Empirical analysis of Binding Precedent efficiency in the Brazilian Supreme Court via Similar Case Retrieval

Raphaël Tinarrage[1], Henrique Ennes[1,4], Lucas E. Resck[1],
Lucas T. Gomes[2], Jean R. Ponciano[3], Jorge Poco[1]

[1]EMAp, Fundação Getulio Vargas, Rio de Janeiro, Brazil
[2]Direito Rio, Fundação Getulio Vargas, Rio de Janeiro, Brazil
[3]ICMC, University of São Paulo, São Carlos, Brazil
[4]INRIA Université Côte d'Azur, Sophia Antipolis, France

July 25, 2024

Abstract   Binding precedents (*Súmulas Vinculantes*) constitute a juridical instrument unique to the Brazilian legal system and whose objectives include the protection of the Federal Supreme Court against repetitive demands. Studies of the effectiveness of these instruments in decreasing the Court's exposure to similar cases, however, indicate that they tend to fail in such a direction, with some of the binding precedents seemingly creating new demands. We empirically assess the legal impact of five binding precedents, 11, 14, 17, 26 and 37, at the highest court level through their effects on the legal subjects they address. This analysis is only possible through the comparison of the Court's ruling about the precedents' themes before they are created, which means that these decisions should be detected through techniques of Similar Case Retrieval. The contributions of this article are therefore twofold: on the mathematical side, we compare the uses of different methods of Natural Language Processing — TF-IDF, LSTM, BERT, and regex — for Similar Case Retrieval, whereas on the legal side, we contrast the inefficiency of these binding precedents with a set of hypotheses that may justify their repeated usage. We observe that the deep learning models performed significantly worse in the specific Similar Case Retrieval task and that the reasons for binding precedents to fail in responding to repetitive demand are heterogeneous and case-dependent, making it impossible to single out a specific cause.

Keywords   Similar Case Retrieval; Binding Precedents; Judicial Efficiency; Brazilian Supreme Court

MSC2020 codes   68T50 (Natural language processing); 68T07 (Deep learning)

1   Introduction

In recent years, the progress in Natural Language Processing (NLP) sparked significant interest in its application within the legal domain. Brazil, renowned for having the world's highest volume of legal cases, is no exception and has already witnessed several implementations of Machine Learning methods in its judiciary system [51]. Notably, these algorithms have the potential to assist legal professionals in finding cases similar to a given one, or a given legal topic — a task commonly referred to as *Similar Case Retrieval*. While these methods are supported by numerous mathematical studies and proofs of concept, one could argue that some of them lack empirical validation. This article seeks to bridge this gap by comparing various document retrieval methods and evaluating their results from a legal perspective.

In particular, this article explores the legal instrument of binding precedent (*Súmula Vinculante*, abbreviated BP), that emerged in Brazil in the 2004 judicial reform. Its purpose was to address the issue of an overwhelming number of cases with repetitive demands inundating the Brazilian Federal Supreme Court (*Supremo Tribunal Federal*, STF), which ideally should handle only a limited number of cases. This situation was partly due to the Brazilian legal system's civil law approach, where Supreme Court decisions don't serve as authority for lower courts, leading to continued case congestion. To mitigate this, Constitutional Amendment No. 45 introduced instruments inspired by the common law system, including binding precedents. They aimed at standardizing jurisprudence, providing normative force over lower instances and the broader public administration.

After its publication, one expects the binding precedent to be cited frequently, until the legal understanding of the courts on the subject calms down, resulting in a significant decrease in the number of citations. However, among the most cited binding precedents, this trend is by no means observed. On the contrary, they show steady growth, as illustrated in Fig. 1 for the five precedents of interest in this paper: 11, 14, 17, 26 and 37. They were chosen for their high number of citations, as well as the variety of legal topics they cover, from administrative to criminal law.

The aim of this article is to shed light on the reasons why these binding precedents have not led, as expected, to a reduction in repeated demands. To this end, we employ Similar Case Retrieval methods to trace the history of these precedents and quantify certain trends. This information is subsequently used to provide a legal analysis. In summary, this article includes two main mathematical and two juridical contributions:

1. The application and comparative analysis of classical algorithms for Similar Case Retrieval on a database of Brazilian legal documents (including TF-IDF-based models, LSTM, BERT and regex);

2. The outline of a methodology for assessing the impact of a law on jurisprudence, through time series of similar cases and features' correlations;

3. Application of the mentioned methodology to five binding precedents emitted by the Brazilian Federal Supreme Court, enabling an empirical study of the juridical mechanisms behind their inefficiency;

4. The identification of five main hypotheses explaining the large number of cases reaching the Supreme Court.
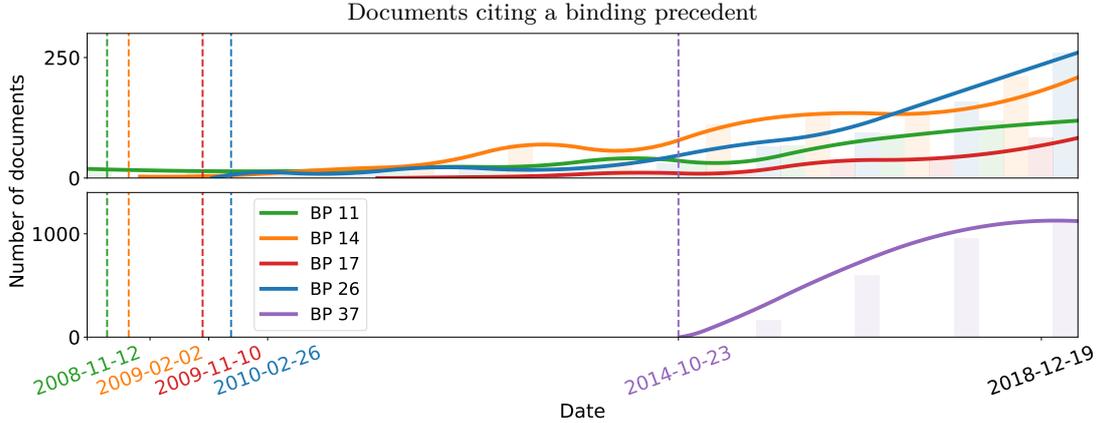

Figure 1: Histograms of the number of cases judged by the Federal Supreme Court citing Binding Precedents 11, 14, 17, 26 or 37, in our collection (Dataset #1). The bins have a length of one year, and the curves are obtained via quadratic spline interpolation. The dashed vertical lines represent the date of publication of each BP. They all exhibit an increasing trend.

Dataset and reproducibility

For our analysis, we will use a set of decisions of the Supreme Court produced between 1989 and 2018, collected and annotated by the project *Supremo em Números* [23]. In fact, these documents were scraped from the official STF website[1], where they can be found one by one. The collection of documents gathered by the project is, however, not publicly available, and has been kindly provided to us. More precisely, we will consider two subsets of this collection, detailed further in Section 3.1: Dataset #1, consisting of all decisions citing a binding precedent (29,743 documents), and Dataset #2, gathering all decisions belonging to the topics "administrative law", "criminal law" or "criminal procedure law" (615,262 documents). We point that the data gathered by *Supremo em Números* have already been used in several works [14, 13, 47]. In addition, other similar databases have been reported, such as the one maintained by the project *Victor* [20].

