JurisBERT: A New Approach that Converts a Classification Corpus into an STS One

Charles F. O. Viegas[1], Bruno C. Costa[1], and Renato P. Ishii[2(✉)]

[1] Alfaneo, Goias, 405, 79020-100 Campo Grande, MS, Brazil
{charles.viegas,bruno.catais}@alfaneo.ai
[2] Federal University of Mato Grosso do Sul, Campo Grande, MS, Brazil
renato.ishii@ufms.br
https://alfaneo.ai, https://www.facom.ufms.br/

Abstract. We propose in this work a new approach that aims to transform a classification corpus into an STS (Semantic Textual Similarity) one. In that sense, we use BERT (Bidirectional Encoder Representations from Transformers) to validate our hypothesis, i.e., a multi-level classification dataset can be converted into an STS dataset which improves the fine-tuning step and evidences the proposed corpus. Also, in our approach, we trained from scratching a BERT model considering the legal texts, called JurisBert which reveals a considered improvement in fastness and precision, and it requires less computational resources than other approaches. JurisBERT uses the concept of sub-language, i.e., a model pre-trained in a language (Brazilian Portuguese) passes through refining (fine-tuning) to better attend to a specific domain, in our case, the legal field. JurisBERT uses 24k pairs of `ementas` with degrees of similarity varying from 0 to 3. We got this data from search mechanisms available on the court websites to validate the model with real-world data. Our experiments showed JurisBERT is better than other models such as multilingual BERT and BERTimbau with 3.30% better precision (F_1), 5 times reduced training time, and using accessible hardware, i.e., low-cost GPGPU architecture. The source code is available at https://github.com/alfaneo-ai/brazilian-legal-text-dataset and the model is here: https://huggingface.co/alfaneo.

Keywords: Semantic Textual Similarity · Retrieving Legal Precedents · Sentence Embedding · Bert

1 Introduction

Searching legal precedents is very important for legal professionals. They use it as a means for either supporting and strengthening their points or exposing opposing arguments. In Brazil, data from the *Conselho Nacional de Justiça*[1] [6]

[1] The Conselho Nacional de Justiça is a public institution that aims to help the Brazilian judiciary. It maintains administrative and procedural control and transparency.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2023
O. Gervasi et al. (Eds.): ICCSA 2023, LNCS 13956, pp. 349–365, 2023.
https://doi.org/10.1007/978-3-031-36805-9_24

shows a significant growth of legal proceedings, confirming that the Brazilian Judiciary System is overly congested, with a big amount of workload, and with an annual influx of millions of proceedings. In this scenario, an approach for efficiently retrieving precedents is very relevant for the Brazilian legal area.

In this context, several information-retrieving applications are using methods for evaluating semantic similarities, a process that is in the Natural Language Processing (NLP) field and involves determining the similarity between two text segments. Recently, models based on Transformers [27] networks and big unlabeled datasets (ex: BERT [8] and RoBERTa [17]) are raising the bar in a lot of NLP tasks, including evaluating the semantic textual similarity. Among some proposed approaches that stand out is Sentence BERT (sBERT) [24], which puts forward a change in the pre-trained BERT network and uses Siamese and triple network structure. This is used to derive semantically relevant sentence embedding, which can be compared by cosine similarity. This approach combines sentence embedding with indexing techniques such as FAISS [12]. It can quickly deal with outstanding amounts of data without losing the precision of transformers models.

However, most of the studies on transformers models are focused on the English language. This presents a challenge to bring technological advances to other languages, like Brazilian Portuguese. Even with popular models, such as BERT [8], having multi-language versions, models trained specifically with Brazilian Portuguese beat them, as shown by BERTimbau [26]. Besides, several works have revealed better results when pre-trained with domain-specialized corpus [3,21]. The BERTLaw study [19] shows that pre-training BERT with legal field specialized vocabulary has better results than using BERT Base. It exposes significant differences in unit vectors and that the intersection between both vocabularies is lower than half of the total vocabulary of each model separated. These differences directly affect the interpretation of a text segment by each model.

In order to validate our approach, we propose a methodology divided into 3 steps: 1) pre-training BERT for Masked Language Modeling (MLM) starting from scratch and using legal field-specific texts, containing laws, decisions, and court votes, besides several legal treatises of the Brazilian law; 2) experiments with the sBERT network using as base pre-trained BERT networks, multilingual BERT (mBERT) and BERTimbau; and 3) fine-tuning sBERT models over dataset generated by our proposed approach, which was prepared for evaluating the effectiveness of experiments and the similarities between the `ementas` of `acordãos`, these concepts will be explained at the Sect. 5. The dataset developed in our research has 24k pairs of `ementas` with a degree of similarity ranging between 0 and 3 got from court websites.[2].

