

Classification on Sentence Embeddings for Legal Assistance

Arka Mitra¹

¹Indian Institute of Technology, Kharagpur, India

Abstract

Legal proceedings take plenty of time and also cost a lot. The lawyers have to do a lot of work in order to identify the different sections of prior cases and statutes. The paper tries to solve the first tasks in AILA2021 (Artificial Intelligence for Legal Assistance) that will be held in FIRE2021 (Forum for Information Retrieval Evaluation). The task is to semantically segment the document into different assigned one of the 7 predefined labels or "rhetorical roles." The paper uses BERT to obtain the sentence embeddings from a sentence, and then a linear classifier is used to output the final prediction. The experiments show that when more weightage is assigned to the class with the highest frequency, the results are better than those when more weightage is given to the class with a lower frequency. In task 1, the team legalNLP obtained a F1 score of 0.22.

Keywords

Deeplearning, Sentence Embeddings, BERT, Classification, Natural Language Processing

1. Introduction

Legal systems in many countries like USA, UK, Canada, India has two main sources- Precedents and Statutes; Precedents are previous similar cases while statutes are written laws that have to be followed in the country. The number of legal cases have been increasing and thus it is quite difficult for a lawyer to go through many of the precedents. Additionally, the legal reports in different countries are structured in different ways. Due to the lack of standardization, it is necessary to have a method that can help the lawyer to identify the different sentences in the report and process the report faster, while obtaining the relevant information quickly. The task 1 of AILA 2021 aims for the semantic segmentation of the document to assist the lawyer to process the information faster.

AILA 2021, held in collocation with FIRE 2021, has several tasks for legal informatics. Legal documents follow certain sections like "Facts of the Case", "Issues being discussed" etc which are called "rhetorical roles". For task 1, the sentences had to be classified into one of the seven different classes. More details on the classes and the dataset have been given in section 3.

The remaining of the paper is divided into the following sections: 2 which goes through the related work done for Rhetorical labelling in legal reports; 3 that provides a detail of the dataset that has been used; 4 describes the methodology that has been used for the paper; 5 showcases

Forum for Information Retrieval Evaluation, December 13-17, 2021, India


✉ thearkamitra@iitkgp.ac.in (A. Mitra)

🌐 <https://n.ethz.ch/~amitra/> (A. Mitra)

🆔 0000-0003-1071-7294 (A. Mitra)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

the results that have been obtained; 6 discusses the results and provides insights on the different models that have been used. 7 talks of the future work that would be done and 8 concludes the paper.

2. Related Work

Text segmentation has been an important task in natural language processing. There have been probabilistic approaches that used Hidden Markov Models [1] and Maximum Entropy Markov Models [2]. Saravanan and Ravindran [3] used Conditional Random Fields (CRF) for the identification of rhetoric labels for the segmentation and summarization of legal documents. Avelka et al. [4] used CRFs on annotated data from the US cyber crime and trade secrets decisions. Bhattacharya et al. [5] used CRF on top of a Bi-LSTM network to classify the sentences into different categories.

3. Dataset

The AILA track started in 2019 [6] and it had focused on Precedent and Statute retrieval. The second version of the same track focused [7] on precedent and statute retrieval as well as rhetoric labelling [8]. The third iteration of the track [9] also includes rhetoric labelling but at the same time, it contains a task for automatic summarization [10].

There are 60 documents in the task 1 dataset with 11285 labelled sentences. Each of the sentences has one of the seven possible labels:

- Facts : Sentences that discuss the facts about the case
- Ruling by Lower Court : The dataset contains Indian Supreme Court cases, which usually have a ruling at a lower court like High Court or Tribunal; the label indicates that the sentences are the decisions given in the lower court
- Argument : Arguments provided by the different parties
- Statute : The statute corresponding to the present case
- Precedent : The precedent corresponding to the present case
- Ratio of the decision : The reasoning given by the Supreme Court for the decision
- Ruling by Present Court: The final decision given by the Supreme Court

4. Method

The first subsection discusses the approach for the task and the next subsection provides the experimental details.

