based on preliminary results indicating convergence between 4-20 epochs. For QLoRA, low-rank matrices were injected into the query, key, value matrices, and linear layers of Llama3, following settings from prior research (Raschka, 2023) (Hu et al., 2021). To fine-tune Llama3, we combined the reference text and golden reference summary into a single sequence, providing Llama3 with the following input sequence:

```
Summarise the following text.
### Text:
{reference text}
### Summary:
{golden reference summary}
```

During prediction, no exemplary summary was given, allowing Llama3 to create a new summary.

A.4 Evaluation metrics details

We implemented ROUGE (Lin, 2004) and BERTScore (Zhang et al., 2020b) using the [HuggingFace evaluate library](#), comparing predictions against reference summaries using F-scores. For BERTScore, we employed the Longformer (Beltagy et al., 2020) architecture for its long context length. BARTScore (Yuan et al., 2021) was implemented with Stanford’s [string2string library](#), using BART(Lewis et al., 2020) fine-tuned on the CNN/Daily Mail dataset (Nallapati et al., 2016). BARTScore calculates precision and recall based on log-likelihood, combined into an F-score, and is limited by BART’s 1024-token context length. We used BLANC-help (Vasilyev et al., 2020a) from the [BLANC package](#), with a gap of two as this best correlates with human evaluation (Vasilyev et al., 2020b). BLANC, using BERT base (Devlin et al., 2019), is limited by its 512-token context.

B Human evaluation results

Individual results for the two expert evaluations on each criterion are shown in Tables 8 and 9. These results are the basis for the averaged results in Section 4.3 in the main paper. Below, we also summarise the main observations from the evaluators’ comments on the summary outputs, for each architecture (architectures are numbered according to the order in the tables).

Architecture 1 The evaluators indicated that the summary is not usable for readers without prior knowledge of the topic due to its incompleteness, factual mistakes, and inaccuracies. While it does touch upon the main principle of CBAM, some of the procedures and rules are described incorrectly.

Architecture 2 The evaluators indicated that the summary is not usable for readers as it places information in the wrong place, describing background details in the ‘key points’ section instead of the main content of the regulation. Additionally, one evaluator mentions that the summary completely misses the main point of what CBAM is, despite the

state information being mostly correct with only a few mistakes.

Architecture 3 One evaluator indicates that the summary correctly grasps the key points of the regulation, making it quite useful. However, the evaluator noted that it is not fully complete and that the fluency and coherence of the sentences could be improved. Despite these shortcomings, the summary is considered a good starting point.

Architecture 4 One evaluator noted that this summary is less flawed than that generated by Architecture 1 but is still unusable due to containing a significant amount of false information and incorrect words.

Architecture 5 Both mentioned that the summary contains excessive repetitions. Although the summary starts well, its usability degrades as more repetitions are encountered.

Architecture 6 One evaluator states that the summary contains quite some useful information. However because the summary contains a lot of repetition, it becomes unusable.

Architecture #	Extr. model	Ratio	Abstr. model	FC	U	Acc	Fl	Coh
1	RoBERTa	Dep.	BART	1	1	1	1	1
2	-	NE	BART	3	1	1	2	1
3	LexLM	Dep.	BART	4	3	3	2	2
4	Longformer	Dep.	BART	1	1	1	1	1
5	-	NE	PegasusX	4	1	2	1	1
6	RoBERTa	Dep.	Llama3	3	1	2	1	1

Table 8: Human evaluation results participant 1. Dep: Dependent ration; NE: No extraction; FC: Factual Correctness; U: Usability; Acc: Accuracy; Fl: Fluency; Coh: Coherence

Architecture #	Extr. model	Ratio	Abstr. model	FC	U	Acc	Fl	Coh
1	RoBERTa	Dep.	BART	3	3	2	2	3
2	-	NE	BART	4	1	3	4	2
3	LexLM	Dep.	BART	4	4	3	4	4
4	Longformer	Dep.	BART	3	3	4	2	3
5	-	NE	PegasusX	3	1	3	5	1
6	RoBERTa	Dep.	Llama3	3	4	3	4	3

Table 9: Human evaluation results participant 2. Dep: Dependent ration; NE: No extraction; FC: Factual Correctness; U: Usability; Acc: Accuracy; Fl: Fluency; Coh: Coherence