

# DeepParliament: A Legal domain Benchmark & Dataset for Parliament Bills Prediction

Ankit Pal

Open Legal AI

openlegalai@gmail.com

## Abstract

This paper introduces DeepParliament, a legal domain Benchmark Dataset that gathers bill documents and metadata and performs various bill status classification tasks. The proposed dataset text covers a broad range of bills from 1986 to the present and contains richer information on parliament bill content. Data collection, detailed statistics and analyses are provided in the paper. Moreover, we experimented with different types of models ranging from RNN to pretrained and reported the results. We are proposing two new benchmarks: Binary and Multi-Class Bill Status classification. Models developed for bill documents and relevant supportive tasks may assist Members of Parliament (MPs), presidents, and other legal practitioners. It will help review or prioritise bills, thus speeding up the billing process, improving the quality of decisions and reducing the time consumption in both houses. Considering that the foundation of the country’s democracy is Parliament and state legislatures, we anticipate that our research will be an essential addition to the Legal NLP community. This work will be the first to present a Parliament bill prediction task. In order to improve the accessibility of legal AI resources and promote reproducibility, we have made our code and dataset publicly accessible at [github.com/monk1337/DeepParliament](https://github.com/monk1337/DeepParliament)

## 1 Introduction

In recent years, Artificial Intelligence(AI) based methods have been employed in the legal field for several uses and within many sub-areas. In the legal field the majority of resources are available in textual format (e.g., contracts, court decisions, patents, legal articles). Therefore considerable efforts have been made at the intersection of Law and Natural Language Processing research. Efforts can be witnessed in the various projects dealing with NLP applications in the legal domain and recently published scientific papers such as legal judgement prediction (Aletras et al., 2016; Xiao

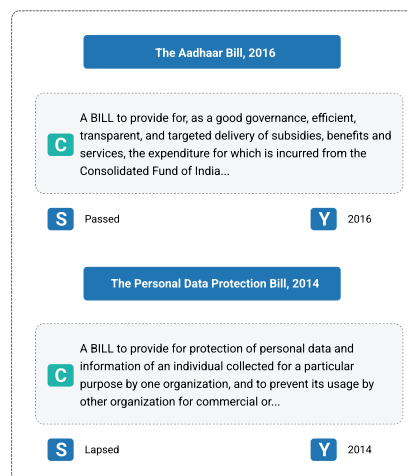


Figure 1: Samples from the DeepParliament dataset. Here C, S, and Y indicate Bill Context, Bill Status, and Bill Year, respectively.

et al., 2018; Chalkidis et al., 2019) legal topic classification (Nallapati and Manning, 2008; Chalkidis et al., 2020), overruling prediction (Zheng et al., 2021). Furthermore, researchers have also explored a variety of Legal AI tasks, including legal question answering (Kien et al., 2020), contract understanding (Hendrycks et al., 2021), court opinion generation (Ye et al., 2018), legal information extraction (Chalkidis et al., 2018), legal entity recognition (Leitner et al., 2019, 2020) and many more. Predictive legal models have the ability to enhance both the effectiveness of decision-making and the provision of services to individuals. Many new datasets have also been proposed in the legal domain to track the recent progress and serve as benchmarks. Recently, there have been initiatives to develop corpora for the India’s judicial system.

The foundation of India’s democracy is the Parliament and state legislatures. Implementing, amending and removing laws is the primary responsibility of Parliament. The Rajya Sabha (Council of States) and the Lok Sabha (House of the People) are the two houses that constitute India’s legislature.

However, the majority of contemporary legal NLP research focuses only on court decisions and cases. This issue, which we refer to and subsequently characterize as "Parliament Bills Prediction", has not been explored. Before qualifying as an act, every bill passes through a long chain of standardized processes, including introduction of the bill, publication in the gazette, first reading, select committee, second reading, and third reading. After the third reading, the bill goes to the other houses, and after the approval of both houses it faces final approval by the president. A significant amount of time and effort is required to pass a bill in either of the Houses of Parliament. Therefore, a lapse of a bill has a negative impact on legislative work.

