

Processing Long Legal Documents with Pre-trained Transformers: Modding LegalBERT and Longformer

Dimitris Mamakas*[†] Petros Tsotsi*[†]
Ion Androutsopoulos[†] Ilias Chalkidis^{‡,◊}

[†] Department of Informatics, Athens University of Economics and Business, Greece

[‡] Department of Computer Science, University of Copenhagen, Denmark

[◊] Cognitiv+, Athens, Greece

Abstract

Pre-trained Transformers currently dominate most NLP tasks. They impose, however, limits on the maximum input length (512 sub-words in BERT), which are too restrictive in the legal domain. Even sparse-attention models, such as Longformer and BigBird, which increase the maximum input length to 4,096 sub-words, severely truncate texts in three of the six datasets of LexGLUE. Simpler linear classifiers with TF-IDF features can handle texts of any length, require far less resources to train and deploy, but are usually outperformed by pre-trained Transformers. We explore two directions to cope with long legal texts: (i) modifying a Longformer warm-started from LegalBERT to handle even longer texts (up to 8,192 sub-words), and (ii) modifying LegalBERT to use TF-IDF representations. The first approach is the best in terms of performance, surpassing a hierarchical version of LegalBERT, which was the previous state of the art in LexGLUE. The second approach leads to computationally more efficient models at the expense of lower performance, but the resulting models still outperform overall a linear SVM with TF-IDF features in long legal document classification.

1 Introduction

Transformer-based models (Vaswani et al., 2017), like BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), and their numerous offspring, currently dominate most natural language processing (NLP) tasks. These models are pre-trained on very large corpora using generic tasks (e.g., masked token prediction) that do not require human annotations, and are then fine-tuned (further trained) on typically much smaller task-specific datasets with manually annotated ground truth. The quadratic complexity of their attention mechanisms, however, imposes limits on the maximum input length

* Equal contribution.




Figure 1: Comparison of examined models presented in Section 3, considering averaged down-stream performance and efficiency (inference time) in LexGLUE long document classification tasks (ECtHR, SCOTUS).

(512 sub-word tokens in BERT, RoBERTa), which are often too restrictive in the legal domain, where longer documents are common. The same restrictions apply to LegalBERT (Chalkidis et al., 2020), a BERT variant pre-trained on legal corpora.

Even the sparse-attention Longformer (Beltagy et al., 2020), a well-known Transformer that increases the maximum input to 4,096 sub-words, severely truncates texts in three of the six datasets (see Fig. 2) of the LexGLUE legal NLP benchmark (Chalkidis et al., 2022). On the other hand, simpler linear classifiers with TF-IDF features (Manning et al., 2008), which were very common before deep learning, can handle texts of any length, at least in text classification tasks, require far less resources to train and deploy, but are nowadays usually outperformed by pre-trained Transformers.

Motivated by these observations, we explore two directions to better cope with long legal texts: (i) we modify a Longformer warm-started from LegalBERT to handle even longer texts (up to 8,192 sub-words), a resource-intensive direction that further increases the parameters and processing time

of large sparse-attention Transformer models; and (ii) we modify LegalBERT to use TF-IDF representations, which allows processing longer texts without increasing the model sizes. The first approach is the best overall in terms of performance, surpassing a hierarchical version of LegalBERT (Chalkidis et al., 2021a), which was the previous state of the art in LexGLUE. The second direction leads to computationally more efficient models at the expense of lower performance, but still outperforms overall a linear Support Vector Machine (SVM) (Cortes and Vapnik, 1995) with TF-IDF features in long document classification.

2 Related Work

2.1 Long Document Processing

Transformer-based models consist of stacked Transformer blocks (Vaswani et al., 2017). Each block builds a revised embedding (vector representation) for each (sub-word) token of the input text, based on the embeddings of the previous block, starting from an initial embedding layer that provides an embedding per vocabulary token. For a block with a single attention head and an input n tokens long, generating a single revised token embedding involves computing a weighted sum (weighted by attention scores) over the n token embeddings of the previous block. Hence, $O(n^2)$ time is required to generate all the n revised token embeddings. With k attention heads, the complexity is $O(k \cdot n^2)$.

Sparse-attention variants of Transformers, like those used in Longformer (Beltagy et al., 2020), ETC (Ainslie et al., 2020), BigBird (Zaheer et al., 2020), generate each revised token embedding by attending (considering) only the previous block’s embeddings for the current, the l previous, and the l next tokens in the input text, i.e., the weighted sum is now over only $2 \cdot l + 1$ (equal to 512 by default) token embeddings of the previous block. The complexity becomes $O(k \cdot n \cdot l)$, linear to n . To better capture long-distance dependencies, these models also use *global attention*. This involves either standard pseudo-tokens, such as the [cls] token at the beginning of each text, or additional pseudo-tokens, e.g., [sep] tokens placed at the end of each paragraph. In both cases, these special global tokens are attended by, and attend all other tokens, allowing information to flow across distant tokens, even when sparse attention is used.

We experiment with Longformer, a well-known and relatively simple sparse-attention Transformer,

which can process texts up to 4,096 sub-words long. ETC and Big Bird use the same maximum input length and are very similar; one difference is that they employ additional pre-training objectives for the global tokens, whereas in Longformer the global tokens are not pre-trained. We also expand Longformer to process texts up to 8,192 sub-words long, and we consider an ETC-like global attention scheme with additional [sep] tokens.