JurisBERT got 3.30% F_1 more than BERTimbau with lower hardware cost for pre-training, confirming the hypothesis that a pre-trained BERT model for a specific domain has superior performance. We also noticed that sBERT with BERTimbau outperforms mBERT, validating that models with specific lan-

[2] STF, STJ, TJRJ, TJMS.

guages do better than the multilingual ones. Also, we could not find other public works of domain-specialized corpus for evaluating text similarity in the Brazilian legal field. In that case, we believe that our research is the first to provide this data publicly and that our contributions will help the development of new studies in this field.

We organized this paper in the following sections: in Sect. 2, we present the main concepts used in the other sections; in Sect. 3, we discuss other similar techniques that inspired us; in Sect. 4, we describe the steps of pre-training and dataset construction; in Sect. 5, we cover the fine-tuning of the models and discuss the results; finally, in Sect. 6, we present the main contributions of our research and point out future works.

2 Background

2.1 Legal Precedent Retrieval

Legal precedents are used to substantiate arguments by attorneys and judges. They use it to reinforce the justification of their decisions [29]. In Brazil, after the Constitution of 1998, precedents became more important since they started to have binding force over decisions made by the Brazilian Supreme Court [13]. For that reason, Brazilian courts have to provide public access to all its decisions over judged proceedings (except the ones that are classified as confidential). However, as the number of proceedings is big and keeps growing every year, more efficient solutions to retrieve precedents are in very high demand.

Worldwide, the retrieval of legal precedents is a very popular theme in the literature, especially the techniques for exploring semantic retrieval. They assist the better understanding of concepts related to contexts and the treatment of linguistic phenomenons, which affects the quality of the retrieval.

In Brazil, the main document used as precedent is named `acordão`. It shows the decisions made by the judging court. Even though it does not have a standard format, most of the time it has the following sections: identification of the concerned parties, the judge who wrote the opinion, the discussed objects, the given facts, the court votes, and the `ementa`, which is similar to the syllabus in the United States law. In Fig. 1 we show an example of `ementa`. We can see a standard in the writing. In the superior part, the text is written in capital letters and in entry, while in the other parts, the text is written in enumerated paragraphs.

Empirically, we see a lot of legal professionals using only the `ementa` section of the `acordão` to decide which precedents to choose. This probably happens because the `ementa` summarizes the decision and can be enough to understand the whole `acordão`. Besides, it is very common to find in lawsuits the full transcription of the `ementa` used to reference the precedent. Such observations helped us make the choice to use the `ementa` as the source material for similarity comparison, which is the goal of our study.

AGRAVO INTERNO – ART. 1.030, § 2º, CPC – ACÓRDÃO ESTADUAL QUE COINCIDE COM A ORIENTAÇÃO FIRMADA PELO SUPERIOR TRIBUNAL DE JUSTIÇA EM SEDE DE RECURSO REPRESENTATIVO DE CONTROVÉRSIA – LIMITAÇÃO DOS JUROS REMUNERATÓRIOS À TAXA MÉDIA DO MERCADO SOMENTE SE VERIFICADA ABUSIVIDADE – CAPITALIZAÇÃO MENSAL DOS JUROS PERMITIDA – EMBARGOS PROTELATÓRIOS – MULTA APLICADA – RECURSO IMPROVIDO. 1–As questões de direito enfrentadas e decididas nos recursos representativos da controvérsia guardam plena identidade ao posicionamento do Tribunal de Origem, pois somente será aplicável a limitação dos juros à taxa média do mercado em caso de comprovada desvantagem ao consumidor, demonstrando abusividade do fornecedor. 2–A capitalização mensal de juros é permitida nos contratos bancários desde que expressamente pactuada e celebrada após após 31.3.2000, ou que haja previsão de taxa de juros anual superior ao décuplo da mensal. 3–Recurso improvido.

Fig. 1. Example of a `ementa` written in Brazilian Portuguese. This text is used to compare similarities.

2.2 Semantic Textual Similarity (STS)

To make precedent-retrieving systems is necessary to use semantic similarity comparison techniques. In NLP, is proposed to use the STS task, which can be considered a regression task, to calculate the *score* that represents the similarity between two sentences. Given a collection of sentences and two sentences, the *score* will have higher values when the similarity is higher and lower values when the similarity is lower. In this area, the *International Workshop on Semantic Evaluation* (SemEval) [1] stands out promoting a series of research on the NLP field to advance the semantic analysis and the creation of high-quality datasets.

In Brazil, the first collection of public data that included semantic similarity between sentences in Portuguese was the ASSIN [9]. Years later, the ASSIN 2 [22] suggested a new data collection based on the SICK-BR [23] collection. However, neither of those collections specializes in legal texts.

2.3 BERT

BERT or *Bidirectional Encoder Representations from Transformers* is a language model used to pre-train deep bidirectional representations from unlabeled texts. It uses a bidirectional approach to model the context to the left and right of the entered sequence tokens. As a result, the pre-trained BERT model is adjusted with only one additional out layer to make models for NLP downstream tasks.