India's first Lok Sabha (1952–1957) passed 333 bills throughout its five-year existence. The average number of bills approved by Lok Sabhas with terms less than three years is 77. Both houses spent about half their time carrying out legislative business. The lapse count of the bill increases at the end of every Lok Sabha. A total of 22 bills lapsed after the 16th Lok Sabha; three bills have been pending for over 20 years; six have been pending between 10-20 years. Legislative activity accounts for a significant portion of Lok Sabha's working hours. To date, 14 bills are still pending between 5-10 years and 10 bills are pending for under five years. Therefore, time is wasted when bills lapse at the end of the Lok Sabha's tenure, as a new Lok Sabha must start over and consider bills from scratch, taking at least two sessions to reconsider the bills. Thus, in order to improve productivity, it is necessary to re-evaluate the rule governing the lapsing of bills in the House of Representatives. This calls for machine learning strategies to enhance the efficiency of the billing process in Parliament.

To the best of our knowledge, a single dataset does not yet exist, which provides a standard benchmark for parliamentary Bills. To facilitate research on bill documents for text classification, we provide DeepParliament, a legal domain Benchmark & Dataset which gathers bill documents and meta data and performs different status classification tasks. The proposed dataset and benchmark are not meant to replace or compete with the decisions of the Houses of Parliament and the President by any means; instead, the proposed solution offers complementing use cases. Models developed for bill documents and relevant supportive tasks may assist members of the Legislative Assembly (MLA),

Members of Parliament (MPs), presidents and other legal practitioners, for example by estimating the likelihood of getting a bill passed, reviewing or prioritizing bills, (thus speeding up the billing process); improving the quality of decisions and finally, reducing the time and energy consumption in both houses. Fig. 1 shows two samples of two parliament bills' context, their corresponding status, and the year from the study dataset.

Applications developed on this dataset, such as automatic summaries, would enable the professionals to decide which documents they should read in detail. Moreover, the model can suggest different sections and acts in need of further exploration by highlighting which areas a new bill falls within. This paper proposes a benchmark and takes initial steps by contributing the dataset and baseline models to the community. Moreover, the plan is to continue to revise and upgrade the DeepParliament dataset in the future.

In brief, the contributions of this study are as follows.

- **New Legal Dataset.** We are proposing a new dataset. To our knowledge, there is no dataset focusing on parliament bills and data. Therefore, this work will be the first to present a parliament bill prediction task having rich information on parliament bills, different acts and laws.
- **Diversity and Difficulty.** The proposed dataset text covers a broad range of bills, including the Government Bill, Private Members Bill, the Money Bill, the Ordinary Bill, the Financial Bill & Constitutional Amendment Bill from 1986 to the present. Moreover, on average, there are 3932.99 tokens per sentence. The documents on the proposed dataset are considerably long. They contain richer information of parliament bill content, testing the reasoning abilities and domain-specific capabilities of language models in the legal domain.
- **Quality Detailed Statistics,** analysis of the dataset, and fine-grained evaluation of different parts of documents are provided. Moreover, we also performed extensive quality experiments to evaluate different types of models ranging from RNN to high-performance pre-trained domain models.




Figure 2: The Distribution of Top India’s Act/Law’s Section mentioned in the dataset.

- **Reproducible Results** We employ the HuggingFace Transformers library (Wolf et al., 2019) to facilitate our experiments. Furthermore, we pre-process and publish datasets on HuggingFace Datasets (Lhoest et al., 2021) to reproduce the results and experiment with new models in the future.
- **Proposing Two Benchmarks** We are proposing two new benchmarks on the DeepParliament dataset: Binary and Multi-Class Bill Status classification. In addition to the code, we also publish the benchmark on PaperwithCode<sup>1</sup> and Open LegalAI<sup>2</sup> to track the progress.

## 2 The DeepParliament Dataset

### 2.1 Task Definition

We model the bill prediction task as a classification problem and design Binary and Multi-Class Classification problem statements on the proposed dataset to evaluate the domain-specific capabilities of language models in the legal domain. For a given collection of labelled bill documents  $X$ , the objective is to learn a classification function:

$$f : x_i \rightarrow y_i \quad (1)$$

Where  $x_i$  is a legal bill document.