Hierarchical Transformers, e.g., SMITH (Yang et al., 2020), use BERT (or other base models) to separately encode each paragraph or other segments of the input text that do not exceed the base model’s maximum input length. The generated paragraph embeddings (e.g., the embeddings of [cls] tokens placed at the beginning of each paragraph) are then passed through additional stacked Transformer blocks, to allow interactions between the paragraph embeddings. The resulting context-aware paragraph embeddings can then be used to classify individual paragraphs or to classify the entire text (e.g., using the first paragraph embedding or by max-pooling over all paragraph embeddings).

In LexGLUE (Chalkidis et al., 2022) a similar hierarchical model (Chalkidis et al., 2021a) was used, with either generic BERT variants (e.g., BERT, RoBERTa, DeBERTa) or LegalBERT as the base model, in three of the benchmark’s tasks (ECtHR Task A and B, SCOTUS) where the average text length was much higher than BERT’s maximum length (Fig. 2). Unlike SMITH, the additional paragraph-level Transformer blocks were not pre-trained. We compare against this hierarchical variant of LegalBERT on LexGLUE.

Recurrent Transformers are another approach to handle long texts (Dai et al., 2019; Yang et al., 2019; Ding et al., 2021). We do not consider them here due to the latency that recurrency introduces.

Bag-of-Word (BoW) models typically represent each text as a (sparse) feature vector $\langle f_1, \dots, f_{|V|} \rangle$, with one feature f_i per vocabulary word. TF-IDF features (Manning et al., 2008) are common. Given a text n words long, each feature f_i becomes:

$$f_i = TF_i \cdot IDF_i = \frac{c_i}{n} \cdot \log \frac{N}{1 + d_i},$$

where c_i is the frequency of the i -th vocabulary word in the text, N is the number of documents in a corpus (in text classification this is often the training set), and d_i counts the documents of the corpus




Figure 2: Distribution of input text length, measured in BERT sub-word tokens, across the six LexGLUE datasets. Copied with permission from Chalkidis et al. (2022).

that contain the i -th vocabulary word.¹ Averaging word embeddings (Jin et al., 2016; Brokos et al., 2016) with or without TF-IDF weighting, is also a BoW representation, but typically performs worse, since averaging leads to very noisy representations.

Such BoW representations discard word order, but are also insensitive to the length of the input text, in the sense that the feature vector always contains $|V|$ features. Combining TF-IDF feature vectors with linear classifiers leads to models that can handle texts of any length and require far less resources to train compared to modern Transformer-based models, at the expense of lower performance.

BoW-BERT: Our attempts to combine TF-IDF features with BERT were inspired by the work of Hessel and Schofield (2021), who reported that shuffling the words of each text during fine-tuning led to a degradation of less than 5 p.p. (F1 or accuracy) of BERT’s performance in most GLUE tasks (Wang et al., 2018).² The resulting model, called BoW-BERT, can be seen as operating on BoW representations, in the sense that word order is lost. Hessel and Schofield (2021) also reported that BoW-BERT performed better than other BoW models on GLUE, including linear models with TF-IDF features. BoW-BERT’s word shuffling, however, does not change the text length, hence it does not address BERT’s maximum input length limit; IDF information is also not considered.

By contrast, we remove multiple occurrences of the same word from each text; to incorporate TF-IDF information, we order the remaining words by TF-IDF and/or we add a TF-IDF embedding layer, both discussed below.

¹In ‘sublinear’ TF-IDF, a logarithm is also applied to TF.

²However, Hessel and Schofield (2021) also cite other work that found word shuffling to have a larger impact on pre-trained Transformers (and LSTMs) in other datasets and tasks. Sinha et al. (2021) and Abdou et al. (2022) investigated the effect of word shuffling on pre-trained Transformers in more detail, considering mostly word shuffling during pre-training.

2.2 Applications in Legal NLP

In the early days of Deep Learning for legal NLP, the community examined the use of the Hierarchical Attention Network (HAN) of Yang et al. (2016) or simpler variants (hierarchical BILSTMs) to encode long documents in applications of legal judgment prediction for Chinese (Zhong et al., 2018) or ECtHR (Chalkidis et al., 2019a) court cases, showcasing improvement over flat RNN-based models, such as stacked BILSTMs followed by a single-head attention layer (Xu et al., 2015). Hierarchical BILSTMs with self-attention were also employed by Chalkidis et al. (2018) for sequential sentence classification in order to identify obligations and prohibitions in contractual paragraphs.

Hierarchical variants of Transformers were initially proposed by Chalkidis et al. (2019a). In their work, document paragraphs are encoded via a shared BERT encoder to produce paragraph embeddings, which are then combined with max-pooling to form the final document embedding. This model outperformed strong RNN-based methods such as the Hierarchical Attention Network (HAN).

Later on, Chalkidis et al. (2021a) presented a new variant, where the paragraph embeddings are fed into additional stacked Transformer blocks, to allow cross-paragraph contextualization. This latter version has also been used in other legal NLP applications, by Niklaus et al. (2021, 2022) in judgment prediction of Swiss court cases, using XLM-R as the underlying encoder, and by Chalkidis et al. (2022) for the long document classification tasks of LexGLUE, using several alternative pre-trained Transformers, alongside Longformer. Xiao et al. (2021) released a Longformer pre-trained on Chinese legal corpora, which outperforms baselines in several legal NLP tasks. More recently, Dai et al. (2022) explored how tunable hyper-parameters of Hierarchical Transformers and Longformer, such

as the size of the local window, affect downstream performance. In experiments on the ECtHR dataset, they found that fewer but larger local windows (paragraphs), e.g., 8×512 , instead of 32×128 , in Hierarchical Transformers improve performance.

