

LegalNLP - Natural Language Processing methods for the Brazilian Legal Language

Felipe Maia Polo¹, Gabriel Caiaffa Floriano Mendonça²,
Kauê Capellato J. Parreira², Lucka Gianvechio², Peterson Cordeiro²,
Jonathan Batista Ferreira³, Leticia Maria Paz de Lima³,
Antônio Carlos do Amaral Maia⁴, Renato Vicente^{2,5}

¹Department of Statistics, University of Michigan
Ann Arbor, Michigan, United States of America

²Institute of Mathematics and Statistics, University of São Paulo
São Paulo, São Paulo, Brazil

³Molecular Sciences, University of São Paulo
São Paulo, São Paulo, Brazil

⁴Tikal Tech
São Paulo, São Paulo, Brazil

⁵Latam Datalab Serasa Experian
São Paulo, São Paulo, Brazil

maiapolo@umich.edu

{gabrielcaiaffa, kauecapellato, luckagg, peterson.cordeiro,

jonathanbf, leticiamaria, rvicente}@usp.br

antonio.maia@tikal.tech

***Abstract.** We present and make available pre-trained language models (Phraser, Word2Vec, Doc2Vec, FastText, and BERT) for the Brazilian legal language, a Python package with functions to facilitate their use, and a set of demonstrations/tutorials containing some applications involving them. Given that our material is built upon legal texts coming from several Brazilian courts, this initiative is extremely helpful for the Brazilian legal field, which lacks other open and specific tools and language models. Our main objective is to catalyze the use of natural language processing tools for legal texts analysis by the Brazilian industry, government, and academia, providing the necessary tools and accessible material.*

1. Introduction

The term Natural Language Processing (NLP) defines the area of research and applications represented by statistical models and algorithms responsible for the analysis and representation of natural language, both phonetic and written. In several countries, the applications of NLP methods in Law are becoming more and more present and represents an increasingly promising future. Classification of legal documents, named entity recognition in legal texts or predictions of legal decisions are just some examples of possible applications.

The first objective of this work is to present and make available pre-trained language models for the Brazilian legal language in addition to a Python package with functions to facilitate the use of these models. The second objective of this work is to present and make available demonstrations¹, which are accessible to the general public, on how to use language models to analyze Brazilian legal texts in real situations using our models.

The organization of this work is as follows: in Sections 2 and 3, we give an overview on NLP methods and mention some related works, besides discussing the importance of this work for the current Brazilian legal context; in Section 4, we briefly present the functions that are available in the first version of our package; in Section 5 we detail the datasets used for training the language models; in Section 6 we present each of the pre-trained language models provided by us and, finally, in Section 7, we present two demonstrations involving our models.

The GitHub repository containing our library (models + package + demonstrations) can be accessed through <https://github.com/felipemaiapolo/legalnlp>.

2. Related Work

2.1. NLP background

Natural Language Processing (NLP) went through a revolution in the last ten years with the advances of Deep Learning methods. Perhaps the first big breakthrough of NLP in the last ten years was the popularization of dense word embeddings with the Word2Vec model [Mikolov et al. 2013a, Mikolov et al. 2013b]. In the way Word2Vec was formulated, the vectors representing words carry the words' meanings derived from the context in which each word can be found. A year later, Doc2Vec was introduced [Le and Mikolov 2014] as a generalization of Word2Vec to whole texts and, in 2017, FastText [Bojanowski et al. 2017] was introduced as a way to obtain more robust word vectors, taking into account also word morphology. Also in 2017, the use of self-attention mechanisms and the Transformer architecture [Vaswani et al. 2017] in text processing shifted to whole field of NLP. In more recent years, BERT [Devlin et al. 2018] and derived models used the Transformer architecture to elevate the state of the art in a variety of NLP tasks such as text classification, question answering, and named entity recognition. With the development of new research and methods in NLP, the availability of models for the general public becomes necessary, and then efforts begin to emerge in different countries and fields to develop their own tools.