Overview

The remainder of this article is organized as follows. In Section 2, we set out the legal context behind the Brazilian's Supreme Court binding precedents, with a particular focus on the five precedents specifically studied in this article, and we give an overview of the mathematical literature surrounding Similar Case Retrieval. Datasets #1 and #2 are introduced in Section 3, as well as the models used throughout this article (TF-IDF, LSTM and BERT), that we train on the former dataset. The application of these models

---

[1]`https://portal.stf.jus.br/`

3

to the latter dataset is the focus of Section 4, where we first describe our methodology (in Section 4.1, summarized in Fig. 2), apply it to the five precedents under consideration (Sections 4.2 to 4.6), and gather our findings in a juridical discussion (Section 4.7). We conclude and address future works in Section 5.
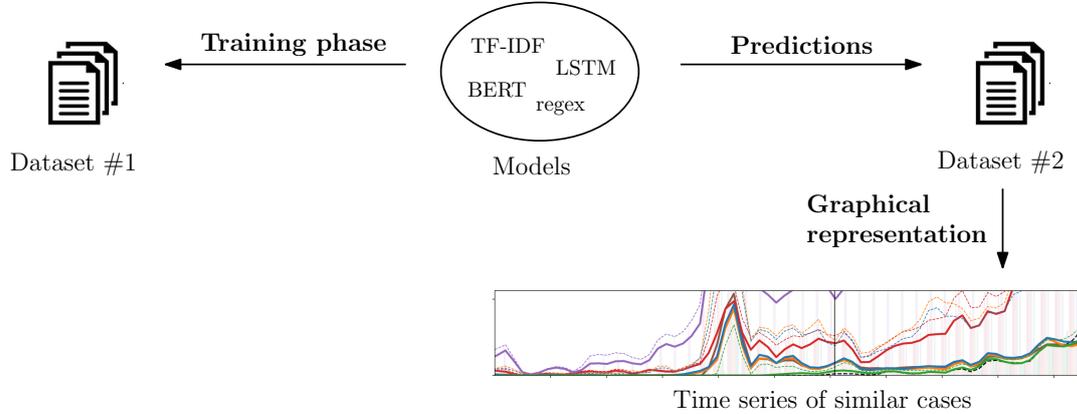

Figure 2: Schematic overview of the article. To understand the dynamics behind the use of a precedent, we train the models on an initial set of labeled documents, then apply these models to a larger set of data, and represent the results as a time series.

Abbreviations

Throughout this article, expressions in Portuguese will be written in italics, and may be followed by a translation in English in parenthesis. In addition, the following list contains the abbreviations that will be used, and also introduced throughout the text.

ADI   *Ação Direta de Inconstitucionalidade* (Direct Action for Unconstitutionality)

ARE   *Recurso Extraordinário com Agravo* (Aggravated Extraordinary Appeal)

BP   *Súmula Vinculante* (Binding Precedent)

HC   *Habeas Corpus* (Habeas Corpus)

Inq   *Inquérito* (Inquiry/Investigation)

LIME   Local Interpretable Model-agnostic Explanations

LSTM   Long Short Term Memory

NLP   Natural Language Processing

Rcl   *Reclamação* (Complaint/Appeal)

RE   *Recurso Extraordinário* (Extraordinary appeal)

regex   Regular Expression

STF   *Supremo Tribunal Federal* (Brazilian Federal Supreme Court)

STJ   *Superior Tribunal de Justiça* (Brazilian Superior Court of Justice)

SVM   Support Vector Machine

TF-IDF   Term frequency-inverse document frequency

2 Related works

2.1 Legal context

The BPs as a juridical tool

The type of precedent known in Brazil as *Súmula Vinculante* (Binding Precedent) emerged in the judiciary reform of 2004, through Constitutional Amendment No. 45 (EC 45/2004, Article 103-A[2]), to standardize decisions, based on the importation of the jurisdiction logic usual in common law [39, p. 827]. The introduction of these instruments can be understood as an attempt to unify decision-making, aiming to achieve equality and legal certainty, i.e., to avoid that identical cases are decided in different ways, violating the constitutional guarantee of equality before the law [8, p. 20]. In fact, the new Federal Constitution Article requires, for the creation of a BP, the fulfillment of three conditions: (i) dealing with a matter under current controversy among judicial bodies, or between them and the public administration; (ii) representing a risk of serious legal uncertainty; and (iii) being the subject of significant multiplication of processes on an identical legal issue. Thus, as a legal instrument, the BP seeks not only to guarantee equality in judicial decisions, but also to increase the efficiency of the Judiciary, by avoiding delays in jurisdictional intervention.

To date, 58 BPs have been published[3]. It should be noted that, in the literature, there is a considerable criticism of the BPs, whether in terms of their institution, their formulation, or their consequences. These issues will be explored in detail throughout the article. We refer the reader to the work of [3] and [43] for a general analysis of BPs.

The precedents 11, 14, 17, 26 and 37

Once a BP is published, is it expected that the subject ceases to be a matter of controversial interpretation. Consequently, the issue should also cease to generate processes with identical demands and, primarily, prevent these from continuing to be brought before the Higher Courts. It is therefore important to establish, empirically, the degree of efficiency of this type of instrument, here understood as the degree of reduction in similar cases reaching the final stage of appeal. In this article, we analyze the impact of the creation of five binding precedents, chosen as those generating many cases, but also for the diversity of legal topics they cover. Three of them belong to criminal law (11, 14 and 26), and the others to administrative law (17 and 37). We give in the list below a brief summary of their content and refer the reader to the corresponding section for a deeper explanation, as well as the results of our analysis.

BP 11, the 2008-11-12 (see Section 4.2): Determines the use of handcuffs acceptable only when some risk is anticipated, predicting disciplinary punishments to the responsible public agents and/or nullity of the penal process in the case of unjustified use.

BP 14, the 2009-02-02 (see Section 4.3): Grants to the investigated individual and their attorneys full access to all documented evidence in ongoing criminal investigations.

BP 17, the 2009-11-10 (see Section 4.4): In Brazil, public administration at all levels, when sentenced to pay debts resulting from juridical decisions, incorporate these

---

[2]https://www.planalto.gov.br/ccivil_03/constituicao/Constituicao.htm#art103a

[3]https://portal.stf.jus.br/textos/verTexto.asp?servico=jurisprudenciaSumulaVinculante

amounts into the public budget by the system of *precatórios* (Court Order Payments), in which orders of payment are issued to beneficiaries. The BP determines that no late payment interests should be applied to the potential delay between the issuing of orders of payment and the actual payment taking place.