Hierarchical Transformers are also used in the work of Malik et al. (2021) in legal judgment prediction of Indian court cases, where their best-performing model uses XLNet (Yang et al., 2019) as the underlying paragraph encoder followed by stacked BiGRUs. Moreover, hierarchical Transformers similar to those of Malik et al. have been also used by Kalamkar et al. (2022) for sequential legal sentence classification in order to segment Indian court cases into topical and coherent parts.

3 Models Considered

We discuss models Chalkidis et al. (2022) evaluated on LexGLUE as baselines, and models we introduce. The LexGLUE baselines also included RoBERTa (Liu et al., 2019), DeBERTa (He et al., 2021), BigBird (Zaheer et al., 2020), and CaseLaw-BERT (Zheng et al., 2021), which are not considered here. Chalkidis et al. (2022) found RoBERTa and DeBERTa to be better than BERT on LexGLUE, but worse than LegalBERT; no legally pre-trained variants of RoBERTa and DeBERTa are available. BigBird and CaseLaw-BERT were found to be overall slightly worse than Longformer and LegalBERT, respectively, on LexGLUE.

3.1 LexGLUE baselines

TFIDF-SVM is a linear SVM with TF-IDF features for the top- K most frequent word n -grams of the training set, where $n \in [1, 2, 3]$.³

LegalBERT (Chalkidis et al., 2020) is BERT pre-trained on English legal corpora (legislation, contracts, court cases). In the long document classification tasks (see Table 1), we deploy its hierarchical variant (Section 2) as in Chalkidis et al. (2022).

Longformer (Beltagy et al., 2020). This is the original Longformer, discussed in Section 2. It extends the maximum input length to 4,096 sub-word tokens. Like BERT and RoBERTa, Longformer uses absolute positional embeddings, i.e., there is a separate positional embedding for each token position up to the maximum input length. Longformer’s positional embeddings were warm-started from the 512 positional embeddings of RoBERTa, cloning

them 8 times (e.g., the embeddings of positions 513–1024 were initialized to the same RoBERTa positional embeddings as positions 1–512). All the other parameters of Longformer (and RoBERTa) are not sensitive to token positions and were warm-started from the corresponding RoBERTa parameters.⁴ After warm-starting, Longformer was further pre-trained for 64k steps on generic corpora.

3.2 Extensions of LegalBERT

TFIDF-SRT-LegalBERT: This is LegalBERT, but we remove duplicate sub-words from the input text and sort the remaining ones by decreasing TF-IDF during fine-tuning. Removing duplicate words is an attempt to avoid exceeding the maximum input length. In ECtHR, for example, the average text length (in sub-words) drops from 1,619 to 1,120; in SCOTUS, from 5,953 to 1,636 (see Fig. 1). If the new form of the text still exceeds the maximum input length, we truncate it (keeping the first 512 tokens). Ordering sub-words by decreasing TF-IDF hopefully allows the model to learn to attend earlier sub-words (higher TF-IDF) more, utilizing BERT’s positional embeddings as TF-IDF ranking encodings. This is a BoW model, since the original word order of the input text is lost.

TFIDF-SRT-EMB-LegalBERT: The same as the previous model, except that we add a TF-IDF embedding layer (Fig. 3). We bucketize the distribution of TF-IDF scores of the training set and assign a TF-IDF embedding to each bucket. During fine-tuning, we compute the TF-IDF score of each sub-word (before deduplication) and we add the corresponding TF-IDF bucket embedding to each token’s input embedding when its positional embedding is also added. The TF-IDF bucket embeddings are initialized randomly and trained during fine-tuning. Hence, this model is informed both about TF-IDF token *ranking* (via word re-ordering) and TF-IDF *scores* (captured by TF-IDF embeddings). This is still a BoW model, since it ignores the original word order, like the previous model.

TFIDF-EMB-LegalBERT: The same as LegalBERT, but we add the TF-IDF layer of the previous model. Token deduplication and ordering by TF-IDF scores are not included. This allows us to study the contribution of the TF-IDF layer on its own by comparing to the original LegalBERT. The resulting model is aware of word-order via its positional

³ $K \in [20k, 30k, 40k]$ is tuned per task on dev. data.

⁴E.g., the dense layers that produce the attention’s query, key, value embeddings are the same for all token positions.

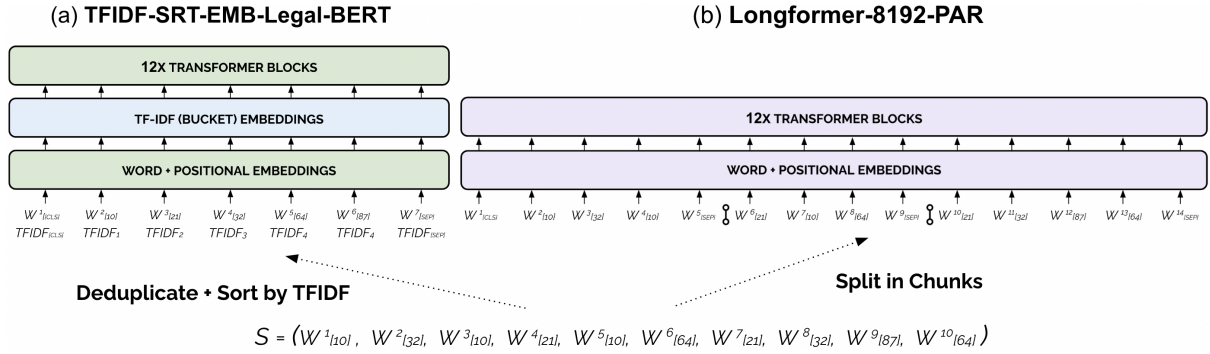


Figure 3: (a) TFIDF-SRT-EMB-LegalBERT, with sub-word token deduplication, re-ordering by TF-IDF, and TF-IDF embedding layer. (b) Longformer-8192-PAR extended to encode up to 8192 sub-word tokens, split into paragraphs separated by [sep] tokens. The original sequence (S) of sub-words (W) is shown in the bottom. Superscripts (W^p) denote positioning in each sequence. Subscripts ($W_{[id]}$) are the indices of the sub-words in the model’s vocabulary. In both models, the resulting contextualized [c1s] token embedding is fed to a linear classifier.

embeddings (like BERT and LegalBERT). For long texts, it addresses the maximum input length limitation via its hierarchical variant, which is similar to LegalBERT’s (Chalkidis et al., 2022).