2.2. NLP for Brazilian Portuguese and Legal applications

In the last years there were some efforts to pre-train language models for the Brazilian Portuguese [Hartmann et al. 2017, Souza et al. 2020] regarding the more traditional word embeddings (Word2Vec, FastText, etc) as well as BERT models. In spite of that, those models were trained with general documents and were not designed to represent Brazilian legal language. The legal language is

¹Full demonstrations are in Google Colab notebooks and are available in our GitHub repository <https://github.com/felipemaiapolo/legalnlp>.

unique and demands its own specific language models in order to solve NLP problems that arise everyday in the legal field. In fact, there is a rapid increase in the use of NLP tools to solve real problems in Law, such as identifying the parties in legal proceedings [Nguyen et al. 2018], classification of legal documents according to their administrative labels [da Silva et al. 2018, Braz et al. 2018, Polo et al. 2021] or predicting the area that a lawsuit belongs to [Sulea et al. 2017]. These works used both classical machine learning paradigm (e.g., TFIDF features + classifier) [Polo et al. 2021, Sulea et al. 2017] or deep learning methods [Nguyen et al. 2018, da Silva et al. 2018, Braz et al. 2018, Polo et al. 2021] in order to accomplish their goals. Therefore, making pre-trained language models for the Brazilian legal language available can be a turning point for the field in Brazil during the next decades.

3. Brazilian legal context and the relevance of this work

Law is a challenging field for NLP applications, especially in Brazil. The difficulties are located in the complexity of communication, because legal discipline is fully based on a proper and singular language, spoken for centuries by a minority, and that resembles but does not coincide with the way ordinary people usually speak. Also, Brazilian judicial landscape adds complexity to this communication system, because Brazilian jurisdiction is bigger than Canada's, Australia's, England's, Wales', Scotland's, Ireland's, Belgium's, Switzerland's and France's, combined. The country covers more than 8 millions square kilometers and there are more than 210 million people living in 26 states, each one with their own regional and federal courts, in addition to the national ones. Regionalism can be an obstacle for NLP solutions as there are different ways to speak Portuguese in different regions of the country even within the legal system. For example, judges for different parts of the country can express themselves using different words to designate the same legal reality.

Given the context of the Brazilian legal language, it is necessary (i) for the legal field to have its own pre-trained language models, in order to properly capture singularities of that reality, and (ii) that those models should be trained using texts from several sources, including different courts and Brazilian regions. Given that our library is built upon and for legal texts, from a diversity of sources, this initiative of making functions and models available is a real help in the Brazilian NLP law landscape. It provides researchers in academia, industry or government the tools with necessary context embarked, eliminating or at least mitigating the burden and costs for their activities, permitting them to approach their issues promptly.

4. Tools for text manipulation

4.1. Text cleaning functions

Our package offers two text cleaning functions, the uses of which are optional. The first one (*clean*) is designed for general use and is compatible with the Phraser, Word2Vec, FastText, and Doc2Vec models presented in Section 6. In addition, it uses regular expression methods (RegEx) that can extract from texts certain entities such as numbers, values, dates, lawyers IDs² etc. On the other hand,

²Known in Brazil as "código OAB".

the second function (*clean_bert*) is much simpler and makes fewer changes to the texts. We recommend using this function, or some adaptation of it, when using BERTikal, which is our BERT-Base (cased) model pre-trained for the Brazilian legal language.

For more information on cleaning functions, please check our library documentation at <https://github.com/felipemaiapolo/legalnlp>.

4.2. Text tokenization

For tokenization purposes, we offer two pre-trained Phraser models which can be used in conjunction with the *split* method in Python. Our Phraser models helps us to identify which tokens should be merged (into bigrams, trigrams, and quadrigrams) during tokenization step and are compatible with our Word2Vec, FastText, and Doc2Vec models. More details on the Phraser models is presented in Section 6.1. Furthermore, BERTikal has its own vocabulary which must be used for tokenization.