BP 26, the 2010-02-26 (see Section 4.5): Determines the unconstitutionality of prohibiting convicts of "heinous crimes" (such as murder, rape, and drug dealing) to receive a reduction of the sentence from a fully closed to semi-open or open carceral system. Moreover, the BP authorizes judges to employ criminological exams to determine the feasibility of this reduction of the sentence.

BP 37, the 2014-10-23, (see Section 4.6): Establishes that the Judiciary cannot increase salaries of public servants based on the principle of isonomy (equality) without prior legislation to authorize such adjustments.

2.2 NLP for legal documents

Text embeddings

The literature surrounding NLP for legal document analysis is rich, reflecting the growing interest in leveraging computational techniques to navigate the complexities of legal texts. Well-studied topics in this area include the analysis of the precedent network [36, 14, 41], Named Entity Recognition [13], summarization of legal texts [28, 26] and prediction of judicial outcome [2, 4, 5]. In addition, of particular interest to us is the detection of similar documents, reviewed further in the next paragraph.

To tackle these problems, a standard technique consists of embedding (i.e., vectorizing) the documents, the most classical methods being bag-of-words, TF-IDF, and *n*-grams. Among recent techniques of words and documents embeddings, we can cite GLoVE [42], contextualized word representation [44], word2vec [12], doc2vec [33], Latent Dirichlet Allocation (LDA) [7], entities and relations-based embedding [72], Universal Sentence Encoder (USE) [10] and TextCNN [11]. Another alternative are the pre-trained language models, such as BERT [21] and its variants [75, 45, 29, 67].

In particular, word2vec is used by [24] to study appellate court modifications in Brazilian legal documents, that is, modifcations by the Supreme Court of the lower Court judge's decision. On the other hand, TF-IDF has been already used for the analysis of binding precedents, and has been reported to outperform other embeddings (such as doc2vec and USE) [47]. We will consider, in this article, TF-IDF embeddings coupled with several classifiers, as well as a recurrent neural network (LSTM) and a Large Language Model (BERT).

Automatic Similar Case Matching and Retrieval

At their most elementary level, although having a specific binding property, BPs are precedent and, as such, the task of searching for documents similar to the object of a BP, in content and language, may be modeled through textual proximity and/or classification techniques. The problem of identifying juridical decisions similar in a corpus of documents, known as automatic Similar Case Matching, has received significant attention in the literature. For example, a multitude of NLP embedding models have been compared in the CAIL2019-SCM dataset, a corpus of thousands of decisions published by the Supreme People's Court of China, in which the task is to determine, for a triple

of cases (A, B, C), whether A is more similar to B or to C [71]. Because of the problem structure, different NLP techniques have been applied to solving the task, including pre-trained models [71], context-based multi-learning [19], ideas from optimal transport [73], casual inference [59], and even techniques combining regex and neural networks [30]. Further techniques for Similar Case Matching were developed and applied to datasets other than CAIL [32, 38], with the work of [6] especially significant for mixing both NLP and citation-based approaches.

Although related in spirit, Similar Case Matching is not exactly the task we are tackling. Instead of computing similarity indexes between pairs of documents or telling which of two documents is more similar to a third one, we need to find in a huge unlabeled corpus of decisions *all* decisions that are similar in content to a particular document. This alternative problem, arguably harder, is known as Similar Case Retrieval and unsurprisingly, has been tackled through different artificial intelligence models [58, 46, 68, 65], including some that we will be using in this work [47]. Surprisingly, though, none of these works seemed to use these techniques to describe the temporal behavior of specific precedents as we do here, restraining themselves to tools for aiding the work of juridical agents.

The reason why we claim the task of case retrieval to be harder than case matching goes beyond the natural computational problem of the sheer size of inputs in the former. Indeed, in this context, when going from the training set (called Dataset #1 in this article), containing actual citations of BPs, to the test set (Dataset #2), where no citations to BPs are to be found, some overfitting of the trained models is expected. However, this bias cannot be directly measured, as there are no documents of this second set which are labeled. Notice that the difficulty is twofold: not only there will be different proportions of positive and negative classes between the two sets (the positive class representing documents citing a specific BP), but also the very textual elements from the positive and negative classes are expected to be distinct, even when regex is used to remove the citations themselves. In conclusion, the problem of Similar Case Retrieval tackled in this article goes beyond other analyses of classification problems found in the literature, where one can, a posteriori, quantify the validity of models through accuracy, recall, or F_1 score. In order to validate our results, we will propose a thorough legal analysis, based in particular on the reading of documents.

Network analysis

Although there is a multitude of NLP models that solve both the similarity matching and retrieval problems, none of these have yet been applied to the actual *juridical* study of the creation, the evolution, and the authority of precedents. On the other hand, empirical legal literature has often investigated these questions using ideas from graph theory. That is, given the natural referential structure inherited from the use of precedents, a substantial part of the empirical study of courts is dedicated to modeling them through directed graphs, generated by the set of all decisions, and where there is an arrow from a decision x towards a decision y if and only if y cites x as a precedent. Drawing from the general theory of citation networks [31], authors have developed metrics to measure the authority of a specific decision as a precedent in a determined court through the decision's (outward) degree in the induced graph, emulating the hypothesis that the more important a precedent, the more cited it is.

These metrics have then been applied to identify the most important precedents in different settings and the evolution of their authority through time. In a fundamental

empirical work on the United States Supreme Court [25], for example, the authors used the (normalized) degree of the decisions to identify historical features of the most important precedents to the court, concluding that for most cases, when an important decision is taken by the court, the number of documents citing it follows a curve which tends to increase with time, reach a peak, and then decrease exponentially. Interestingly, the authors indicate that although this seems to be a general trend, internal forces within the court, such as the reversion of previous decisions, may significantly impact these curves. The authors of [69] point that even external sources, such as the court's political composition impact the citation curves of precedents in the Supreme Court and, in a study considering all instances in the American justice, [55] identifies the probability distributions generating the precedent graphs. Similar conclusions regarding the change of authority over time are found about the International Criminal Court [60], the Canadian legal system [40], and the Court of Justice of the European Union [61]. In this last study, the authors had access to a list of the most important precedents identified by the court's legal experts and attested that, in practice, there exist important cases, called symbolic, which will not have high authority scores, even considering their doctrinal impact.

In an attempt to replicate these studies to the Dutch legal system, [70] discusses whether the use of citation networks is appropriate to a civil law setting, in which precedents play a minor role compared to international courts or common law systems. Interestingly, the authors conclude in favor of the usefulness of network analysis also in this alternative setting, although results should be differently interpreted.

Network analysis of Brazilian system

Another application of network analysis to a civil law system can be found in [14], who applied graph techniques to study precedents of the Brazilian Federal Supreme Court. The authors argue that, given the recent reformations of the Brazilian system aimed toward a higher observance of precedents, graph-based methods become even more useful to this setting. Moreover, historical analysis of authority scores over time indicates citation curves of precedent similar to what was found by [25], although here, the topic of the decision — i.e., constitutional, criminal, etc. — has a strong influence on the shape of the curve. Nonetheless, the authors point out that even in the more "mixed" Brazilian system, some complications arise from the non-exclusivity of the precedent paradigm: the lack of a common law culture makes Justices' citations to other documents missing standardization, which not only complicates the whole construction of the networks but also "contaminate" the documents with only marginal references to precedent. By marginal preferences to precedent, we mean citations to documents that do not impact the merit of the decision, only referring to some procedural aspect of previous decisions.