3.3 Extensions of Longformer

Longformer-8192: This is the same as the original Longformer (Beltagy et al., 2020), which was warm-started from RoBERTa (Section 3.1), but we extend the maximum input length to 8,192 sub-words. We warm-start the positional embeddings from those of Longformer, cloning them once (positions 4,097–8,192 get the same initial embeddings as positions 1–4,096). To keep the computational complexity under control, we decrease the local attention window size from 512 to 128 sub-words.⁵ All parameters, including positional embeddings, are updated during fine-tuning, again as in the original Longformer. We did not perform any additional pre-training, however, beyond that of the original Longformer, lacking computing resources. All Longformer variants are aware of word order.

Longformer-8192-PAR: This is the same as the previous model, but we place a global token (Section 2), specifically a [sep] token, at the end of each paragraph (Fig. 3). By contrast, the original Longformer and Longformer-8192 use the single [c1s] token at the beginning of the input text as a single global token for classification tasks.⁶ As in the previous model, we decrease the local attention

window size from 512 to 128 sub-words.

Our intuition was that using more global tokens, and synchronizing them with paragraph breaks would allow information to flow more easily across paragraphs, viewed as discourse segments. Previous work by Zaheer et al. (2020) also suggests that such ETC-like global attention layouts lead to better results. Again, all parameters are updated during fine-tuning, but we did not perform any additional pre-training to better adjust the model to the new global attention layout, lacking resources.

LegalLongformer: Similar to Longformer, but warm-started from LegalBERT. We clone the positional embeddings of LegalBERT eight times to cover positions 1–4,096 (instead of 1–512 in LegalBERT) and update them during fine-tuning. All other parameters are also warm-started from LegalBERT and are updated during fine-tuning. Following Beltagy et al. (2020), we warm-start the global attention parameters of LegalLongformer with the (local) attention parameters of LegalBERT. Again, no additional pre-training was performed.

LegalLongformer-8192: Similar to Longformer-8192, but again warm-started from LegalBERT. In this case, we clone the positional embeddings of LegalBERT 16 times to cover positions 1–8,192. Again, no additional pre-training was performed.

LegalLongformer-8192-PAR: The same as the previous model, but with global tokens at the end of each paragraph, as in Longformer-8192-PAR.

⁵Table 4 shows that despite this counter-measure, the expansion to 8,192 sub-words leads to almost 2× inference time and 30% increase in memory.

⁶Additional global tokens were used by Beltagy et al. (2020) in other tasks, e.g., question answering.

Dataset	Source	Text Length (Original/Unique)		Instances			Classes
		Average	Maximum	Train	Dev.	Test	
LONG DOCUMENT CLASSIFICATION TASKS							
ECtHR (Task A)	(Chalkidis et al., 2019a)	1.6k (1.1k)	35.4k (27.2k)	9,000	1,000	1,000	10+1
ECtHR (Task B)	(Chalkidis et al., 2021b)	1.6k (1.1k)	35.4k (27.2k)	9,000	1,000	1,000	10+1
SCOTUS	(Spaeth et al., 2020)	6.0k (1.6k)	88.6k (12.1k)	5,000	1,400	1,400	14
SHORT DOCUMENT CLASSIFICATION TASKS							
EUR-LEX*	(Chalkidis et al., 2021a)	1.1k (341)	140.1k (10k)	55,000	5,000	5,000	100
LEDGAR	(Tuggener et al., 2020)	113 (65)	1.2k (484)	60,000	10,000	10,000	100
UNFAIR-ToS	(Lippi et al., 2019)	33 (25)	441 (181)	5,532	2,275	1,607	8+1

Table 1: LexGLUE statistics. Lengths in sub-word tokens, before and after (in brackets) deduplication. *EUR-LEX is treated as a short document task in our work, since using only the first 512 tokens (what we do) has comparable performance with using the full texts (Chalkidis et al., 2019b). +1 denotes an extra class for no-label instances.

4 Experiments

4.1 Datasets

LexGLUE (Chalkidis et al., 2022) is a collection of six simplified English legal NLP datasets that are used to evaluate the performance of NLP methods across seven legal text understanding tasks. Inspired by GLUE (Wang et al., 2018) and SuperGLUE (Wang et al., 2019), LexGLUE was designed to push towards generic-pretrained models that can cope with multiple legal NLP tasks with limited extra training (fine tuning) for each one.