### 2.2 Task 1: Binary Classification

In equation (1)  $y_i \in \{0, 1\}$  is target binary label of corresponding status Passed, Failed of classification task.

### 2.3 Task 2: Multi-Class Classification

Task 2, the coarse-grained classification task, had a total of 5 classes. In equation (1)  $y_i \in \{1, \dots, K\}$  is the multi-class label of the corresponding status Passed, Negatived, Lapsed, Removed, Withdrawn.

<sup>1</sup><https://paperswithcode.com/dataset/deepparliament>

<sup>2</sup>[openlegalai.github.io/DeepParliament](https://openlegalai.github.io/DeepParliament)

## 2.4 Bill and Lawmaking Procedure

The foundation of India’s democracy is the parliament and state legislatures. Making, changing, and removing laws is the primary responsibility of Parliament. The method by which a legislative proposal is turned into an act is referred to as the legislative process or the lawmaking process in relation to Parliament. The procedure of a new act starts with identifying the need for a new law or an amendment to a current part of the legislation. Following the legal requirement, the relevant ministry writes a text for the proposed legislation, known as a Bill. Other relevant ministries are informed about this bill so that they can make any alterations or amendments.

A bill, which is draft legislation, cannot become law until it has been approved by both houses of Parliament and the President of India. Furthermore, the bill is introduced in Parliament after receiving cabinet approval. Prior to becoming an act, every bill passes through several readings in both houses. After both houses have approved a bill of Parliament, it is forwarded to the president for his or her approval. However, the president can request information and an explanation about the bill. The bill may be returned to Parliament for further consideration. The bill is declared an act with the president’s assent. Moreover, the bill is then made into law, and the responsible ministry draughts and submits to Parliament the rules and regulations necessary to carry out the Act.

## 3 Dataset Collection & Preprocessing

We constructed the DeepParliament corpus from raw data collected from the official<sup>3</sup> & open website<sup>4</sup> which put together all the parliament bills from 1986 to the present. In addition to the raw data, additional metadata is also provided, i.e. the title, type of bill such as government or private, source of the bill, pdf URL and status of the bill. We used pdfminer3<sup>5</sup> to extract bill content from each PDF. Some old pdfs are in image format; we applied an OCR system to convert images to text. The pdf content & metadata were converted into CSV format and combined into a single dataset. Next, we eliminate bill documents with a single token and duplicates. The cleaning pipeline involves removing the special characters, extra spaces etc.

<sup>3</sup><https://loksabhaph.nic.in>

<sup>4</sup><https://prsindia.org/billtrack>

<sup>5</sup><https://pypi.org/project/pdfminer3>

All bill documents in this dataset are specifically in the English language.

#### 4 Dataset statistics

The statistics of our proposed dataset, including the train and the test corpus, are shown in Table 1. Total documents are 5,329, where 4223 are in the train and 1106 are in the test dataset. Each bill document contains many sentences in both cases, and the document’s length varies greatly. The perfor-




Figure 3: Distribution & Cumulative Frequency Graph of Passed and Failed Bill status in the last 20 years.

mance of the models is influenced by the amount of vocabulary, which is a good indicator of the linguistic and domain complexity associated with a text corpus; Fig 5 shows the distribution of unique tokens of the train & test set. As shown in the table, this dataset has 284103 tokens in total, and the train split contains 243393 tokens where the test vocabulary size is 86616. On average, there are 3932.99 tokens per sentence.

#### 5 Dataset Analysis

The documents on the proposed dataset are considerable long and contain richer information on bill content. As described before, the dataset has been categorized into two settings: Binary Classification and Multi-Class classification. The most frequent category status is Lapsed, which occupies 50.6%. Fig 4 shows the percentage of each status type. Lapsed, Passed, and Withdrawal is the dataset’s top three common statuses. The proposed dataset text covers a broad range of bills, including Government Bill, Private Member Bill, Money




Figure 4: Relative sizes of documents per bill status in Dataset

Bill, Ordinary Bill, Financial Bill and Constitutional Amendment Bill. We used word cloud to visualize the top Commonly occurring legal words in the dataset shown in Fig 7.