Moreover, the non-existence of a standardized citation system for the STF allows for what the authors call indirect citations, that is, references to documents that cannot be captured through parsing techniques based on regex. In practice, this means that the interpretation of the authority of a precedent as the number of citations to it at a certain interval of time is insufficient to capture their impact on the legal system.

Finally, the situation is aggravated by the defensive posture assumed by the court to cope with its excessively high demands, with many of the most cited precedents being only procedural decisions of the STF denying to arbitrate on the merit of cases and directing them to other courts. This distinction between decisions in which the

precedent use has a significant juridical value in the judge's juridical reasoning, which we shall call *citations by merit*, versus those in which the citation is only a procedural justification for dictating the next formal steps of the case, which we call *procedural citations*, impose a noteworthy constraint in the interpretation of the number of citations of a document as metric for its authority, for a particularly juridically insignificant decision may be frequently cited due to some incidental procedural extract. These constraints are particularly challenging to deal with in the study of the STF: although the vast majority of court decisions cite precedent, most of them lie in the procedural category [34], limiting the ordinary interpretation of authority scores.

The infant precedent citation culture in Brazil and the predominance of procedural-only citation discourage the use of graph techniques in the study of BPs. Additionally, the very use of these precedents as a tool to *decrease* the frequency of repetitive cases arriving at the STF indicates that degree-based scores are insufficient to describe the effectiveness of the BPs as defense devices, suggesting the use of embedding-based methods.

3 Datasets and methods

3.1 Datasets

Documents

Cases decided at the STF are consolidated in a document, containing the decision's text, as well as the publication date, the Justice that conducted the process (Justice rapporteur), and the document type (e.g., Complaint or Extraordinary Appeal). The research project *Supremo em Números* gathered, in text format, more than 2,500,000 of these documents [23]. In this collection, we shall consider the subset of documents that cite a binding precedent, among the 58 precedents edited by STF until 2018. This amounts to 29,743 documents, which we refer to as Dataset #1 (see Table 1). In addition, we shall consider the so-called Dataset #2, consisting of those documents, in the whole collection, labeled as "administrative law", "criminal law" or "criminal procedure law", gathering 615,262 documents. These are the categories of the five particular BPs we will study (already presented in Section 2.1). We draw the reader's attention to the fact that Dataset #1 is not a subset of Dataset #2, though the latter contains approximately 36% of the former.

Dataset #1	Number of docs.	Dataset #2	Number of docs.
BP 11	830	Administrative law	463,497
BP 14	1601	Criminal law	46,628
BP 17	856	Criminal procedure law	105,137
BP 26	954	Total	615,262
BP 37	3984		
Other BPs	21,518		
Total	29,743		

Table 1: Composition of the datasets involved in our study.

Labels

In the context of Similar Case Retrieval, Dataset #1 is particularly convenient. Indeed, since the 58 binding precedents cover different themes, one decides whether a document deals with a specific theme by reading the corresponding BP it cites. In opposition, no information is available in Dataset #2 regarding the juridical content of the cases. For this reason, we will employ the former to train our models. In the rest of this section, only Dataset #1 will be considered, while the application of our models to Dataset #2 is the content of Section 4.

3.2 Models for Similar Case Retrieval

In this section, we delve into the methodologies employed for analyzing our dataset, focusing on four distinct approaches: TF-IDF, LSTM, BERT, and regex. Together, they faithfully represent some of the most common NLP methods.

TF-IDF

Term Frequency-Inverse Document Frequency (TF-IDF) is a vector embedding technique divided into two parts. In the "term-frequency" part, each document d is embedded in the space $\mathbb{R}^n$, where n is the number of different words on the corpora, and where the i-th component of the vector is the frequency that the i-th word of the corpora appears in the document. We denote it by $\text{TF}(i, d)$. The "inverse document frequency" part is another $\mathbb{R}^n$ vector, with the i-th coordinate given by the logarithm

$$\text{IDF}(i, d) = \log \frac{\# \text{ documents in the corpora}}{\# \text{ documents in the corpora that } i \text{ appears}}$$

Finally, the TF-IDF vector has the i-th coordinate given by the product

$$\text{TF-IDF}(i, d) = \text{TF}(i, d) \times \text{IDF}(i, d).$$

By itself, TF-IDF is only an embedding, so it must be coupled with some classification algorithm for us to perform Similar Case Retrieval. In this article, we will use three kinds of classification models for the embedded documents: linear Support Vector Machine (SVM), logistic regression, and random forests. We will not explicitly describe these classifiers here, but refer the interested reader to [27].

LSTM

A recurrent network architecture was also implemented for this task, given the well-known capacity of these models to work with textual data, especially for longer texts [27, 50]. The implemented neural networks were composed of three layers: one embedding layer, followed by two layers consisting of 128 LSTM (Long Short-Term Memory) neurons each, all using `relu` activation functions. For the embedding layer, `tensorflow` was used, where the vocabulary (i.e., set of unique words in the corpora) was restricted to the 10,000 most common terms, using 1,000 out-of-bucket embedding for the rest of them. Moreover, because the architecture expects a constant input size of documents, only the first 512 tokens of each document were considered, using additional masking for shorter documents. Although naive, this approach is common when dealing with long text analysis in neural networks. Finally, the output consisted of a single neuron followed by a sigmoid non-linearity. We stress that no pre-training of the network's weights was used.

BERT

Bidirectional Encoder Representations from Transformer (BERT) [21] is a common Transformer-based model [66] used for several NLP tasks that can benefit from a deep embedding of a word or a sequence, e.g., a text classification. In this model, each token (word or part of a word) is vectorized and their vectors are sequentially updated to incorporate their neighbors' information. BERT models are released with their weights already optimized in a large amount of text in a self-supervised manner. The output of the model is a vector for each token and one for the entire sentence (corresponding to the first `CLS` token). These vectors can be used as features for any classification model, e.g., a logistic regression. One can also extend this neural network by adding a downstream neural network layer, like a linear classification layer, and then both the new layer and the rest of the model are trained together in the downstream task. Just like the LSTM implementations, BERT models have a text input size limited to 512 tokens. Because of that, texts were truncated to their first 510 tokens (to include other BERT special tokens).