Here, we experiment with six of the seven tasks of LexGLUE, excluding CaseHOLD (Zheng et al., 2021), a multiple choice question answering task about holdings of US court cases. The other six tasks are all framed as text classification problems. While our work targets the long document classification tasks (ECtHR Tasks A and B, SCOTUS), we also experiment with tasks that involve short texts (EUR-LEX, LEDGAR, UNFAIR-ToS), for completeness. Table 1 lists the sources of the datasets we experiment with and provides key statistics. ECtHR Task A and B require deciding which articles of the European Convention of Human Rights were violated, or allegedly violated, respectively; both tasks use the same dataset in LexGLUE. SCOTUS requires classifying opinions of the US Supreme Court into issue areas (e.g., Criminal Procedure, Civil Rights). EUR-LEX requires labeling European laws with concepts from a European Union taxonomy. LEDGAR requires assigning topical categories to contract provisions. UNFAIR-ToS requires detecting unfair terms in terms of service. Consult Chalkidis et al. (2022) and the work cited in Table 1 for further information.

4.2 Evaluation measures

Following Chalkidis et al. (2022), for each task we report macro-F1 ($m-F_1$), which assigns equal importance to all classes, and micro-F1 ($\mu-F_1$), which assigns more importance to frequent classes.

4.3 Experimental setup

Across all experiments, we use Adam (Kingma and Ba, 2015) with initial learning rate $3e-5$. We train models up to 20 epochs using early stopping, monitoring $\mu-F_1$ on the development data. We run all experiments with 5 different random seeds and report test results for the seeds with the best development scores. For the TF-IDF bucket embedding layer, we search in $\{16, 32, 64, 128\}$ for the number of buckets that maximizes $\mu-F_1$ on the development data, separately for each task.

4.4 Experimental results

Table 2 lists the test results of all models across the six tasks considered. Table 3 aggregates the test results over the three long-document classification tasks (ECtHR Tasks A and B, SCOTUS) we are mainly interested in (see also see Table 1). We use the harmonic mean over the scores of the three tasks, following Shavrina and Malykh (2021).

BoW models: The results of the two BoW variants of LegalBERT (TFIDF-SRT-LegalBERT, TFIDF-SRT-EMB-LegalBERT) in Table 2 are mixed. In the two ECtHR tasks and EUR-LEX, both models outperform the TFIDF-SVM baseline, a much simpler linear BoW model. Contrary, both models are outperformed by TFIDF-SVM in SCOTUS and LEDGAR. In UNFAIR-ToS, the three models perform overall on par. While the original word

Method	ECtHR (A)*		ECtHR (B)*		SCOTUS*		EUR-LEX		LEDGAR		UNFAIR-ToS	
	μ -F ₁	m-F ₁	μ -F ₁	m-F ₁	μ -F ₁	m-F ₁	μ -F ₁	m-F ₁	μ -F ₁	m-F ₁	μ -F ₁	m-F ₁
BoW models (word order lost)												
TFIDF-SVM	62.6	48.9	73.0	63.8	<u>74.0</u>	<u>64.4</u>	63.4	47.9	<u>87.0</u>	<u>81.4</u>	94.7	75.0
TFIDF-SRT-LegalBERT	<u>69.8</u>	62.8	78.5	71.9	73.4	61.8	69.6	53.7	86.9	80.8	95.3	<u>80.6</u>
TFIDF-SRT-EMB-LegalBERT	68.7	<u>63.1</u>	<u>79.0</u>	<u>72.5</u>	73.9	63.6	<u>69.7</u>	<u>53.9</u>	86.5	80.3	<u>95.8</u>	78.7
LegalBERT variants that retain word order												
LegalBERT	<u>70.0</u>	<u>64.0</u>	<u>80.4</u>	<u>74.7</u>	<u>76.4</u>	<u>66.5</u>	<u>72.1</u>	<u>57.4</u>	88.2	83.0	96.0	83.0
TFIDF-EMB-LegalBERT	<u>70.0</u>	61.9	79.4	73.5	74.9	64.7	71.6	56.9	<u>88.7</u>	<u>83.4</u>	95.9	82.1
Longformer variants (all retain word order)												
Longformer	69.9	64.7	79.4	71.7	72.9	64.0	71.6	<u>57.7</u>	88.2	83.0	95.5	<u>80.9</u>
Longformer-8192	70.9	62.1	79.2	73.9	73.7	63.6	(Not considered for short-document tasks.)					
Longformer-8192-PAR	70.8	62.3	79.0	73.1	73.9	<u>66.0</u>	>>					
LegalLongformer	<u>71.7</u>	63.6	80.5	76.4	76.6	66.9	<u>72.2</u>	56.5	88.8	83.5	<u>95.7</u>	80.6
LegalLongformer-8192	71.2	64.3	81.4	74.2	77.5	67.3	(Not considered for short-document tasks.)					
LegalLongformer-8192-PAR	71.4	68.4	79.6	73.9	76.2	66.3	>>					

Table 2: Test results across LexGLUE tasks considered. In starred tasks, we use the hierarchical variant of LegalBERT. We do not consider extended Longformers in short document classification tasks (last three; see also Table 1), which are included for completeness. Best scores per group are underlined, and best overall are in **bold**.

order is lost in all three models, TFIDF-SVM relies on n -grams up to 3 words long, which allows it to retain local word order in features that represent multi-word terms, like ‘civil rights’ or ‘federal taxation’ in the case of SCOTUS. We suspect that such multi-word terms are more important in SCOTUS and LEDGAR, which would explain the fact that TFIDF-SVM outperforms the other two BoW models in these tasks. Future work could add a TFIDF-SVM variant with only unigram features to check this hypothesis; there should be a large performance drop in the three tasks. One could also explore ways to use TF-IDF information about n -grams (not just unigrams) in the BoW variants of LegalBERT.