We visualized the top sections mentioned in the entire dataset. A section is a specific provision of a legal code or body of laws, often laying out a specific legal obligation. Most sections in the dataset come under the Indian panel code and income tax act. Fig. 2 shows the top Indian Act/Law’s Section mentioned in the dataset. To understand the dataset better, we also visualize the bill status of the last top 20 years. Fig. 3 shows the visualization. We can see that year 2015 and 2016 has the most significant failure ratio in the last 20 years, while in 2019, most bills were passed compared to other years.

	Train	Test	Total
Documents #	4223	1106	5329
Vocab	243393	86616	284103
Max D tokens	219378	227407	227407
Max T tokens	36	36	36
Avg D tokens	3932.99	4080.97	3963.70
Avg T tokens	11.15	11.48	11.22

Table 1: DeepParliament dataset statistics, where D, T represents the Documents and Title, respectively

### 6 Methods

Our study considers ten text classification models ranging from long-range RNN and CNN to Transformer-based methods increasing recency and sophistication.

#### 6.1 Sequence & Convolutional models

In this category, we experimented with standard long short-term memory (LSTM) (Hochreiter and Schmidhuber, 1997), Bidirectional LSTM (BiLSTM) and convolutional neural network (CNN) (Kim, 2014) models for bill prediction tasks. We decided to use a shallow CNN model for glove since research has shown that deep CNN models do not consistently outperform other algorithms for text classification tasks. Initially, we utilized Xavier weight initialization (Glorot and Bengio, 2010) for both models’ embedding matrices. Later we leverage this by initializing word vectors using pre-trained GloVe embedding (Pennington et al., 2014) of length 300.

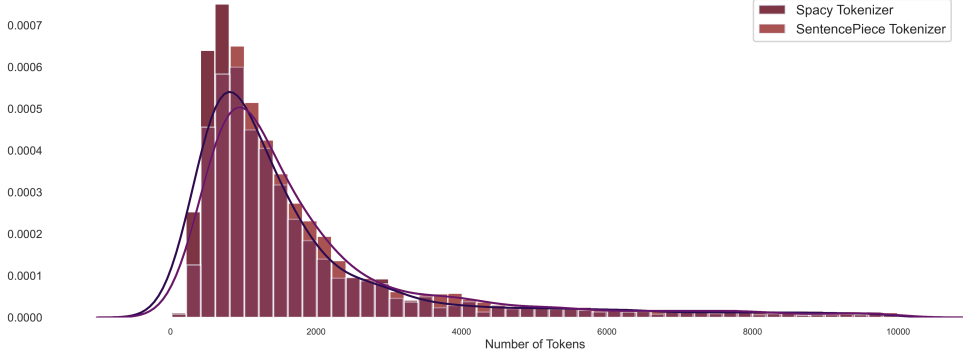


Figure 5: Distribution of the bill document’s length in common words (using the spacy tokenizer) and sub-word units (generated by the SentencePiece tokenizer used in BERT)

## 6.2 General Domain Pre-trained Models

We experiment with Transformer based model BERT (Devlin et al., 2019) and its variants. RoBERTa (Liu et al., 2019) and ALBERT (Lan et al., 2020) are an extension of the standard BERT model. RoBERTa uses a byte-level BPE as a tokenizer and a different pre-training scheme where ALBERT’s model has a smaller parameter size than corresponding BERT models.

## 6.3 Legal Domain Pre-trained Models

Recent research has also demonstrated that language representation models trained on massive corpora and precisely adjusted for a particular domain task perform much better than models trained on task-specific data. This method of transfer learning is beneficial in legal NLP. We thus evaluated three Pre-trained Language models trained from scratch with legal documents, including LegalBERT (Chalkidis et al., 2020), Legal-RoBERTa and Custom Legal-BERT (Zheng et al., 2021).