In Fig. 2 we show the relations between BERT versions with pre-training and fine-tuning. The pre-trained version is the base for the fine-tuning versions that are adjusted to perform downstream tasks (ex: STS, Entity Named Recognition, Text classification, etc.). The pre-training uses unlabeled data, while the fine-tuning uses labeled ones.

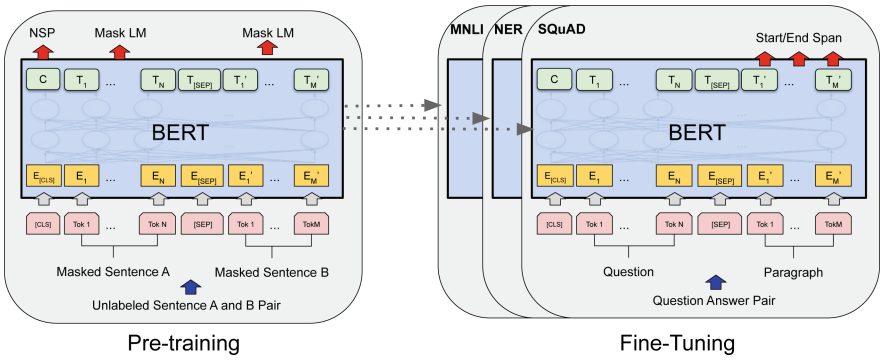

Fig. 2. BERT training process [8].

In the original BERT paper, there are two goals during pre-training: Masked Language Modeling (MLM) and Next Sentence Prediction (NSP). Within MLM, the number of random tokens in the input sequence is replaced by the special tokens [MASK] that are predicted using cross-entropy loss, 15% of the input tokens are evenly selected for possible replacements, of these tokens, 80% are in fact replaced by the token [MASK], 10% are unchanged and the remaining 10% are replaced by another random token of the vocabulary. In NSP, is used a binary classification loss to predict if two text segments follow one another in the original text. Positive examples are made with consecutive sentences and negative ones by paring text segments of different documents, both are created in equal proportions.

2.4 Transformers

The BERT architecture is based on the encoder part of the transformers [27] network architecture, which is considered a neural network of encoder-decoder type. The transformer network does not use either Recurrent Neural Networks structures or Convolution ones. Its main characteristics are being capable of reading sequential entries (from the left to the right or from the right to the left) in a single time. This characteristic allows for considering contexts from the right and the left. Also, it promotes better parallelization and requires less training time.

The transformer network is composed of an encoder that receives as input a sequence of words and transforms it into a vector sequence (internal representation). Next, a decoder, out of the internal representation, makes a sequence

of words, one by one. To learn more about transformer architecture, read the original paper [27].

2.5 Sentence BERT

Even though BERT models have been raising the bar in evaluating semantic similarity, they use a cross-encoder where both text sentences are sent to the transformer network. This produces a huge computational cost. For example, it would be necessary roughly 50 million in computational inferences (65 processing hours) to find the most similar sentences in a collection of 10,000. This turns BERT into an unviable option for information retrieval systems.

A way to address this type of problem is to map each sentence in a vector space, where semantic similar ones lay close together. For this, fixed sentence embeddings can be gotten by sending individual sentences to the BERT network. The most used technique gets the fixed vector embeddings from the average of values generated in the output layer of BERT (known as BERT embeddings) or using the output of the first token (or [CLS] token), though both techniques have lower performances than older ones like GloVe embeddings [20].

To fix this problem, sBERT was developed. It uses Siamese network architecture, which means using two identical networks with shared weights. As shown in Fig. 3, sBERT allows for deriving fixed-size vectors that, using a similarity measurement (ex: cosine similarity and Manhattan distance), can calculate the similarity between two sentences. To make sure that the generated embeddings have fixed sizes, there is a pooling[3] operation on the BERT output.

The sBERT model is computationally efficient. In the previously discussed example, the authors say that sBERT can reduce the computational cost to find the most similar pair of sentences, from 65 hours to approximately 5 seconds.

3 Related Works

The sBERT is basically a sentence embedding technique, a field with several studies and methods proposed. Unsupervised methods based on encoder-decoder techniques look to be dominating recent research. Skip-Thought [16] proposes a model with encoder-decoder architecture that uses neighbor sentences to codify the embedding. However, this model needs a corpus made with continuous texts for training. The Universal Sentence Encoder [4] suggests a general model that uses the encoder part of the transformer network, it utilizes the attention mechanism to compute the words that are sensitive to the context, that way, getting the sentence embedding from the average of the internal state of the codified tokens. It is also suggested that transferring knowledge at the level of sentences is more efficient than at the level of words. More recently, PromptBERT [11] applies prompts to reduce biases found in the internal state of the tokens for making the original BERT layers more effective.

[3] It is an operation that reduces the dimensionality of data by applying an aggregation of type max average.