4.1. Approach

The dataset contained seven different classes but the distribution among those classes are quite skewed. Table 1 shows that the number of samples with the label “Ratio of decision” is about




Figure 1: Overall Methodology

12 times the number of samples with the label “Ruling by Present Count”. It is important to keep the number of samples almost equal- that would allow the model to learn meaningful information from each of the classes.

There are three main methods to do that. In the first method, the sampling from the dataset can be done in such a way that more samples from the lower frequency class are taken and lower samples are taken from the class with higher frequency. The downside to this method is that we are leaving several samples from the dataset which would decrease the performance of the model. In another method, one can do the sampling in such a way that the same example of the class with lower frequency is included into the batch multiple times such that the total distribution in the new dataset has the same number for each of the classes. However, since the same example is chosen multiple times, it is increasing the chances of overfitting and also increases the computation overhead. The last method keeps the computational cost about the same as the first method but at the same time does not change the dataset size. In this method, the loss is changed so that false predictions for the class with lower number of samples is penalized more. The new loss, as shown in Eqn. 1, is a modified version of the cross-entropy loss. The weight multiplied with the cross-entropy can be considered as the number of times the sample has been considered. A class with a higher number of samples in the dataset needs to have a lower weight associated with it.

$$loss(logits, class) = weight[class] * \left(-\log\left(\frac{\exp(logits[class])}{\sum_j \exp(logits[j])}\right) \right) \quad (1)$$

Table 1

The distribution of the different classes

Label	Count of Label
Ratio of the decision	4211
Facts	2622
Precedent	1787
Argument	939
Statute	902
Ruling by Lower Court	483
Ruling by Present Court	341

The author preprocessed the document and combined the sentences in all the documents and the associated labels to create the dataset that has been used. BERT[11] was used for creating the sentence embeddings for the sentences. BERT has been trained on a lot of data and thus would be able to create a condensed representation of the sentence. The output of the “CLS” token from the BERT output was considered to be the sentence embedding for the sentence and then that was sent through a linear layer. The logits obtained from the linear layer were considered and the maximum of them was selected to be the predicted class of the section. The overall methodology has been described in Figure 1.

4.2. Experimental Details

The cased and uncased BERT model have been implemented with the help of the Huggingface library [12]. Pytorch has been used for the framework. The batch size of 8 has been considered and the model has been trained for 4 epochs. 80% of the data has been considered as the training set and the rest is considered as the validation set. The model weights which gave the best results on the validation set had been saved and used for inference on the test set. AdamW optimizer [13] with an initial learning rate of 2e-5 is used for training. The max length for the padding was kept at about 0.98 percentile which is around 120 tokens. The codes are publicly available on github¹. The random seed has been set to 42.

5. Results

There are three runs which have been submitted finally. The macro-F1 score, precision and recall of the different runs are given in Table. 2.

The description of the three runs are as follows:

- The first run uses base cased BERT with weights to modify the cross entropy loss
- The second run uses base uncased BERT with the same weights as the previous run.
- The third and final run uses base cased BERT but here the weights are inverted such that the class with higher number of samples is given more weights.

¹<https://github.com/thearkamitra/LegalNLP>

Table 2

The results of the different runs

Run number	Precision	Recall	F1
legalNLP_1	0.197	0.217	0.196
legalNLP_2	0.198	0.215	0.192
legalNLP_3	0.225	0.227	0.22

6. Discussion

The results show that the cased BERT model performed better than the uncased model. This can be explained by the fact that there may be some phrases in legal reports that have different meanings when used in uppercase vs lowercase and the BERT cased model is able to capture the contextual information. Due to the better performance of cased BERT, the author performed the same experiment with the same random seeds but with different weight for the cross-entropy loss. Comparison between the first and third run shows that the model performed better when more weightage was given to classes that existed more abundantly. This contradicts the belief that a model trained with skewed distribution would perform worse than one without. A possible explanation might be that the test data has more sentences with labels corresponding to the higher class. As a consequence, the metric reports a higher score for the third metric.