In our work, we use a Portuguese BERT (BERTimbau BASE[4]) model fine-tuned in approximately 16 GB of Brazilian court decision raw data for eight epochs. This data gather decisions from the five most important Brazilian courts: *Supremo Tribunal Federal*, *Superior Tribunal de Justiça*, *Tribunal de Justiça de São Paulo*, *Tribunal de Justiça do Rio de Janeiro*, and *Tribunal Superior Eleitoral*.

regex

Regular expression (regex) stands out as a commonly used tool among professionals for retrieving similar cases. In particular, it constitutes the search engine of the official STF website. We designed a regex search for each of the five BPs by selecting, with the help of legal expertise, the most important words in their statements (please refer to Sections 4.2 to 4.6 for the statement and juridical context of each BPs). We draw the reader's attention to the fact that, as our models are intended to detect, in a subsequent section, the documents *before* creation of the BPs, we do not use the words *Súmula Vinculante* or any specific mention of a law. More precisely, we consider:

BP 11: the words *algemas* or *algemado* (handcuffs, handcuffed),

BP 14: the expression *acesso aos elementos/autos/documentos* (access to elements, records or documents),

BP 17: the word *precatório* and the expression *juros de mora* (court orders, late payment interest),

BP 26: the expressions *exame criminológico* and *progressão de regime* (criminological examination, regime progression),

BP 37: the expressions *isonomia*, *vencimentos* and *servidores públicos* (isonomy, salaries, public servants).

---

[4]Available at `https://huggingface.co/neuralmind/bert-base-portuguese-cased`.

3.3 Validation

Our first analysis focuses on Dataset #1, a collection of documents annotated by the binding precedent they cite. We are, therefore, dealing with a binary classification task. As we will see in this section, and as it is expected from the NLP literature, all the models perform well on this task. However, we remind the reader that the main objective of this article lies in the analysis of Dataset #2 (in Section 4), which is unannotated, and where the "scores" will be quite different.

Preparation of the data

Some preprocessing was applied to the documents before training. For both the TF-IDF and LSTM models, this encompassed the removal of Portuguese-specific accents, and of usual stop words (e.g., *a*, *o*, *um*, *mas*). Because BERT does not expect this sort of preprocessing, raw texts were given.

Additionally, we want to avoid the models from learning textual cues that mark direct citations to BPs, such as the very term *Súmula Vinculante* or similar. Therefore, regular expressions were used to substitute these terms by empty strings. Additionally, as our interest consists in detecting possible uses of BPs before their publications, dates were also removed. Once more, BERT models receive a slightly different masking process than the others: each BP citation is identified using regular expressions and is replaced by `[MASK]` tokens, keeping the same length as the original text.

Our models are trained on Dataset #1 (described in Section 3.1). We randomly split the data set into training and test data, while preserving the distribution of the BPs among the documents. For such, 10% was dedicated for testing, and from the training data, a further 10% was used for validation. Although the pre-processing of input texts was not the same for each class of models, we ensured that all shared the same documents for training, testing, and validation. Lastly, given a BP, we attributed to the documents the labels "1" or "0", depending on whether the document cites or not the BP, and trained the models for this binary classification task.

Scores

To assess the performance of our binary classifiers, we use a handful of score metrics. By denoting P (number of class 1 examples), N (class 0), TP (class 1 predicted class 1), FP (class 0 predicted class 1), TN (class 0 predicted class 0), and FN (class 1 predicted class 0), we consider the following metrics:

$$\text{Accuracy} = (\text{TP} + \text{TN})/(\text{P} + \text{N}),$$
$$\text{Precision} = \text{TP}/(\text{TP} + \text{FP}),$$
$$\text{Recall} = \text{TP}/\text{P}$$
$$F_1 = 2\,\text{Precision} \cdot \text{Recall}/(\text{Precision} + \text{Recall}).$$

Intuitively, accuracy measures the ratio of correct predictions, while precision measures the accuracy of pointing to class 1, and recall measures how much of class 1 is being correctly predicted. The F_1 score is the harmonic mean of recall and precision.

Most models return a continuous output, e.g., a probability. In this case, defining when the model is predicting classes 0 or 1 depends on the classification threshold. The *Area Under the Precision Recall Curve* (AUPRC) swaps a classification threshold

extreme values, measuring precision and recall for each threshold, and creates a precision vs. recall curve. The area under the precision recall curve establishes a common performance metric. The results are gathered in Table 2.

BP	Model	F_1	Precision	Recall	AUPRC
11	TF-IDF + SVM	**93.9**	98.6	**89.6**	**88.6**
	TF-IDF + logistic	92.3	**100**	85.7	86.1
	TF-IDF + forest	85.1	**100**	74.0	74.7
	LSTM	87.8	98.4	79.2	78.5
	BERT	90.5	94.4	87.0	82.5
	regex	91.7	98.5	85.7	84.8
14	TF-IDF + SVM	96.3	97.0	**95.6**	92.9
	TF-IDF + logistic	**97.0**	98.5	**95.6**	**94.3**
	TF-IDF + forest	94.1	**100**	88.9	89.4
	LSTM	95.5	96.9	94.1	91.5
	BERT	92.1	93.2	91.1	85.3
	regex	73.1	97.5	58.5	59.0
17	TF-IDF + SVM	96.7	98.7	94.9	93.7
	TF-IDF + logistic	**98.0**	**100**	**96.2**	**96.3**
	TF-IDF + forest	90.1	**100**	82.1	82.5
	LSTM	93.9	**100**	88.5	88.8
	BERT	90.5	95.7	85.9	82.6
	regex	88.1	86.4	89.7	77.8
26	TF-IDF + SVM	**98.6**	**100**	**97.3**	**97.4**
	TF-IDF + logistic	**98.6**	**100**	**97.3**	**97.4**
	TF-IDF + forest	95.8	**100**	91.9	92.1
	LSTM	93.7	97.1	90.5	88.2
	BERT	94.4	97.1	91.9	89.5
	regex	86.2	100	75.7	76.3
37	TF-IDF + SVM	**95.2**	97.4	**93.1**	**91.4**
	TF-IDF + logistic	94.5	97.0	92.1	90.3
	TF-IDF + forest	85.4	**100**	74.5	77.4
	LSTM	91.7	95.3	88.4	85.5
	BERT	90.6	93.6	87.7	83.5
	regex	83.4	93.7	75.2	73.2

Table 2: For each BP and each model, we indicate the scores F_1, precision, recall, and AUPRC (percentage). The best score is shown in bold.

Discussion

One sees from Table 2 that the TF-IDF models, equipped with the classifiers SVM and logistic regression, consistently achieve the best performance across all BPs, measured by the F_1 score. On the other hand, the deep learning models LSTM and BERT show slightly lower scores. This phenomenon has already been reported in [47], using, as we do Brazilian legal documents (more precisely, data from *Supremo em Números*). In addition, it has been observed that the TF-IDF-based models outperform certain more modern embeddings (such as Doc2vec, Universal Sentence Encoder, and Longformer). In this context, TF-IDF's superior performance is attributed to its ability to leverage

specific words that allow it to identify the BP. This is supported by the regex search: although being a rather simple way of answering the problem of Similar Case Retrieval, its performance isn't too bad, even outperforming deep learning models for BP 11.