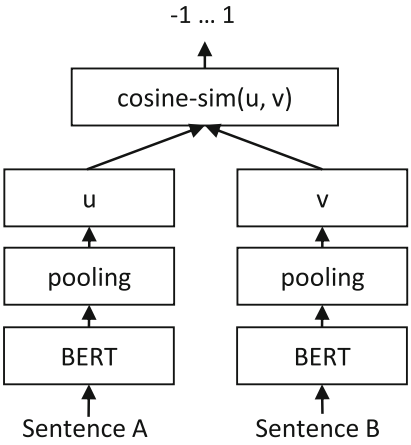

Fig. 3. sBERT architecture [24].

In [30], the authors conclude transformers model pre-trained with domain-specific contexts performs better than general models. This suggests that the best approach is merging both general and domain-specific models, by continuing pre-training using the last checkpoint of the general model in the domain-specific corpus. The corpus size was compared and concluded that it makes little difference between themselves.

Other works propose approaches that aim to transpose barriers from human annotation, which motivated us to develop our method that converts a classification dataset to STS to eliminate the annotation step. Meshram and Anand Kumar [18] propose a deep contextual long semantic textual similarity network. Deep contextual mechanisms for collecting high-level semantic knowledge are used in the LSTM network, and such a model performs better than human annotation. Keskar et al. [14] propose a unified span-extraction approach that allows combinations of question-answering and text classification datasets in intermediate-task training that outperform using only one or the other. Banerjee et al. [2] provide an unsupervised modeling approach that allows for a more flexible representation of text embeddings, transforming such a problem into an optimization problem over the spherical manifold. Results show improvements in document classification, document clustering, and semantic textual similarity benchmark tests.

4 JurisBERT

In this work, JurisBERT is the model trained from scratch using the BERT model and considering the legal texts. Our proposed approach aims to transform a classification corpus into an STS one. To validate our hypothesis, a multi-level classification dataset was converted into an STS dataset which improves the

fine-tuning step and evidences the proposed corpus. The JurisBERT is a specific model for semantic textual similarity of the `ementas` of `acordãos` with an sBERT network. Our study puts forward pre-training BERT with Brazilian legal field domain-specific texts. For this, we constructed two corpora (for training and fine-tuning), discussed in more detail in Sects. 4.1 and 5. For evaluating the experiments we have two steps: first, we iterate over each model with all the analyzed `ementas` to create the embedding vectors[4]. Second, we calculate the degree of similarity between each pair of embeddings through the method of cosine similarity (as suggested by the authors of sBERT [5,24]). The score got from the cosine similarity can only have values between 0 and 1, the closer to 1, the higher the similarity, and the closer to 0, the lower the similarity. We stabilized a threshold number to optimize the division into two groups. That way, the score values over the threshold are similar and the ones under are not similar.

These two models chosen for comparison were pre-trained with general domain texts and with a much larger corpus than JurisBERT. The mBERT was pre-trained with 104 different languages, provided by the authors of BERT. BERTimbau, in turn, was pre-trained in Brazilian Portuguese with data from BrWac [28], which is a corpus got from Brazilian websites. We made two variations of pre-training: one from scratch, called JurisBERT, and the other, coming from the BERTimbau checkpoint, we apply further pre-training with legal texts called BERTimbau further. The goal is to determine if there are significant differences between both methods.

All BERT models are based on the $\text{BERT}_{\text{BASE}}$ architecture ($L = 12$, $H = 768$, $A = 12$, Total parameters $= 110M$). We chose this architecture because, even though it is not the most performative, it has a lightweight model that is more suitable for our hardware.

4.1 Pre-training Corpus

To pre-train JurisBERT, we had to create our own corpus, because we could not find one done with the amount of data and variety needed. Then, the first step was developing web scrappers to read and retrieve public documents from many sources, for example, court websites, public agencies, and universities. During the construction of this corpus, we aimed for the primary sources of the legal field, such as laws, precedents, treatises, analogies, general principles of law, and equity [7]. So, among the retrieved documents are laws and federal decrees, `súmulas`[5], decisions, `acordãos` and court votes, besides treatises of different legal fields. In Fig. 4, we show a chart with the participation of each type of legal source. The predominating of laws happens because of the plentifulness of this kind of document on governmental websites.

[4] It is a dense vector of floating points that aims to capture the semantic of the text in the vector space.

[5] The `súmulas` summarizes the dominant precedent of a given court.

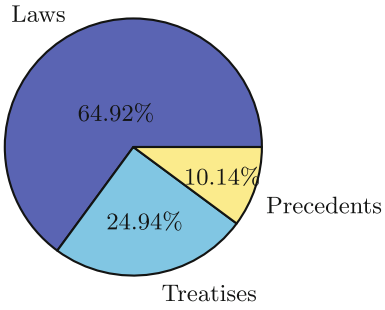

Fig. 4. Participation of each type of legal source in the pre-training corpus.