7. Future Work

In the present work, the sentences from the documents had been extracted and aggregated to form the dataset. However, there is a relation between the labels and where the sentence is located in the document (for example, Ruling by Present Court is always present in the final ending of the documents). Also, the author has not considered the co-occurrence of the different label. For that, Hidden markov model or some probabilistic state machine could be used to further improve the accuracy of the model.

8. Conclusion

The paper describes the modified cross-entropy loss and the use of BERT models for rhetoric role labelling in legal documents. The three runs that had been submitted obtained a score of 0.196, 0.192 and 0.22 respectively.

Acknowledgments

The author thanks the organizers of Artificial Intelligence for Legal Assistance for creating this task. The author would also like to acknowledge Google Colab for providing the computational resources needed. The BERT model is built on the library made by Huggingface [12].

References

- [1] V. Borkar, K. Deshmukh, S. Sarawagi, Automatic segmentation of text into structured records, SIGMOD record 30 (2001) 175–186.
- [2] A. McCallum, D. Freitag, F. C. Pereira, Maximum entropy markov models for information extraction and segmentation, in: ICML, 2000.
- [3] M. Saravanan, B. Ravindran, Identification of rhetorical roles for segmentation and summarization of a legal judgment 18 (2010) 45–76. URL: <https://doi.org/10.1007/s10506-010-9087-7>. doi:10.1007/s10506-010-9087-7.
- [4] J. Savelka, K. D. Ashley, Segmenting u.s. court decisions into functional and issue specific parts, in: JURIX, 2018.
- [5] P. Bhattacharya, S. Paul, K. Ghosh, S. Ghosh, A. Z. Wyner, Identification of rhetorical roles of sentences in indian legal judgments, ArXiv abs/1911.05405 (2019).
- [6] P. Bhattacharya, K. Ghosh, S. Ghosh, A. Pal, P. Mehta, A. Bhattacharya, P. Majumder, Fire 2019 aila track: Artificial intelligence for legal assistance, Proceedings of the 11th Forum for Information Retrieval Evaluation (2019).
- [7] P. Bhattacharya, P. Mehta, K. Ghosh, S. Ghosh, A. Pal, A. Bhattacharya, P. Majumder, Overview of the fire 2020 aila track: Artificial intelligence for legal assistance, in: FIRE (working notes), 2020.
- [8] P. Bhattacharya, S. Paul, K. Ghosh, S. Ghosh, A. Wyner, Identification of rhetorical roles of sentences in indian legal judgments, in: Proc. International Conference on Legal Knowledge and Information Systems (JURIX), 2019.
- [9] V. Parikh, U. Bhattacharya, P. Mehta, A. Bandyopadhyay, P. Bhattacharya, K. Ghosh, S. Ghosh, A. Pal, A. Bhattacharya, P. Majumder, Overview of the third shared task on artificial intelligence for legal assistance at fire 2021, in: FIRE (Working Notes), 2021.
- [10] V. Parikh, U. Bhattacharya, P. Mehta, A. Bandyopadhyay, P. Bhattacharya, K. Ghosh, S. Ghosh, A. Pal, A. Bhattacharya, P. Majumder, Fire 2021 aila track: Artificial intelligence for legal assistance, in: Proceedings of the 13th Forum for Information Retrieval Evaluation, 2021.
- [11] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, in: NAACL, 2019.
- [12] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, A. M. Rush, Huggingface’s transformers: State-of-the-art natural language processing, 2020. arXiv:1910.03771.
- [13] I. Loshchilov, F. Hutter, Decoupled weight decay regularization, in: ICLR, 2019.