5. Data used to pre-train language models

In total, we used four text datasets. In the process of training Phraser, Word2Vec, Doc2Vec, and FasText we used only the first dataset (Data 1) and for training BERTikal³ we used the other three datasets (Data 2, 3, 4) and started training from the checkpoint provided by a recent work [Souza et al. 2020]. This difference is arbitrary and depended on the context in which the models were trained.

The first and last datasets (Data 1 and 4) are composed exclusively of publications (*publicações*), also known as clippings (*recortes*), from several Brazilian courts for the years 2019 and 2020. On the other hand, the second dataset (Data 2) is composed of longer legal documents mainly from the Court of Justice of São Paulo (TJSP). This is a sample of the dataset used in a master thesis work [Massoni 2021] and was kindly provided to us by its author. The third dataset (Data 3) is composed exclusively of motions (*movimentações*) from several Brazilian courts for the years 2019 and 2020.

In Table 1, we have information about the number of texts and the approximate size, in gigabytes, of each of the text datasets:

Table 1. Number of texts/documents in each dataset and their approximate size

Data sources	Number of texts	Approximated size (GB)
Data 1	1772351	3,14
Data 2	34369	0,74
Data 3	5246350	0,77
Data 4	705521	1,12

In order to bring more information about each of the datasets, we detail their sources. Unfortunately, given the way in which we obtained the data, it was not possible to directly know which courts each of the texts came from. Therefore, for

³Our BERT-Base (cased) model pre-trained for the Brazilian legal language.

each of the datasets, we separated a random sample of a maximum of 50k texts and searched for the court of origin for their IDs (Número Único de Processo - NUP) that appeared in the body of the text. In cases where more than one NUP appeared in the texts, we assumed that the first one was the correct one. Unfortunately, not all texts consulted had a NUP present: for Data 1, 49,006 texts had a NUP; for Data 2, 34,175 texts had a NUP; for Data 3, 1,034 texts had a NUP; finally, for Data 4, 48,802 texts had a NUP. In Table 2, we can see the courts of origin of the texts with their respective relative frequencies, calculated from the random sample already described. We show the fifteen courts with the highest number of texts:

Table 2. Datasets’ sources. We show the fifteen courts with the highest number of texts present in each dataset

Data 1		Data 2		Data 3		Data 4	
Tribunal	%	Tribunal	%	Tribunal	%	Tribunal	%
TJSP	36,05	TJSP	99,67	TJSP	86,80	TJSP	43,20
TRT15	10,83	TRT82	0,05	TJAL	5,87	TRT02	6,36
TRT02	8,13	TJRJ	0,04	TJRN	1,66	TJBA	5,19
TRT03	4,00	TRE82	0,03	TJSC	1,08	TRT15	4,47
TJBA	3,86	TJMG	0,03	TJRJ	0,78	TRT01	3,72
TRT01	3,58	TRF82	0,03	TJMG	0,59	TRT09	2,78
TJSC	3,03	TJCE	0,02	TRT15	0,49	TJSC	2,45
TRT04	2,94	TJRS	0,02	TRF03	0,49	TJMS	2,39
TRT09	2,50	TJPR	0,02	TJBA	0,49	TRT03	2,36
TRT05	1,99	TJBA	0,01	TJMS	0,39	TRF03	2,35
TRF03	1,88	TJGO	0,01	TRT02	0,20	TRT05	2,24
TJMS	1,83	TJMS	0,01	TRT03	0,20	TJDF	2,13
TJDF	1,81	TRE55	0,01	TJAC	0,20	TJSE	1,97
TRT12	1,47	TJSC	0,01	TRT01	0,20	TRT04	1,90
TJRJ	1,41	TRE17	0,01	TRF02	0,10	TJRJ	1,83

6. Pre-trained language models

In this section, we go deeper into our pre-trained models. Our focus is to give more details about the models we are making available, also going through the parameters used for the training phase.