After retrieving the documents, we pre-processed them to remove special characters, excesses of spaces, and blank lines, though we maintained the accentuation and the letter casing because the networks we used are case-sensitive. Also, we split the documents into paragraphs and dispose of the ones with less than 10 tokens and the duplicated ones. We made this choice, intending to make the paragraphs more similar to the average size of the `ementa`. The result we got was 1.5 million of sentences(paragraphs), 99% of those with less than 384 tokens, as shown in Fig. 5, repeating a proportion like the one in the corpus of the *fine-tuning*, detailed in Sect. 5. After the pre-processing, the corpus reached 410 MB of raw text, a significantly lower number than the ones of BERTimbau (17 GB) and RoBERTa (160 GB). Then, we divided the corpus into two parts: training (95%) and evaluation (5%).

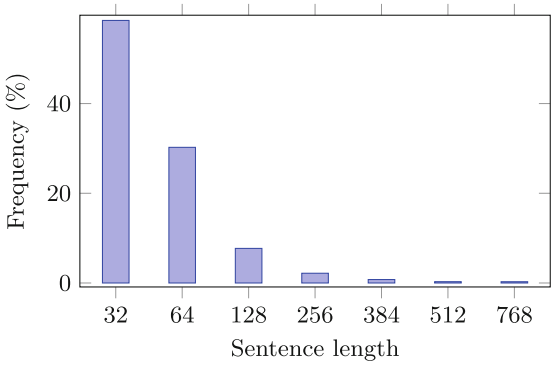

Fig. 5. Frequency of each sentence length.

4.2 Vocabulary Generation

For training JurisBERT, we generated an uncased vocabulary of 30k units of sub-words using the training corpus. The chosen tokenizer was WordPiece [25] of the huggingface library.

4.3 Pre-training

We used only the MLM goal in the pre-training since recent papers [17] have suggested that the NSP goal is not effective. During training, we optimized JurisBERT with Adam [15] using the following parameters: $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-6$ e L_2 weight decay of 0.01. The learning rate in the first $10,000$ steps with the apex of $1e-4$ and linear decay. We also used a *dropout* of 0.1 in all levels, alongside GELU's [10] activation function. The model was pre-trained with 20 epochs ($220k$ steps, 5 times smaller than the BERTimbau), with a batch size of 128 and sequences with a maximum length of 384 tokens. We used two `NVIDIA GeForce RTX 3080` with 12 GB `GDDR6X` each. The training time of models JurisBERT and BERTimbau further was 7 and 21 days, respectively.

5 Evaluating Semantic Textual Similarity

We evaluated the performance of the models after their fine-tuning. This step is essential to confirm our hypothesis, in other words, if the dataset generated by fine-tuning step provides better results in STS, this proves that our approach is adequate to construct qualifiers STS datasets. For this, we also needed to develop our own dataset composed of many paired sentences showing whether they were similar. The term sentence, in our study, means the content of the `ementa` of an `acordão`, which usually is made of a few paragraphs.

We constructed our dataset through search mechanisms available on the websites of the courts: STF, STJ, TJRJ, and TJMS. There, the courts provide one or more `acordãos` for each theme considered stable, which means that it has a standard understanding in the legal context. They also provide notorious cases such as racial quotas, the use of embryonic stem cells, and party infidelity.

In particular, STF's search mechanism has an interesting characteristic to define degrees of similarity, it is organized hierarchically at three levels of grouping. The first level is divided by fields of law: Administrative, Civil, Constitutional, Electoral, Criminal, Retirement, Civil Procedural, Criminal Procedural, and Tax. The second themes, for example, in the field of Administrative Law, the divided themes are Public Job Applications, Pharmaceutical Assistance Programs, Liability of Public Administration, and Public Sector Employees. The third and final level is divided by legal discussions, for example, in the Public Sector Employees theme we got the following discussions: Teacher's Special Retirement, Public Employee Payment Discount for Striking, Vested Right of Probationary Period, Judgment Deadline for Legality of Retirement.

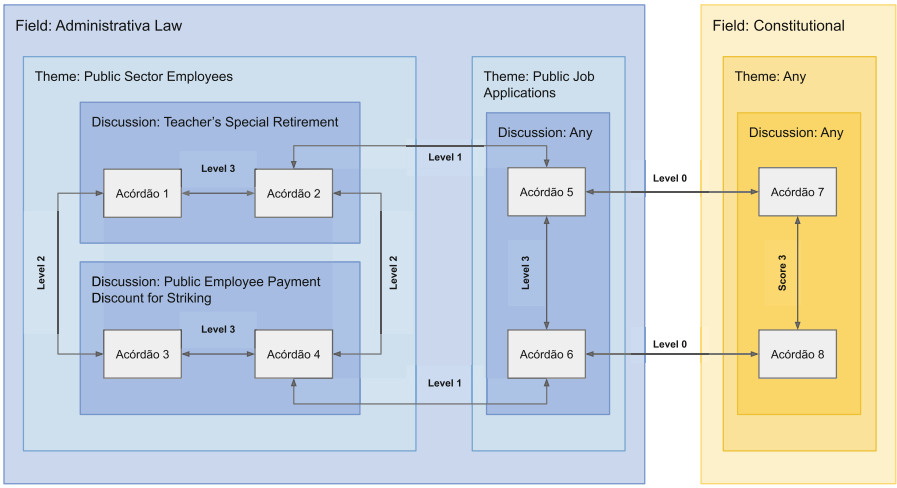

Fig. 6. Example of the method used to define the similarity levels during STS dataset generation.