Switching to the aggregated results of the long document tasks of Table 3, we observe that both BoW variants of LegalBERT outperform TFIDF-SVM. Table 3 also shows that TFIDF-SRT-EMB-LegalBERT (which includes the TF-IDF embeddings layer) performs slightly better than TFIDF-SRT-LegalBERT in terms of m-F₁ (1 p.p. improvement), but there is almost no difference in μ -F₁, and the results of Table 2 show no clear winner between the two methods across tasks.

LegalBERT variants that retain word order:

Table 2 shows that adding the TF-IDF embeddings layer to LegalBERT (TF-IDF-EMB-LegalBERT), without word deduplication and retaining the original word order, leads to lower performance in 5

Method	μ -F ₁	m-F ₁
TFIDF-SVM	69.5	58.1
TFIDF-SRT-LegalBERT	<u>73.7</u>	65.2
TFIDF-SRT-EMB-LegalBERT	73.6	<u>66.1</u>
LegalBERT	<u>75.4</u>	<u>68.1</u>
TFIDF-EMB-LegalBERT	74.6	66.3
Longformer	73.9	66.6
Longformer-8192	74.4	66.1
Longformer-8192-PAR	74.4	66.8
LegalLongformer	76.1	68.6
LegalLongformer-8192	76.5	68.4
LegalLongformer-8192-PAR	75.6	69.4

Table 3: Test results aggregated (harmonic mean) over the long-document classification tasks (ECtHR Tasks A and B, SCOTUS) of LexGLUE. Best scores per group are underlined, and best overall are in **bold**.

out of 6 tasks compared to the original LegalBERT; LEDGAR is the only exception, with small improvements. The aggregated results of Table 3 also show that TF-IDF-EMB-LegalBERT is worse than the original LegalBERT. We can only hypothesize that TF-IDF-EMB-LegalBERT is in most cases unable to learn how to use the additional information from the additive TF-IDF embeddings, which are added only during fine-tuning (they were not present during pre-training). This hypothesis

is based on the positive (albeit small) impact of the TF-IDF embeddings layer on LEDGAR, the largest dataset with 60k training examples. All other datasets contain fewer than 10k training examples (Table 1), with the exception of EUR-LEX (55k), which does not support our hypothesis.

Given appropriate computing resources, one could further pre-train TFIDF-EMB-LegalBERT to help it learn how to exploit the newly introduced TF-IDF embeddings. The same applies to both BoW variants of LegalBERT, although in that case appropriate BoW pre-training objectives should be considered, since Masked Language Modeling (MLM) is not reasonable when the original word order is lost. Predicting the TF-IDF bucket id when masked, or predicting masked words given their TF-IDF bucket ids seem better alternatives.

Longformer variants: Comparing the original Longformer with Longformer-8192, a variant capable of processing even longer documents, the results are mixed (Table 2) across the 3 long document classification tasks (ECtHR Tasks A and B, SCOTUS), i.e., μ -F₁ is improved at the expense of m-F₁, or vice-versa. Aggregating the results (Table 3), we observe the very same trade-off (+0.5 p.p. in μ -F₁, -0.5 p.p. in m-F₁). Considering the additional global tokens in Longformer-8192-PAR, we have comparable results in ECtHR tasks and improved results in SCOTUS, the dataset with the longest documents in LexGLUE (Table 1). Aggregating the results (Table 3), we observe that the extra global tokens do not improve μ -F₁ further (74.4), but lead to the best m-F₁ (66.8) of all the Longformer variants that have not been pre-trained on legal corpora. Based on the aforementioned observations, we believe that the additional positional embeddings and adding more global tokens are in the right direction when seeking better long document performance with Longformer.

Moving on to Longformer variants warm-started from LegalBERT, Table 2 shows that LegalLongformer outperforms the original generic Longformer (Beltagy et al., 2020) in most cases, which highlights the importance of domain-specific models as already noted in the literature (Chalkidis et al., 2022; Zheng et al., 2021). We observe notable improvements in long document classification tasks (ECtHR A and B, SCOTUS), with approx. +2.0 p.p. in both μ -F₁ and m-F₁ in the aggregated results of Table 3. These results are impressive considering that LegalLongformer was warm-started

from LegalBERT, but no additional pre-training was conducted; hence several parameters of the model (e.g., additional positional embeddings and global attention matrices) may be far from optimal. By contrast, the original Longformer was warm-started from RoBERTa and was pre-trained for 64k additional steps on generic long documents.

Considering the last two variants of LegalLongformer (-8192, -8192-PAR), the results are mixed (trade-off between μ -F₁ and m-F₁ in Table 2, as with the generic Longformer) and share the best aggregated results across all examined methods in long document classifications tasks (Table 3).

Based on the above, we believe that the proposed extensions (warm-start from a legally pre-trained model, additional positional embeddings, additional global tokens) are in the right direction, already producing better results compared to the generic Longformer, and state-of-the-art results in several LexGLUE tasks (ECtHR A&B and LEDGAR). Given appropriate resources, one could further pre-train LegalLongformer-8192-PAR for a limited number of steps (e.g., 64k) on long legal documents (e.g., the training subsets of ECtHR, and SCOTUS) to optimize the newly introduced parameters and expect further improvements.

4.5 Efficiency considerations

In Table 4, we present important information with respect to efficiency. As expected, TFIDF-SVM has the fewest parameters (200× fewer than LegalBERT variants) and is substantially faster and less memory-intensive compared to all other neural methods, while achieving state-of-the-art results in two tasks (SCOTUS and EUR-LEX, Table 2).

Our proposed BoW variants of LegalBERT are substantially less memory intensive; approx. 25% less GPU memory across the long document classification tasks (starred), and approx. 50% less GPU memory across others with much shortened texts compared to LegalBERT. The TF-IDF embeddings do not affect memory or inference time (storing and looking up TF-IDF embeddings are negligible).