6.1. Phraser

Phraser is a statistical method proposed in the natural language processing literature [Mikolov et al. 2013b] for identifying which words, when they appear together, can be considered as unique tokens. This method application is able to identify the relevance of the occurrence of a bigram against the occurrence of the words that make it up separately. Thus, we can identify that a bigram like "São Paulo" should be treated as a single token, for example. If the method is applied a second time in sequence, we can check which are the relevant trigrams and quadrigrams. Since the two applications should be done with different Phraser models, it can be the case that the second application identifies bigrams that were not identified by the first model.

As based on Gensim's package⁴ for the Python language, we used two Phraser models, which should be applied sequentially. The first model trains for bigrams and the second trains for trigrams, quadrigrams and also finds bigrams that were not found by the first model - an example is given in Section 6.1.2. The version of the package that was used to train the models is the 3.8.3, though our models can be used with other versions. In case there is any problem resulting from a conflict of versions, we suggest that the user uses the same version as us.

6.1.1. Data, text preprocessing and models' parameters

The data used for training the two models are the Data 1 presented in Section 5. The text preprocessing phase, performed before training the models, was composed by text cleaning and lower casing steps, performed using the general cleaning function presented in Section 4.1. Finally, as model training parameters, we kept the default configuration proposed by the Gensim package, as can be seen in its webpage⁴. It is worthy noting that we did not remove stopwords.

6.1.2. Usage example

In order to illustrate how Phraser works in practice, we extracted a piece of text from our data. First, we pass the snippet through our general cleaning function, then we tokenize the text, and apply the Phraser models. Then we put the tokens together to reconstruct the texts. Phraser models combine different tokens into unique tokens using the "_" character:

- **Clean text:** "(...) direito do consumidor origem : bangu regional xxix juizado especial cível ação : [processo] - - recte : fundo de investimento em direitos creditórios (...)"
- **Applying Phraser 1x:** "(...) direito do consumidor origem : bangu regional xxix juizado_especial cível_ação : [processo] - - recte : fundo de investimento em direitos_creditórios (...)"
- **Applying Phraser 2x:** "(...) direito do consumidor origem : bangu_regional xxix juizado_especial_cível_ação : [processo] - - recte : fundo de investimento em direitos_creditórios (...)"

It is possible to see that the first Phraser model merged the tokens "juizado" "especial" resulting in "juizado_especial" and "cível" "ação" resulting in "cível_ação", for example. The second Phraser model merged "juizado_especial" "cível_ação" into "juizado_especial_cível_ação" and "bangu" "regional" into "bangu_regional".

6.2. Word2Vec/Doc2Vec

Our first models for generating vector representation for tokens and texts (embeddings) are variations of the Word2Vec [Mikolov et al. 2013b, Mikolov et al. 2013a] and Doc2Vec [Le and Mikolov 2014] methods. In short, the

⁴See https://radimrehurek.com/gensim_3.8.3/models/phrases.html - Accessed on 07/30/2021.

Word2Vec methods generate embeddings for tokens⁵ and that somehow capture the meaning of the various textual elements, based on the contexts in which these elements appear. Doc2Vec methods are extensions/modifications of Word2Vec for generating whole text representations.

The Word2Vec and Doc2Vec methods are presented together in this section as they were trained together using the Gensim⁶ Python package. The Gensim version used to train the models is 3.8.3, though our models can be used with other versions. In case there is any problem resulting from a conflict of versions, we suggest that the user uses the same version as us.

6.2.1. Data, text preprocessing and models' parameters

The dataset used for training the two set of models are the "Data 1" presented in Section 5. The text preprocessing phase, performed before training the models, was composed by a text cleaning and lower casing, performed using the general cleaning function presented in Section 4.1, and by tokenization alongside with the application of the Phraser models, described in Section 6.1, twice in sequence. In that way, we train representations for unigrams, bigrams, trigrams, and quadrigrams.