## 7 Experiments

In this section, we evaluate the mentioned models on the proposed Dataset, describe the executed experiments, and examine the results.

### 7.1 Experimental Settings

In all sequence models, the batch size was set to 64, and the number of epochs was set to 50. At the same time, we iterate through 50 epochs with a batch size of 8 for all Bert-based models.

We used Tensorflow’s Keras API (Abadi et al., 2016) to build sequence models. The BERT-based model follows the base configuration, consisting of 12 layers, 786 units, and 12 attention heads. We developed these models using Pytorch (Paszke

et al., 2019) and obtained pretrained checkpoints from the HuggingFace library (Wolf et al., 2019).

We evaluate the models in two settings: Binary Classification and Multi-Class Classification. For the first setting, the classification layer consists of a dense layer with 1 unit as output, with sigmoid activation. The loss was calculated using binary cross-entropy.

$$\text{Loss}_{\text{bce}} = -\frac{1}{N} \sum_{n=1}^N [y_n \log \hat{y}_n + (1 - y_n) \log (1 - \hat{y}_n)] \quad (2)$$

In The Multi-Class setting, we used a dense layer with five units as output, with softmax activation. In this case, categorical cross-entropy was used for loss calculation.

$$\text{Loss}_{\text{cce}} = -\sum_{i=1}^N y_i \cdot \log \hat{y}_i \quad (3)$$

We perform five runs with different seeds for each method and report the average scores. All the experiments were conducted on the Google Colab Pro and used the default GPU Tesla T4 16GB. The proposed dataset & code are available at [github.com/monk1337/DeepParliament](https://github.com/monk1337/DeepParliament) for reproducibility.

### 7.2 Evaluation Metrics

We assessed the baseline and other models based on their Macro-averaged F1 scores, accuracy, and recall in multi-class and Binary environments. Before calculating the average across labels, macro-averaging computes the metric inside each label.

## 8 Results & Discussion

Table 4 & Table 5 shows the performance of all models in Macro-Precision, Macro-Recall and

Model	Starting Tokens			Middle Tokens			End Tokens		
	Precision	Recall	F1-score	Precision	Recall	F1-score	Precision	Recall	F1-score
ALBERT <sub>Base</sub>	92.21	92.54	92.28	90.13	90.38	90.19	89.06	89.31	88.08
Custom Legal-BERT <sub>Base</sub>	92.47	92.83	92.47	89.51	89.80	89.56	89.03	89.31	89.10
RoBERTa <sub>Base</sub>	92.74	93.06	92.83	89.30	89.53	89.36	89.42	89.70	89.47
Legal-RoBERTa <sub>Base</sub>	92.89	93.17	92.92	90.24	90.51	90.32	90.42	90.63	90.45
Bert <sub>Base</sub>	92.92	93.23	93.01	<b>90.68</b>	<b>90.95</b>	<b>90.73</b>	90.02	90.06	90.19
Legal-BERT <sub>Base</sub>	<b>93.11</b>	<b>93.49</b>	<b>93.11</b>	90.62	90.93	90.61	<b>90.42</b>	<b>90.64</b>	<b>90.49</b>

Table 2: Performance of all Transformer based baseline models in Macro-Precision, Macro-Recall and Macro-F1 (%) on Binary test set under different tokens settings.

Model	Starting Tokens			Middle Tokens			End Tokens		
	Precision	Recall	F1-score	Precision	Recall	F1-score	Precision	Recall	F1-score
Custom Legal-BERT <sub>Base</sub>	51.49	49.12	49.42	42.79	39.96	40.58	41.10	41.21	43.25
ALBERT <sub>Base</sub>	56.14	51.55	52.87	52.55	46.28	47.73	50.71	45.05	45.99
RoBERTa <sub>Base</sub>	56.07	54.85	54.89	49.00	47.05	47.60	45.57	45.11	45.03
Legal-RoBERTa <sub>Base</sub>	61.40	54.74	57.44	48.08	44.02	45.53	50.89	46.98	49.06
Bert <sub>Base</sub>	60.55	55.92	57.86	<b>63.75</b>	<b>52.61</b>	<b>55.64</b>	46.13	46.17	46.34
Legal-BERT <sub>Base</sub>	<b>62.96</b>	<b>56.96</b>	<b>58.79</b>	55.47	49.15	49.86	<b>53.50</b>	<b>47.20</b>	<b>49.68</b>

Table 3: Performance of all Transformer based baseline models in Macro-Precision, Macro-Recall and Macro-F1 (%) on Multi-Class test set under different tokens settings.