Considering LegalLongformer, we observe an approx. 50% increase in the number of parameters and approx. 25% increase in GPU memory. With respect to inference time, there is a 10× increase compared to LegalBERT models in long document processing tasks, and larger in the other tasks with much shorter documents, which makes hierarchical Transformers a faster alternative.

Method	Params.	ECtHR*		SCOTUS*		EUR-LEX		LEDGAR		UNFAIR-ToS	
		Mem.	Time	Mem.	Time	Mem.	Time	Mem.	Time	Mem.	Time
BoW models (word order lost)											
TFIDF-SVM	0.5M	0.1	.001	0.1	.001	0.1	.001	0.1	.001	0.1	.001
TFIDF-SRT-LegalBERT	110M	0.9	.012	0.9	.012	0.9	.012	0.9	.007	0.9	.007
TFIDF-SRT-EMB-LegalBERT	110M	0.9	.012	0.9	.012	0.9	.012	0.9	.007	0.9	.007
LegalBERT variants that retain word order											
LegalBERT	110M	1.3	.014	1.3	.014	1.9	.012	1.9	.007	1.9	.007
TFIDF-EMB-LegalBERT	110M	1.3	.014	1.3	.014	1.9	.012	1.9	.007	1.9	.007
Longformer variants (all retain word order)											
LegalLongformer	148M	1.7	.164	1.7	.164	1.3	.033	1.3	.033	1.3	.033
LegalLongformer-8192	151M	2.2	.318	2.2	.318	(Not considered for short-document tasks.)					
LegalLongformer-8192-PAR	151M	2.2	.331	2.2	.331	>>					

Table 4: Model parameters, memory footprint (GBs/sample), and inference time (sec/sample). In starred tasks, we use the hierarchical variant of LegalBERT. For ECtHR Tasks A and B, the information of this table is identical.

Moving to the extensions of LegalLongformer that are able to encode longer documents (LegalLongformer8192) and use extra global tokens (LegalLongformer-8192-PAR), there is an approx. 30% increase in GPU memory compared to the standard Longformer (encoding up to 4,096 sub-words), and 2× increase in inference time. In other words, there is no free lunch when seeking performance improvements.

5 Conclusions and Future Work

Concluding, we presented BoW variants of LegalBERT, which remove duplicate words and consider TF-IDF scores by reordering the remaining words and/or by employing a TF-IDF embedding layer. These variants are more efficient than the original LegalBERT and still overall outperform a TF-IDF-based SVM in long legal document classification.

We also modified Longformer to handle even longer texts (up to 8,192 sub-words), use additional global tokens, and also showed the positive effect of warm-starting it from LegalBERT. Unlike the BoW models, this is a resource-intensive direction, with substantial improvements compared to the original Longformer (up to 4,096 sub-words, a single global token, warm-started from RoBERTa) in long legal document classification. The new LegalLongformer (and its variants) are the new state of the art in the long document tasks of LexGLUE.

In future work, we would like to further pre-train the proposed BoW variants of LegalBERT, LegalLongformer, and variants on legal corpora, to help them better optimize the newly introduced modifications (e.g., TF-IDF embeddings, additional posi-

tional embeddings, updated attention scheme with additional global tokens). We would also like to experiment with long documents from other domains (e.g., long business documents).

Acknowledgments

This research has been co-funded by the European Regional Development Fund of the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH – CREATE – INNOVATE (T2EAK-03849). This work is also partly funded by the Innovation Fund Denmark (IFD)⁷ under File No. 0175-00011A.

References

- Mostafa Abdou, Vinit Ravishankar, Artur Kulmizev, and Anders Søgaard. 2022. Word order does matter and shuffled language models know it. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6907–6919, Dublin, Ireland.
- Joshua Ainslie, Santiago Ontanon, Chris Alberti, Václav Cvicek, Zachary Fisher, Philip Pham, Anirudh Ravula, Sumit Sanghai, Qifan Wang, and Li Yang. 2020. [ETC: Encoding long and structured inputs in transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 268–284, Online.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *CoRR*, abs/2004.05150.