Both Word2Vec and Doc2Vec were trained with three different sizes (100, 200, 300), with window equal to 15, number of epochs equal to 20 and using its two most known versions: Skip-Gram (SG) and Continuous Bag of Words (CBOW) for Word2Vec [Mikolov et al. 2013b, Mikolov et al. 2013a] and Distributed Memory (DM) and Distributed Bag of Words (DBOW) for Doc2Vec [Le and Mikolov 2014]. Furthermore, we used the option *dm_mean=1* when we trained Doc2Vec DM and Word2Vec CBOW, ignored tokens that appear less than 50 times in our corpus and set all the other parameters to Gensim's default values.

6.3. FastText

The FastText [Bojanowski et al. 2017] methods, like Word2Vec, form a class of models for creating vector representations (embeddings) for tokens. Unlike Word2Vec, which disregards the morphology of the tokens and allocates a different vector for each one of them, the FastText methods consider that each one of the tokens is formed by n-grams of characters or substrings. In this way, the representation of tokens which do not appear in the training set can be inferred from the representation of substrings. Also, rare tokens can have more robust representations than those returned by the Word2Vec methods.

The implementation used to train our FastText models is from the Gensim⁷ Python package. The Gensim version used to train the models is 4.0.1, though our models can be used with other versions. In case there is any problem resulting from a conflict of versions, we suggest that the user uses the same version as us.

⁵Most of the time they are words, n-grams or punctuation.

⁶See https://radimrehurek.com/gensim_3.8.3/models/doc2vec.html - Accessed on 07/30/2021.

⁷<https://radimrehurek.com/gensim/models/fasttext.html> - Accessed on 07/30/2021.

6.3.1. Data, text preprocessing and models' parameters

The dataset used for training the two set of models are the "Data 1" presented in Section 5. The text preprocessing phase, performed before training the models, was composed by a text cleaning and lower casing, performed using the general cleaning function presented in Section 4.1, and by tokenization alongside with the application of the Phraser models, described in Section 6.1, twice in sequence. In that way, we train representations for unigrams, bigrams, trigrams and quadrigrams.

The FastText methods were trained in the same way as Word2Vec and Doc2Vec.

6.4. BERTikal

We call BERTikal our BERT-Base model [Devlin et al. 2018] (cased) for Brazilian legal language. BERT models are models based on neural network architectures called Transformers, which in turn are based on the concept of self-attention [Vaswani et al. 2017]. BERT models are trained with large sets of texts using the self-supervised paradigm, which is basically solving unsupervised problems using supervised techniques. A pre-trained BERT model is capable of generating representations for entire texts and can be adapted for a supervised task, e.g., text classification or question answering, using the fine-tuning mechanism. Fine-tuning consists of training a model that solves a supervised learning task using labeled data and the pre-trained BERT model as its starting point [Devlin et al. 2018].

BERTikal was trained using the Python package Transformers⁸ from the company Hugging Face in its 4.2.2 version and its checkpoint made available by us is compatible with PyTorch⁹ 1.9.0. Although we expose the versions of both packages, more current versions can be used in applications of the model, as long as there are no relevant version conflicts.

6.4.1. Data, text preprocessing and model parameters

The datasets used for training the model are Data 2, 3, 4 presented in Section 5. Moreover, the text preprocessing phase, carried out before the training of the models, was composed of a specific text cleaning for BERTikal, performed using the *clean_bert* function presented in Section 4.1. That function makes few changes to the texts and keeps the uppercased letters.

Our model was trained from the checkpoint made available in Neuralmind's Github repository¹⁰ by the authors of a recent research [Souza et al. 2020]. In the training phase, we (i) kept the configuration of the model and vocabulary used by the authors, (ii) used the Masked Language Model (MLM) objective with masking probability 0.15, (iii) used one epoch, (iv) batch size equals to 4 texts, and (v) made use of a Tesla T4 GPU. The optimizer settings have been set as the

⁸See <https://huggingface.co/transformers/> - Accessed on 07/30/2021.