The case of TF-IDF coupled with random forest is interesting: it achieves perfect precision across all BPs. That is to say, it is highly reliable when it comes to positive predictions. On the other hand, its recall is the lowest in almost every case. Its high precision will be of great importance in Section 4, when we apply the models to Dataset #2. Indeed, we will observe that the models suffer from a serious generalization problem, more precisely, they detect too many documents, which are not related to the topic of the BP. The model TF-IDF+random forest, on the other hand, will mainly detect documents that we have checked as relevant. This comes at the cost of potentially missing documents; we will therefore be using the other TF-IDF models in a complementary way.

It should be noted that the performances of LSTM and BERT may suffer from the fact that their inputs are limited to 512 tokens. One way around this problem, although not considered in our work, is to cut the input documents into several parts, apply the models to each of these parts, and agglomerate the results — a method known as pooling [76]. In addition, we have used the BERT model as is, trained previously on Brazilian legal documents, but not fine-tuned to our specific dataset. The implementation of these two corrections could significantly improve the performance of these models, and we intend to study them in a subsequent work.

3.4   Explainability

To understand further what the models have learned, we move on to study their most important features. For the TF-IDF models, common measures of importance can be employed and directly computed from the models. For BERT, however, there is no direct way of identifying the most important features, thus we will use the explainability algorithm LIME [48]. Note that, although not presented here, a similar analysis could be made with LSTM.

Explainability with TF-IDF

For the classifiers SVM, logistic regression and random forest, based on the TF-IDF vectorization, common measures of importance of features are respectively the weights of the linear kernel, the coefficients in the decision function, and the standard deviation of impurity decrease in the trees. We compute these quantities via the native functions of `scipy`. We inspect, for each model, the top features, and gather those common to all models. Table 3 presents the importance of the words selected this way.

BP	Classifier	Features				
		algemas	*uso*	*algemado*	*audiência*	*nulidade*
		(handcuffs)	(use)	(handcuffed)	(court hearing)	(nullity)
11	SVM	1.0	0.43	0.39	0.33	0.32
	logistic	1.0	0.44	0.39	0.37	0.33
	forest	1.0	0.49	0.3	0.42	0.25

		acesso	*investigação*	*criminal*	*inquérito*	*documentados*
		(access)	(investigation)	(criminal)	(inquiry)	(documented)
14	SVM	1.0	0.67	0.61	0.56	0.45
	logistic	1.0	0.64	0.55	0.56	0.42
	forest	1.0	0.45	0.21	0.58	0.64
		Guarujá	*precatório*	*juros*	*mora*	*cálculos*
		(Guarujá)	(court order)	(interest)	(late)	(calculations)
17	SVM	1.0	0.66	0.62	0.5	0.41
	logistic	0.81	1.0	0.96	0.75	0.57
	forest	0.01	1.0	0.72	0.52	0.14
		criminológico	*progressão*	*execuções*	*regime*	*realização*
		(criminological)	(progression)	(executions)	(regime)	(realization)
26	SVM	1.0	0.82	0.6	0.58	0.49
	logistic	1.0	0.86	0.62	0.58	0.45
	forest	1.0	0.59	0.38	0.25	0.4
		isonomia	*cabe*	*aumentar*	*vencimentos*	*públicos*
		(isonomy)	(apply)	(increase)	(salaries)	(public)
37	SVM	1.0	0.67	0.61	0.44	0.42
	logistic	1.0	0.59	0.56	0.55	0.42
	forest	1.0	0.27	0.78	0.46	0.29

Table 3: For each BP and each TF-IDF model, we computed the importances of the features, and selected the five highest ones. The values are normalized so that the largest importance is equal to one.

The results of Table 3 are not surprising: the most important words are already contained in the wording of the BPs, or are related grammatically, except two (*audieêcia* and *Guarujá*, studied below). We refer the reader to Sections 4.2 to 4.6 for the wordings, as well as a more detailed analysis of the results. In addition, it is remarkable that all three classifiers display similar importance for each of the words. This shows that they have all faithfully learned the statements of the BPs.

The first feature to stand out is *audiência* (audience/trial), not directly connected to BP 11. However, as explained in Section 4.2, this precedent was born out of the controversy surrounding the use of handcuffs (*algemas*) in trials. It is therefore consistent that this word has been recognized as important by the models.

Similarly, now in the context of BP 17, the word *Guarujá* (municipality in the state of São Paulo) is associated with high importance for the models SVM and logistic (1.0 and 0.81) but not for random forest (0.01). As it turns out, this word is only marginally relevant when it comes to identifying the binding precedent: only 9 documents of the training set contain *Guarujá*, among which 5 cite BP 17. The importance given to this feature is, therefore, a mere consequence of training the models to maximize the F_1 score, and carries little legal information.

Explainability with LIME

In the context of explainability, instead of directly assessing the weights of an interpretable model (e.g., a logistic regression), we can treat the model as a black box and verify its behavior when we disturb its input. For instance, one could be interested in

how much the output probability changes if one removes certain words from the text. If it changes significantly, these words are deemed important for the model decision.

This is precisely the underlying mechanism of LIME (*Local Interpretable Model-agnostic Explanations*) [48], a popular model-agnostic machine learning explainability method. More precisely, to explain the model decision for a specific text sample, LIME perturbs the input of the classifier by randomly removing words from the text and measures the model output probability. Then, after thousands of disturbances, LIME fits a linear regression to predict the model probability from the word presence. From the linear regression coefficients, we can assign an importance score for each word in the input text regarding their contribution to the model decision.

In practice, we use LIME in the following way: for each document in the collection, we calculate with LIME the N most important words (i.e., words which contribute most to the positive class), where N is a hyperparameter, chosen as $N = 5$. We then search, among all the documents belonging to a BP, for the most frequent LIME features, and the number of times they appear. The N most important features are represented in Table 4.

BP	Features				
11	*algemas* (handcuffs) 57.0	*algemado* (handcuffed) 15.8	*audiência* (court hearing) 5.7	*algemados* (handcuffed) 4.7	*prisão* (prison) 4.3
14	*acesso* (access) 29.7	*inquérito* (inquiry) 10.0	*criminal* (criminal) 8.4	*defesa* (defense) 7.1	*investigação* (investigation) 4.4
17	*precatório* (court order) 31.1	*juros* (interests) 24.1	*100* (100) 12.7	*mora* (late) 12.0	*precatórios* (court orders) 10.7
26	*criminológico* (criminological) 33.2	*progressão* (progression) 16.4	*regime* (regime) 6.5	*RAJ* (Judicial Region) 5.1	*execuções* (executions) 4.4
37	*isonomia* (isonomy) 18.3	*reajuste* (readjustment) 9.2	*vencimentos* (salaries) 7.4	*2015* (2015) 4.8	*equiparação* (equivalent) 4.5

Table 4: For each document, LIME detects the most important features. For each BP, we then select the five most frequent features, and indicate the percentage of documents in which LIME detected them.