⁷<https://innovationsfonden.dk/en>

- Georgios-Ioannis Brokos, Prodromos Malakasiotis, and Ion Androutsopoulos. 2016. [Using centroids of word embeddings and word mover’s distance for biomedical document retrieval in question answering](#). In *Proceedings of the 15th Workshop on Biomedical Natural Language Processing*, pages 114–118, Berlin, Germany. Association for Computational Linguistics.
- Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019a. [Neural legal judgment prediction in English](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4317–4323, Florence, Italy. Association for Computational Linguistics.
- Ilias Chalkidis, Ion Androutsopoulos, and Achilleas Michos. 2018. [Obligation and prohibition extraction using hierarchical RNNs](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 254–259, Melbourne, Australia. Association for Computational Linguistics.
- Ilias Chalkidis, Emmanouil Fergadiotis, Prodromos Malakasiotis, and Ion Androutsopoulos. 2019b. [Large-scale multi-label text classification on EU legislation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6314–6322, Florence, Italy. Association for Computational Linguistics.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. [LEGAL-BERT: The muppets straight out of law school](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2898–2904, Online.
- Ilias Chalkidis, Manos Fergadiotis, Dimitrios Tsarapat-sanis, Nikolaos Aletras, Ion Androutsopoulos, and Prodromos Malakasiotis. 2021a. [Paragraph-level rationale extraction through regularization: A case study on European court of human rights cases](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 226–241, Online. Association for Computational Linguistics.
- Ilias Chalkidis, Manos Fergadiotis, Dimitrios Tsarapat-sanis, Nikolaos Aletras, Ion Androutsopoulos, and Prodromos Malakasiotis. 2021b. [Paragraph-level rationale extraction through regularization: A case study on european court of human rights cases](#). In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics*, online.
- Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Martin Katz, and Nikolaos Aletras. 2022. [Lexglue: A benchmark dataset for legal language understanding in english](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, pages 4310–4330, Dublin, Ireland.
- Corinna Cortes and Vladimir Vapnik. 1995. [Support-vector networks](#). *Maching Learning*, 20(3):273–297.
- Xiang Dai, Ilias Chalkidis, Sune Darkner, and Desmond Elliott. 2022. [Revisiting transformer-based models for long document classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, Abu Dhabi, UAE. Association for Computational Linguistics.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy.
- 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 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota.
- SiYu Ding, Junyuan Shang, Shuohuan Wang, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2021. [ERNIE-Doc: A retrospective long-document modeling Transformer](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2914–2927, Online.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [DeBERTa: Decoding-enhanced BERT with disentangled attention](#). In *International Conference on Learning Representations*.
- Jack Hessel and Alexandra Schofield. 2021. [How effective is BERT without word ordering? implications for language understanding and data privacy](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 204–211, Online.
- Peng Jin, Yue Zhang, Xingyuan Chen, and Yunqing Xia. 2016. [Bag-of-embeddings for text classification](#). In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJ-CAI’16*, page 2824–2830. AAAI Press.
- Prathamesh Kalamkar, Aman Tiwari, Astha Agarwal, Saurabh Karn, Smita Gupta, Vivek Raghavan, and Ashutosh Modi. 2022. [Corpus for automatic structuring of legal documents](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4420–4429, Marseille, France. European Language Resources Association.

- D. P. Kingma and J. Ba. 2015. [Adam: A method for stochastic optimization](#). In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, San Diego, CA, USA.
- Marco Lippi, Przemysław Pałka, Giuseppe Contissa, Francesca Lagioia, Hans-Wolfgang Micklitz, Giovanni Sartor, and Paolo Torroni. 2019. [CLAUDETTE: an automated detector of potentially unfair clauses in online terms of service](#). *Artificial Intelligence and Law*, pages 117–139.
- 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](#). *arXiv preprint arXiv:1907.11692*.
- Vijit Malik, Rishabh Sanjay, Shubham Kumar Nigam, Kripabandhu Ghosh, Shouvik Kumar Guha, Arnab Bhattacharya, and Ashutosh Modi. 2021. [ILDC for CJPE: Indian legal documents corpus for court judgment prediction and explanation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4046–4062, Online. Association for Computational Linguistics.
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press, USA.
- Joel Niklaus, Ilias Chalkidis, and Matthias Stürmer. 2021. [Swiss-Court-Predict: A Multilingual Legal Judgment Prediction Benchmark](#). In *Proceedings of the 3rd Natural Legal Language Processing Workshop Workshop*, Online.
- Joel Niklaus, Matthias Stürmer, and Ilias Chalkidis. 2022. [An empirical study on cross-x transfer for legal judgment prediction](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (ACL-IJCNLP 2022)*, Online. Association for Computational Linguistics.
- Tatiana Shavrina and Valentin Malykh. 2021. [How not to lie with a benchmark: Rearranging NLP leaderboards](#).
- Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela. 2021. [Masked language modeling and the distributional hypothesis: Order word matters pre-training for little](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2888–2913, Online and Punta Cana, Dominican Republic.
- Harold J. Spaeth, Lee Epstein, Jeffrey A. Segal Andrew D. Martin, Theodore J. Ruger, and Sara C. Benesh. 2020. [Supreme Court Database, Version 2020 Release 01](#). Washington University Law.
- Don Tuggener, Pius von Däniken, Thomas Peetz, and Mark Cieliebak. 2020. [LEDGAR: A large-scale multi-label corpus for text classification of legal provisions in contracts](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1235–1241, Marseille, France. European Language Resources Association.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6000–6010, Long Beach, California, USA.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium.
- 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*, 2:79–84.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. [Show, attend and tell: Neural image caption generation with visual attention](#). In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2048–2057, Lille, France. PMLR.
- Liu Yang, Mingyang Zhang, Cheng Li, Michael Bendersky, and Marc Najork. 2020. [Beyond 512 tokens: Siamese multi-depth transformer-based hierarchical encoder for long-form document matching](#). In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*, page 1725–1734, Virtual Event, Ireland.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhudinov, and Quoc V. Le. 2019. [XLNet: Generalized autoregressive pretraining for language understanding](#). In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Red Hook, NY, USA. Curran Associates Inc.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. [Hierarchical attention networks for document classification](#). In

Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 1480–1489, San Diego, California. Association for Computational Linguistics.

Manzil Zaheer, Guru Guruganesh, Kumar 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*, pages 17283–17297, online.

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. In *International Conference on Artificial Intelligence and Law*.

Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Chaojun Xiao, Zhiyuan Liu, and Maosong Sun. 2018. [Legal judgment prediction via topological learning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3540–3549, Brussels, Belgium. Association for Computational Linguistics.

A Additional Material

In Figures 4–5, we show boxplots of the average text length (in sub-words) across LexGLUE datasets before and after word deduplication.




Figure 4: Average text length (in sub-words) across LexGLUE datasets before word deduplication.




Figure 5: Average text length (in sub-words) across LexGLUE datasets after word deduplication.