So, to STF's search mechanism, we applied a scale of similarity between `acordãos` considering the group hierarchy, because `acordãos` inside the same discussion is more similar than `acordãos` of different themes and fields. Thus, we automatically annotated `acordãos` of the same discussion with a similarity of 3, `acordãos` of different discussion but of the same theme with a similarity of 2, `acordãos` of different themes but of the same field with similarity of 1 and `acordãos` of different fields with similarity of 0 as shown in Fig. 6.

Applying this strategy with STF, we got a total of 24,926 pairs of `acordãos` applying a combination two by two. Table 1 shows an example of `acordãos` pairs generated using the same examples shown in Fig. 6.

We dedicate 80% of generated pairs to training and 20% to evaluation and the other search mechanisms (STJ, TJRJ, and TJMS) got only a single level of grouping, we only used those for testing and comparing models. The strategy used for automatically annotating these searches was considering `acordãos` of the same group as similarity of 1 and the rest as similarity of 0. Therefore, we generated 19,027 pairs from the STJ, 8,626 from the TJRJ, and 6,305 from the TJMS.

5.1 Fine-tuning

In the beginning, we had defined the maximum sequence length after analyzing the corpus; it showed that 91.33% of the `ementa` texts had lengths lower than 384 tokens. However, we trained models with low epochs varying only in the maximum sequence length to validate this hypothesis. As shown in Table 2, the value 384 had better results.

Table 1. Example pairs of `acordãos` from the generated STS dataset.

Sentence 1	Sentence 2	Similarity level
Acordão 1	Acordão 2	3
Acordão 1	Acordão 3	2
Acordão 1	Acordão 4	2
Acordão 1	Acordão 5	1
Acordão 1	Acordão 6	1
Acordão 1	Acordão 7	0
Acordão 1	Acordão 8	0
Acordão 2	Acordão 3	2
Acordão 2	Acordão 4	2
Acordão 2	Acordão 5	1
Acordão 2	Acordão 6	1
Acordão 2	Acordão 7	0
Acordão 2	Acordão 8	0
Acordão 3	Acordão 4	3
Acordão 3	Acordão 5	1
Acordão 3	Acordão 6	1
Acordão 3	Acordão 7	0
Acordão 3	Acordão 8	0
Acordão 4	Acordão 5	1
Acordão 4	Acordão 6	1
Acordão 4	Acordão 7	0
Acordão 4	Acordão 8	0
Acordão 5	Acordão 6	3
Acordão 5	Acordão 7	0
Acordão 5	Acordão 8	0
Acordão 6	Acordão 7	0
Acordão 6	Acordão 8	0
Acordão 7	Acordão 8	3

After having defined the maximum sequence length, we made the fine-tuning training with 3 epochs, hyper-parameters, batch size, and maximum sequence length of 8 and 384, respectively. The cosine similarity loss was used. The training took roughly 3 hours using the same hardware as described in Sect. 4.3. To measure and compare our experiments, we considered only the checkpoint of the models that got the best performance during the training epochs.

Table 2. Experiments varying the maximum sequence length.

Sequence length	$F_1(\%)$
64	71.30
128	76.72
256	81.93
384	**83.27**
512	82.70

5.2 Discussions

We evaluated the models with $F_1(\%)$ metric, as shown in Table 3. In general, we can see that all the models that went through fine-tuning performed better than the ones that did not. These results suggest that our approach to constructing the dataset and making the fine-tuning was, in fact, effective. Afterward, these results demonstrated that our proposed approach could be useful in other areas not only in the law field. We believe that the multi-level characteristic from the classification dataset was crucial to infer the distance degrees for the sentence pairs.

Also, we can notice that models pre-trained with specific texts of a sub-language, in other words, texts specialized in a specific domain, performed better than the others. Models BERTimbau and JurisBERT showed better general results. This means there are advantages in training from scratch and in doing more training on pre-existing models. Our training approach proved effective even with a corpus size 42 times smaller and with 5 times fewer pre-training steps than other methods.

In that regard, we can say that our work addresses an important problem in NLP, dealing with domain-specific topics (healthcare, law, and others). These have plenty of data pulverized through different mediums, but few of those are annotated for a specific task. Even without fine-tuning, our approach proved to be a viable option, beating even fine-tuned mBERT and BERTimbau. These results infer it is also possible to have good performances in other downstream tasks.