Some interesting observations can be made from Table 4. First, and as expected, we can see that many words are shared with the TF-IDF models (presented in Table 3). However, in the latter models, we noticed that the important features were almost exclusively words already present in the BP statement. In the case of BERT, LIME brings to light features which, although highly relevant to the identification of the legal theme covered by the BP, do not appear in its wording. This is the case of *inquérito* (inquiry/investigation) for BP 14, or *reajuste* (readjustment) for BP 37; we refer the reader to Sections 4.3 and 4.6 respectively for a legal analysis.

On the other hand, one sees that these important words are only detected by LIME in a limited number of documents. This is at most 57%, for the word *algema* (handcuffs) in BP 11. For BP 37, the most important word is *isonomia* (isonomy), which is detected

in only 18.3% of documents. One might have expected these words to have been detected more frequently. This could explain the relatively low performance of the BERT models shown in Table 2: the models are not sufficiently fit for the classification task.

4   Results and legal analysis

This section is concerned with the application of our models, trained on the first dataset (documents citing a BP), to the analysis of the second dataset (whole collection of documents emitted by the High Court). From a legal point of view, we are interested in uncovering the juridical mechanisms behind the observed increase in cases citing a binding precedent, visualized in Fig. 1. Besides, from a Machine Learning perspective, we aim to verify the ability of the models to generalize their results, when applied to a larger dataset. We shall start with a general description of our methodology for empirical evaluation of binding precedents (Section 4.1), then give a detailed analysis of each BP (Sections 4.2 to 4.6), and finally draw juridical conclusions (Section 4.7).

4.1   Description of our methodology

Time series of similar cases

To evaluate the effect of a binding precedent on jurisprudence, we argue that one should not perform a mere reading of the cases citing that precedent, but also compare them to similar cases published before the precedent. This information can be visualized through the curve which associates, to each timestamp, the number of similar documents. In our context, we expect such a curve to behave as follows: the existence of a peak of documents just before the publication of the BP, followed by a steady decrease, until reaching a stable state, hopefully consisting of few documents. This behavior would indicate that the BP served its main purpose, which is to "settle" an increasingly contentious legal issue, clarifying for interested parties and the lower courts what the final position of STF is, thus reducing the likelihood of future conflicts involving the same issue.

If this behavior is not observed, as it is the case for the five precedents studied in this article, then one can add metadata information to the curve, to reveal the reasons for this deviant phenomenon. Namely, we will consider the type of process (e.g., habeas corpus or appeal), the state of provenance, the decision of the court (accepted, rejected or deferred), and the presence of certain words (specific to the legal subject considered).

We point out that to obtain a thorough analysis of the impact of a BP, it would have been appropriate to detect similar documents issued not only by the STF but by lower courts as well. This analysis, however, is beyond the scope of this article, given the current unavailability of such a dataset.

Generalization of the models

As discussed in Section 3.1, this article deals with Datasets #1 and #2, the former consisting of documents citing a binding precedent (29,743 documents), and the latter of all the documents in the categories "administrative law", "criminal law" and "criminal procedure law" (615,262 documents). In Section 3.3, we trained the models on Dataset #1 and presented the results (see Table 2). We will, in this section, apply these models on Dataset #2. In doing so, we face a significant problem: the models have been tuned to

the first dataset, and are likely to give poor results on the second. As a matter of fact, we observed that the models tend to predict way more documents than intended, that is to say, they are underfitted.

To remedy this problem, we will not consider the binary output of the models, but rather their "probabilities". As a native property of the models we considered, each document is associated with a likeliness of belonging to the positive class. In the case of SVM, this number is a distance; for the other models, it is a value between 0 and 1. Internally, a document is assigned the positive class if this value exceeds a certain threshold. To adapt the models to the new dataset, one should increase this threshold, discarding the false positives.

However, as already pointed out in Section 3.1, Dataset #2 does not come with indications of whether a document is related to a certain BP. Thus, this dataset cannot be used directly to tune the thresholds. Instead, we will take the following observation as a reference: if a document, issued *after* the publication of the BP, is related to the subject of this BP, then it probably cites the BP. In other words, we can take the annotations of Dataset #1 as ground truth, valid only after the publication of the BP. A well-tuned model is expected to follow faithfully the curve of citation of the BP.

In practice, we chose the new thresholds as the largest such a value that the recall of the predictions, after publication of the BP, is greater or equal to a fixed value $p \in [0, 1]$. That is, we impose that at least a proportion p of the groundtruth documents are detected. Note that the larger p is, the more documents are detected. Specifically, we chose the values $p = 0.9$ for BPs 11, 14 and 26, and $p = 0.8$ for BPs 17 and 37. Although arbitrary, these values have been chosen in accordance with the models' performance in Table 2, and because they yield satisfactory results. We indicate in Table 5, for each model and each BP, both the initial and the tuned value of the thresholds. For all cases, the tuned one is higher, indicating a model with fewer positive predictions. This will be also illustrated in the Figs. 3, 6, 10, 15 and 19, which represent the time series of estimated documents by the models, both for the initial thresholds (dashed curves) and the tuned thresholds (solid curves).

Model	Threshold	BP 11	BP 14	BP 17	BP 26	BP 37
TF-IDF+	Dataset #1	-0.7390	-0.6639	-0.9233	-0.6967	-0.7248
SVM	Dataset #2	0.8867	0.9661	0.9101	0.9601	0.9428
TF-IDF+	Dataset #1	0.5	0.5	0.5	0.5	0.5
logistic	Dataset #2	0.9967	0.9990	0.9978	0.9983	0.9649
TF-IDF+	Dataset #1	0.5	0.5	0.5	0.5	0.5
forest	Dataset #2	0.59	0.662	0.611	0.642	0.538
LSTM	Dataset #1	0.9999991	0.9970	0.99992	0.9997	0.9848
	Dataset #2	0.9999999	0.9999	0.99995	0.9999	0.9999
BERT	Dataset #1	0.8431	0.9676	0.7256	0.1393	0.9543
	Dataset #2	0.9906	0.9928	0.9995	0.9955	0.9998

Table 5: For each model and each BP, the threshold over which a document is assigned the positive class is indicated. The first line refers to the initial threshold (computed from Dataset #1) and the second line to the tuned one (adapted to Dataset #2).

Besides, to evaluate the quality of the predictions, another tool will be employed: the correlations between the words. Given a specific word (for instance, a relevant word associated with the BP or from its wording), the presence or absence of it in a document

can be seen as a binary variable. From this point of view, the correlation between two words indicates how often they are found combined. Since there is little difference in the way two documents on a similar subject are written — it is often observed that some decisions are merely copies of another —, the collection of documents predicted by a model is expected to show similar correlations compared to the collection of documents citing the BP. Comparing these measurements will allow us to check the models' accuracy and to find out why some do not perform well.

Common hypotheses

A variety of legal, administrative, and political factors may explain the increase in the number of cases concerning a given subject. We list below some of these hypotheses, which will be embodied by the binding precedents in the next sections. We will conclude on the adequacy of these hypotheses in Section 4.7.