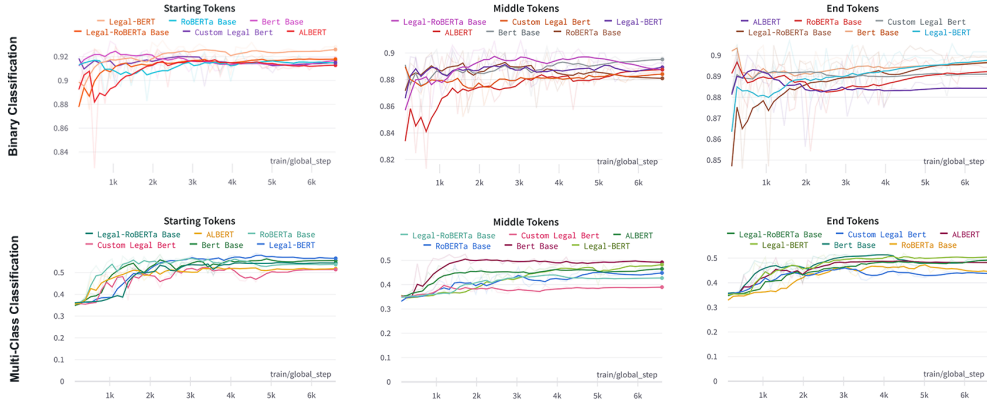


Figure 6: Macro-F1 scores of different Transformer based models on the test dataset.

Macro-F1 on Binary & Multi-Class test set respectively under full token settings. Under the Sequence & Convolutional models category, LSTM performed better than Vanilla CNN in both the Binary and Multi-Class Bill Prediction tasks.

It is observed that there is a significant improvement in the model’s performance when Glove embedding is used as word vectors compared to other embeddings results. BiLSTM + Glove performed best in sequential and convolutional models. CNN + glove gave the second-best results in this category. In the General Domain of Pre-trained Models, transformer models outperform sequential & Convolutional models. Bert<sub>Base</sub> performed best in both Binary and Multi-Class settings while RoBERTa<sub>Base</sub> and ALBERT<sub>Base</sub> are a close second with better f1-score of all the models based in Multi-Class and Binary settings, respectively.

Our assumption was that legal domain models would not perform well on the proposed dataset as India’s legal systems are completely different. However, our assumption did not hold true. In a few settings, Domain-specific models performed well compared to general domain models; This is likely because most words in the proposed dataset are legal domain-specific. The high frequency of unique, domain-specific terminologies appears in the dataset but not in the vocabulary of the Transformer Models trained on the general domain. It is observed that Bert<sub>Base</sub> performs best in terms of precision, while In legal domain trained models, the Legal-BERT<sub>Base</sub> model performs better in recall and f1 score in the Multi-Class classification task. On the other hand, in the Binary classification task, Legal-RoBERTa<sub>Base</sub> & Custom Legal-BERT<sub>Base</sub> performs better than other models.



prioritize adding new bill documents from official sources; introduced in either house of Parliament, i.e. the Lok Sabha or the Rajya Sabha. Documents in the dataset are long and unstructured. Current Transformers models are limited by their input size and cannot process full documents at once. Extended sequence models such as Longformer (Beltagy et al., 2020) and BigBird (Zaheer et al., 2020) are not currently evaluated on the dataset. We leave the investigation of those models on the proposed dataset for other groups to experiment with and publish the results.