⁹See <https://pytorch.org/> - Accessed on 07/30/2021.

¹⁰<https://github.com/neuralmind-ai/portuguese-bert> - Accessed on 07/30/2021.

default for the Transformers package from the company Hugging Face¹¹ and the full training took approximately one week to be completed.

7. Demonstrations

Now we present some demonstrations of uses of the pre-trained models. The full demos are available on GitHub¹² and are intended to help the practitioner to make real applications of the models and functions that we make available. The first demonstration presented below consists of an illustration of the vector representations for tokens returned by our 100-sized Word2Vec CBOW model; Secondly, we present solutions to a classification problem, carried out on a legal dataset used in the literature [Polo et al. 2021] and also made available on Kaggle¹³. We should clarify that the purpose of those classification experiments are not to extensively compare models' performances, but only show an example of real application.

7.1. Visualizing tokens

In this first example, we perform a dimensionality reduction step on the embeddings returned by our Word2Vec model. To reduce the dimension of the embeddings, we apply principal component analysis (PCA), reducing the dimension from 100 to 2. In this way, it is possible to see the vectors in a bidimensional scatter plot. From there, we draw the graph in Figure 1.




Figure 1. Bidimensional representation of tokens generated by Word2Vec + PCA. The legend shows the central word, the tokens of the same color are the 5 ones with the greatest cosine similarity between their embedding and the central word. You can see how similar words tend to cluster together.

¹¹See https://github.com/huggingface/transformers/blob/master/src/transformers/training_args.py - Accessed on 10/22/2020.

¹²See <https://github.com/felipepolo/legalnlp>

¹³See <https://www.kaggle.com/felipepolo/brazilian-legal-proceedings>.

This kind of visualization is useful to exemplify the property of tokens to approximate other similar tokens, that is, those used in similar contexts. In Figure 1, the dots of the same color are the tokens with the greatest cosine similarity between their embedding and the tokens' embeddings in the legend. We can observe that the tokens form clusters. Female names, city names, tokens related to "juiz" (judge, in English) are clustered together, for example.

7.2. Predicting legal proceedings status

In this case study, we use the dataset "Brazilian Legal Proceedings"¹⁴ which has data from 6449 brazilian legal proceedings, each one classified as archived (47.14%), active (45.23%) or suspended (7.63%). Our objective in this study is to, using the most recent text in each proceeding, perform a classification task. This is in contrast with previous work, which the five last texts were used [Polo et al. 2021].

To that end, we have combined different language models and classifiers. Firstly, we use a Convolutional Neural Network (CNN) architecture with frozen embeddings (100-sized Word2Vec CBOV models) from our library and from the NILC/USP repository [Hartmann et al. 2017]. Secondly, we combined CatBoost [Prokhorenkova et al. 2017] with our 100-sized Doc2Vec DM. Finally, we combined CatBoost with BERTikal and BERTimbau [Souza et al. 2020] (cased BERT-base) embeddings for texts. We did not fine tune the BERT models.

In order to train and test the different approaches, we randomly split our dataset in training (70%) and test (30%) sets. To train the CNNs, we used batch learning with batch size equals 500, 50 epochs, and an early stop with 15 rounds of patience and a validation split of 10% from training set, we have used 32 filters with kernel size equals 3. For the CatBoost classifier, we used the CatBoost package default hyperparameters' values and an early stop with 100 rounds of patience and a validation split of 10% from training set.

The experiments' results obtained can be observed in Table 3.