Our pre-training had a hardware cost lower than the other models, the training time was lower and the graphics card used was inferior to those used in BERTimbau and mBERT. The `P100` graphics card, by `NVIDIA`, is more expensive than the `RTX` used in our study. So, we showed it is technically and financially viable to replicate these experiments in other sub-language domains, considering the accessible price of current `RTX` models. For instance, the price of a graphic card used in our study is about $\$1,000$ dollars.

Table 3. Comparing F_1(%) results.

Model	TJMS	TJRJ	STJ	Mean
No Fine-tuning				
mBERT	65.05	63.78	30.39	53.08
BERTimbau	71.20	63.88	35.55	56.88
BERTimbau further	75.48	64.06	41.75	60.43
JurisBERT (ours)	**80.25**	**73.21**	**52.00**	**68.49**
Fine-tuning approaches				
mBERT	69.51	64.21	39.49	57.73
BERTimbau	81.77	71.34	50.75	67.95
BERTimbau further	84.39	73.46	47.37	68.41
JurisBERT (ours)	**85.16**	**77.65**	**50.95**	**71.25**

6 Conclusion

This work's main contribution is a new approach that transforms a classification corpus into an STS one. We use BERT to validate our hypothesis, i.e., a multi-level classification dataset can be converted into an STS dataset which improves the fine-tuning step and evidences the proposed corpus.

Moreover, we contribute to creating a corpus for unsupervised pre-training, including many laws, treatises, and decisions of several branches of Brazilian law. Similarly, the pre-training from scratching is 5 times faster than other approaches in training time, and it can be deployed in a low-cost environment considering a usual graphics card, i.e., with a price of less than $1,000. Using our dataset alongside the data from `acordãos` contributes to many other NLP tasks like clustering and topic modeling. Further, the source code of the web scrappers and parsers we used in the construction of the datasets is available in a public repository[6].

Finally, JurisBERT got 3.30% F_1 more than BERTimbau with lower hardware cost for pre-training, confirming the hypothesis that a pre-trained BERT model for a specific domain has superior performance. We also noticed that sBERT with BERTimbau outperforms mBERT, validating that models with specific languages do better than the multilingual ones. Also, we could not find other public works of domain-specialized corpus for evaluating text similarity in the Brazilian legal field. In that case, we believe that our research is the first to provide this data publicly and that our contributions will help the development of new studies in this field.

[6] Proposed dataset and web scrappers are available here: https://github.com/alfaneo-ai/brazilian-legal-text-dataset, and the models, here: https://huggingface.co/alfaneo.

Acknowledgments. We thank the support of the UFMS (Universidade Federal de Mato Grosso do Sul), FUNDECT, and Finep. We also thank the support of the INCT of the Future Internet for Smart Cities funded by CNPq, proc. 465446/2014-0, Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001, and FAPESP, proc. 2014/50937-1 and 2015/24485-9.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of FUNDECT, Finep, FAPESP, CAPES, and CNPq.

References

1. Agirre, E., Cer, D., Diab, M., Gonzalez-Agirre, A.: SemEval-2012 task 6: A pilot on semantic textual similarity. In: *SEM 2012: The First Joint Conference on Lexical and Computational Semantics - Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012), pp. 385–393. Association for Computational Linguistics, Montréal, Canada (7–8 Jun 2012), https://aclanthology.org/S12-1051
2. Banerjee, S., Mishra, B., Jawanpuria, P., Shrivastava, M.: Generalised spherical text embedding (2022)
3. Beltagy, I., Lo, K., Cohan, A.: Scibert: A pretrained language model for scientific text. In: EMNLP (2019)
4. Cer, D., et al.: Universal sentence encoder for English. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. pp. 169–174. Association for Computational Linguistics, Brussels, Belgium (Nov 2018). https://doi.org/10.18653/v1/D18-2029, https://aclanthology.org/D18-2029
5. Choi, H., Kim, J., Joe, S., Gwon, Y.: Evaluation of bert and albert sentence embedding performance on downstream nlp tasks (2021). https://doi.org/10.48550/ARXIV.2101.10642, https://arxiv.org/abs/2101.10642
6. CNJ: Justiça em números 2020: ano-base 2019. Tech. rep., Conselho Nacional de Justiça (2020), https://www.cnj.jus.br/pesquisas-judiciarias/justica-em-numeros/
7. Cunha, A.S.: Introdução ao estudo do direito. Saraiva (2012)
8. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. CoRR abs/1810.04805 (2018). http://arxiv.org/abs/1810.04805
9. Fonseca, E.R., Borges dos Santos, L., Criscuolo, M., Aluísio, S.M.: Visão geral da avaliação de similaridade semântica e inferência textual. Linguamática 8(2), 3–13 (Dez 2016). https://linguamatica.com/index.php/linguamatica/article/view/v8n2-1
10. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus) (2016). https://doi.org/10.48550/ARXIV.1606.08415, https://arxiv.org/abs/1606.08415
11. Jiang, T., et al.: Promptbert: Improving bert sentence embeddings with prompts (2022). https://doi.org/10.48550/ARXIV.2201.04337, https://arxiv.org/abs/2201.04337
12. Johnson, J., Douze, M., Jégou, H.: Billion-scale similarity search with gpus (2017). https://doi.org/10.48550/ARXIV.1702.08734, https://arxiv.org/abs/1702.08734
13. Júnior, H.: Curso de direito processual civil, vol. I. Editora Forense (2019)
14. Keskar, N.S., McCann, B., Xiong, C., Socher, R.: Unifying question answering, text classification, and regression via span extraction (2019)

15. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2014). https://doi.org/10.48550/ARXIV.1412.6980, https://arxiv.org/abs/1412.6980
16. Kiros, R., Zhu, Y., Salakhutdinov, R., Zemel, R.S., Torralba, A., Urtasun, R., Fidler, S.: Skip-thought vectors (2015). https://doi.org/10.48550/ARXIV.1506.06726, https://arxiv.org/abs/1506.06726
17. Liu, Y., et al.: Roberta: A robustly optimized bert pretraining approach (2019). https://doi.org/10.48550/ARXIV.1907.11692, https://arxiv.org/abs/1907.11692
18. Meshram, S., Anand Kumar, M.: Long short-term memory network for learning sentences similarity using deep contextual embeddings. International Journal of Information Technology, pp. 1633–1641 (2021)
19. Nguyen, H.T., Nguyen, L.M.: Sublanguage: A Serious Issue Affects Pretrained Models in Legal Domain. arXiv e-prints arXiv:2104.07782 (Apr 2021)
20. Pennington, J., Socher, R., Manning, C.: GloVe: Global vectors for word representation. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 1532–1543. Association for Computational Linguistics, Doha, Qatar (Oct 2014). https://doi.org/10.3115/v1/D14-1162, https://aclanthology.org/D14-1162
21. Rasmy, L., Xiang, Y., Xie, Z., Tao, C., Zhi, D.: Med-bert: pre-trained contextualized embeddings on large-scale structured electronic health records for disease prediction (2020). https://doi.org/10.48550/ARXIV.2005.12833, https://arxiv.org/abs/2005.12833
22. Real, L., Fonseca, E., Gonçalo Oliveira, H.: The assin 2 shared task: A quick overview. In: Quaresma, P., Vieira, R., Aluísio, S., Moniz, H., Batista, F., Gonçalves, T. (eds.) Computational Processing of the Portuguese Language, pp. 406–412. Springer International Publishing, Cham (2020)
23. Real, L., et al.: SICK-BR: A Portuguese Corpus for Inference: 13th International Conference, PROPOR 2018, Canela, Brazil, September 24–26, 2018, Proceedings, pp. 303–312. Springer, Cham (01 2018). https://doi.org/10.1007/978-3-319-99722-3_31
24. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks (2019). https://doi.org/10.48550/ARXIV.1908.10084, https://arxiv.org/abs/1908.10084
25. Song, X., Salcianu, A., Song, Y., Dopson, D., Zhou, D.: Fast wordpiece tokenization (2020). https://doi.org/10.48550/ARXIV.2012.15524, https://arxiv.org/abs/2012.15524
26. Souza, F., Nogueira, R., Lotufo, R.: Bertimbau: Pretrained bert models for brazilian portuguese. In: Intelligent Systems: 9th Brazilian Conference, BRACIS 2020, Rio Grande, Brazil, October 20–23, 2020, Proceedings, Part I. p. 403–417. Springer-Verlag, Berlin, Heidelberg (2020). https://doi.org/10.1007/978-3-030-61377-8_28
27. Vaswani, A., et al.: Attention is all you need (2017)
28. Wagner Filho, J.A., Wilkens, R., Idiart, M., Villavicencio, A.: The brwac corpus: A new open resource for brazilian portuguese. In: Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018) (2018)

29. Weber, R.: Intelligent jurisprudence research: A new concept. In: Proceedings of the 7th International Conference on Artificial Intelligence and Law. p. 164–172. ICAIL '99, Association for Computing Machinery, New York, NY, USA (1999). https://doi.org/10.1145/323706.323791, https://doi.org/10.1145/323706.323791
30. Zhang, G., Lillis, D., Nulty, P.: Can Domain Pre-training Help Interdisciplinary Researchers from Data Annotation Poverty? A Case Study of Legal Argument Mining with BERT-based Transformers. In: Proceedings of the Workshop on Natural Language Processing for Digital Humanities (NLP4DH), pp. 121–130. Association for Computational Linguistics (2021), https://rootroo.com/downloads/nlp4dh_proceedings_draft.pdf