Despite the limitations as mentioned above, we believe that the dataset will be helpful to many researchers. as it takes the initial steps to establish a well-defined benchmark to evaluate legal domain models in this field. Models developed on this dataset may assist MPs, presidents, and other legal practitioners.

## Ethics Statement

This study focuses on proposing the first dataset on Parliament Bill status prediction, adheres to the ethical guidelines outlined in the ACL code of Ethics and examines the ethical implications. DeepParliament gathers its data from two public sources. There is no privacy concern since all bill documents are collected against open-access databases. Moreover, the documents do not include personal or sensitive information, except minor information provided by authorities, such as the names of the presidents, Union Council of Ministers, and other official administrative organisations.

The details of dataset collection and statistics are provided in Sections 3 and 4. The model trained to utilise our dataset is mainly meant to support decision-making during bill analysis, not to replace the human specialists.

## References

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Z. Chen, Craig Citro, Gregory S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian J. Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Józefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Gordon Murray, Christopher Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul A. Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda B. Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016.

Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *ArXiv*, abs/1603.04467.

Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preotiuc-Pietro, and Vasileios Lamos. 2016. [Predicting judicial decisions of the european court of human rights: a natural language processing perspective](#). *PeerJ Computer Science*, 2:e93.

Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *ArXiv*, abs/2004.05150.

Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. [Neural legal judgment prediction in english](#). *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4317–4323.

Ilias Chalkidis, Ion Androutsopoulos, and Achilleas Michos. 2018. [Obligation and prohibition extraction using hierarchical rnns](#). *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 254–259.

Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. [Legal-bert: The muppets straight out of law school](#). *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2898–2904.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). *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.

Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. *AISTATS*.

Dan Hendrycks, Collin Burns, Anya Chen, and Spencer Ball. 2021. [Cuad: An expert-annotated nlp dataset for legal contract review](#). *ArXiv*, abs/2103.06268.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Comput.*, 9:1735–1780.

Phi Manh Kien, Ha-Thanh Nguyen, Ngo Xuan Bach, Vu Tran, Minh Le Nguyen, and Tu Minh Phuong. 2020. [Answering legal questions by learning neural attentive text representation](#). *Proceedings of the 28th International Conference on Computational Linguistics*, pages 988–998.

Yoon Kim. 2014. Convolutional neural networks for sentence classification. *EMNLP*.

Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2020. [Reformer: The efficient transformer](#). *ArXiv*, abs/2001.04451.



- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. Albert: A lite bert for self-supervised learning of language representations. *ArXiv*, abs/1909.11942.
- Elena Leitner, Georg Rehm, and Julian Moreno-Schneider. 2020. A dataset of german legal documents for named entity recognition. *LREC*.
- Elena Leitner, Georg Rehm, and Julián Moreno Schneider. 2019. Fine-grained named entity recognition in legal documents. *SEMANTiCS*.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario vSavsko, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clement Delangue, Théo Matussiere, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander M. Rush, and Thomas Wolf. 2021. Datasets: A community library for natural language processing. *ArXiv*, abs/2109.02846.
- 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*, abs/1907.11692.
- Ramesh Nallapati and Christopher D Manning. 2008. Legal docket classification: Where machine learning stumbles. *EMNLP*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. *EMNLP*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Chaojun Xiao, Haoxiang Zhong, Zhipeng Guo, Cunchao Tu, Zhiyuan Liu, Maosong Sun, Yansong Feng, Xianpei Han, Zhen Hu, Heng Wang, and Jianfeng Xu. 2018. Cail2018: A large-scale legal dataset for judgment prediction. *ArXiv*, abs/1807.02478.
- Hai Ye, Xin Jiang, Zhunchen Luo, and Wenhan Chao. 2018. [Interpretable charge predictions for criminal cases: Learning to generate court views from fact descriptions](#). *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1854–1864.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33.
- Lucia Zheng, Neel Guha, Brandon R Anderson, Peter Henderson, and Daniel E Ho. 2021. When does pre-training help? assessing self-supervised learning for law and the casehold dataset. *Proceedings of the 18th International Conference on Artificial Intelligence and Law*.