Table 3. Results obtained for each model in the test set (score \pm bootstrap std. error)

Model	Classifier	Accuracy	F1 (Macro Avg.)	F1 (Weighted Avg.)
Ours	—	—	—	—
Word2Vec	CNN	0.84 \pm 0.01	0.80 \pm 0.01	0.84 \pm 0.01
Doc2Vec	CatBoost	0.86 \pm 0.01	0.82 \pm 0.01	0.85 \pm 0.01
BERTikal	CatBoost	0.86 \pm 0.01	0.82 \pm 0.01	0.86 \pm 0.01
NILC/USP	—	—	—	—
Word2Vec	CNN	0.85 \pm 0.01	0.82 \pm 0.01	0.85 \pm 0.01
BERTimbau	—	—	—	—
BERT-Base	CatBoost	0.84 \pm 0.01	0.79 \pm 0.01	0.84 \pm 0.01

From the results presented in Table 3 we can see that our models got similar results¹⁵ at accuracy and F1 scores with the two benchmark models. This is

¹⁴See <https://www.kaggle.com/felipepolo/brazilian-legal-proceedings>.

¹⁵The error bars give us an idea on how results can fluctuate.

somewhat expected given the dataset we used. As seen in previous work, different approaches have similar results in classification tasks using this dataset [Polo et al. 2021]. That can happen because classification is not a complex task in this dataset and the presence of some specific words say a lot about the actual class of a document [Polo et al. 2021]. As this experiment was not intended to show which model performs best in different scenarios, but only give an example of application, we point that direction as a possible future path to go. Future work can be done in order to better understand how our models compare to alternatives in different tasks and datasets.

8. Conclusion

In this work, we presented and made available pre-trained language models, functions and demonstrations specific for the Brazilian legal language and field. Our main objective is to catalyze the use of natural language processing tools for legal texts analysis by the Brazilian industry, government and academia, providing the needed tools and accessible material.

References

- Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Braz, F. A., da Silva, N. C., de Campos, T. E., Chaves, F. B. S., Ferreira, M. H., Inazawa, P. H., Coelho, V. H., Sukiennik, B. P., de Almeida, A. P. G. S., Vidal, F. B., et al. (2018). Document classification using a bi-lstm to unclog brazil's supreme court. *arXiv preprint arXiv:1811.11569*.
- da Silva, N. C., Braz, F., Gusmão, D., Chaves, F., Mendes, D., Bezerra, D., Ziegler, G., Horinouchi, L., Ferreira, M., Inazawam, P., et al. (2018). Document type classification for brazil's supreme court using a convolutional neural network.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Hartmann, N., Fonseca, E., Shulby, C., Treviso, M., Rodrigues, J., and Aluisio, S. (2017). Portuguese word embeddings: Evaluating on word analogies and natural language tasks. *arXiv preprint arXiv:1708.06025*.
- Le, Q. and Mikolov, T. (2014). Distributed representations of sentences and documents. In *International conference on machine learning*, pages 1188–1196. PMLR.
- Massoni, G. (2021). Análise de textos por meio de processos estocásticos na representação word2vec.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013a). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. (2013b). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Nguyen, T.-S., Nguyen, L.-M., Tojo, S., Satoh, K., and Shimazu, A. (2018). Recurrent neural network-based models for recognizing requisite and effectuation parts in legal texts. *Artificial Intelligence and Law*, 26(2):169–199.
- Polo, F. M., Ciochetti, I., and Bertolo, E. (2021). Predicting legal proceedings status: approaches based on sequential text data. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law*, pages 264–265.
- Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A. V., and Gulin, A. (2017). Catboost: unbiased boosting with categorical features. *arXiv preprint arXiv:1706.09516*.
- Souza, F., Nogueira, R., and Lotufo, R. (2020). BERTimbau: pretrained BERT models for Brazilian Portuguese. In *9th Brazilian Conference on Intelligent Systems, BRACIS, Rio Grande do Sul, Brazil, October 20-23 (to appear)*.
- Sulea, O.-M., Zampieri, M., Malmasi, S., Vela, M., Dinu, L. P., and Van Genabith, J. (2017). Exploring the use of text classification in the legal domain. *arXiv preprint arXiv:1710.09306*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.