New avenue to STF or procedural issues	The adoption of a BP creates a new gateway for parties to bring their cases to the STF. Since the existing legislation restricts access to the STF, these cases might otherwise have received their final decision from a lower court. Conversely, losing litigants have an incentive to take advantage of a new opportunity to reach STF in an attempt to change the lower court's decision. This also creates procedural disputes regarding the proper use of specific classes of appeals and what cases fall under the new BP.
Resistance by a group of litigants or regional specificity	A group of recurrent litigants or judges deviates from the pacified interpretation of the BP. Because of the current organization of the Brazilian Judiciary, it is not uncommon for regions and states to develop diverging positions on specific legal issues. When a new BP conflicts with such a position, many cases that reach STF might stem from this disagreement, following an attempt from specific groups to influence the interpretation of the BP.
Vague wording	An element of the precedent is open to interpretation, leading to cases that seek to clarify its concrete meaning. Vagueness might create, for instance, confusion related to which cases fall under the BP and which do not. It may also create disputes over how the content of the BP should be interpreted and applied in a specific case. In all these instances, many new cases might reach the STF as a byproduct of legal uncertainty.
New theme	The BP introduced a new theme, resulting in a variety of cases discussing this previously unmentioned topic in STF. It should be noted that, in every legal system, judicial decisions entail a creative element, particularly heightened when a court possesses the authority to establish general norms through precedents. Thus, a BP might also be wielded as an instrument of judicial activism, whenever STF encounters a situation where current social practices are deemed contradictory to the Constitution.
External conjecture not directly related to the precedent	A new development unrelated to the issuing of the BP has brought a specific legal issue into the spotlight, causing, as a side-effect, a significant number of decisions to refer to the precedent. In this context, the growth in the number of citations should not be seen as a direct effect of the precedent, but rather as a natural trend in the law, where subjects of interest are dynamic and dependent on political, economic, and social conjectures.

4.2 BP 11

Juridical context

During the trial of HC (*Habeas Corpus*) 91.952[5] (08/07/2008), the defense requested the removal of handcuffs from the defendant, due to concerns about the negative perception that the sight of handcuffs could convey to the jury. This case triggered a highly publicized debate, which then shifted to discussing the use of handcuffs for public exposure, and more generally the sensationalization of criminal prosecution in the media. Based on the presumption of innocence, individual freedom, and the dignity of the human person, the High Court voted a few days later on the following:

> **Binding Precedent 11.** "The use of handcuffs is only permitted in cases of resistance and a well-founded fear of escape or danger to the physical integrity of the prisoner or others, justified in writing, under penalty of disciplinary, civil, and criminal liability of the agent or authority and nullity of the arrest or the procedural act to which it refers, without prejudice to the civil liability of the State." (STF, 08/2008)

> **Súmula Vinculante 11.** "Só é lícito o uso de algemas em casos de resistência e de fundado receio de fuga ou de perigo à integridade física própria ou alheia, por parte do preso ou de terceiros, justificada a excepcionalidade por escrito, sob pena de responsabilidade disciplinar, civil e penal do agente ou da autoridade e de nulidade da prisão ou do ato processual a que se refere, sem prejuízo da responsabilidade civil do Estado." (STF, 08/2008)

This precedent asserts that the use of handcuffs is allowed only when explicitly justified. In particular, their use is prohibited in the context of trials or media exposure, which was deemed as humiliating by part of the public discourse. It has been considered a notable accomplishment in advancing the principles of the Democratic Rule of Law over those of a Police State [16, 53].

It is worth mentioning certain controversies surrounding the precedent. First, the nature of the nullity (absolute or relative) is not explicitly specified. Currently, the Supreme Court understands that the nature of the nullity is relative, as it depends on a demonstration of concrete harm to the defendant. However, [53] note that lower courts across the country have been deciding similar cases under the assumption of absolute nullity. Besides, as pointed out by [54] and [17], a question remains regarding the exact modalities of *justification* of the fear of escape or danger evoked in the text.

As already visualized in Fig. 1, the number of documents citing BP 11 has been rising steadily since 2008, with a notable jump in 2016. We aim, through our analysis, to uncover some of the reasons for this increase.

Mathematical results

We show in Fig. 3 the predictions of the five models on our whole collection of documents (Dataset #2), as well as a regex search for documents containing the words *algemas* or *algemado* (in English, "handcuffs" and "handcuffed"). As explained in Section 4.1, the thresholds of the models, trained on Dataset #1, are not adapted for Dataset #2: way too many documents are predicted. Therefore, we choose new thresholds with the following rule: the highest value such that at least 90% of the documents citing BP 11

---

[5]HC 91.952 `https://redir.stf.jus.br/paginadorpub/paginador.jsp?docTP=AC&docID=570157`

are predicted. In the figure, the initial and the tuned prediction curves are visualized as dashed or solid curves, respectively.
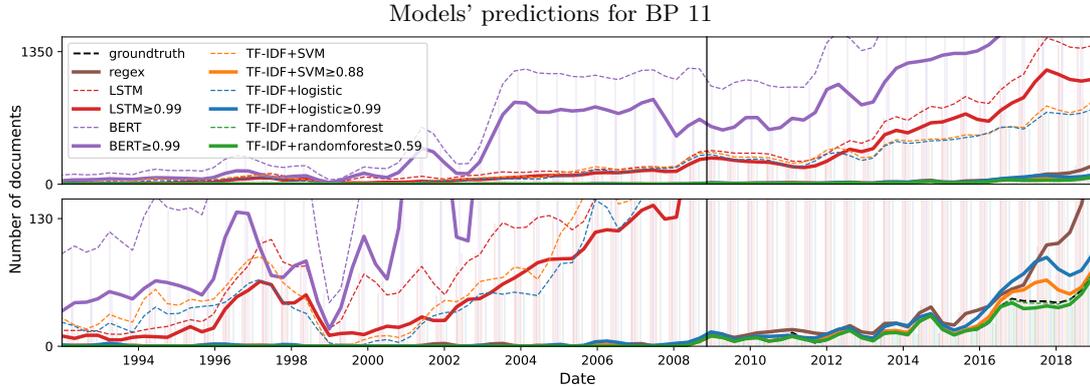

Figure 3: Number of documents predicted by each model for BP 11 in Dataset #2, represented as a histogram (window length of six months) and interpolated via quadratic spline. We give the predictions for the thresholds adapted to Dataset #1 (dashed) and #2 (solid), as described in Table 5. Two views are given, the second one zooming in on the ordinate axis.

As it turns out, there are almost no documents, before the publication of BP 11, containing the words of the regex query. This phenomenon is also illustrated by the models TF-IDF+SVM, TF-IDF+logistic, and TF-IDF+randomforest, which closely follow the groundtruth curve. As a matter of fact, among the 853 documents of our dataset containing the words *algemas* or *algemado*, respectively 422, 433, and 420 have been predicted by the model, accounting for 79.9%, 63% and 99.1% of their predictions. In these documents, these words are often used in conjunction with *uso* and *nulidade* (use and nullity), also present in the official wording of BP 11.

The other models, however, estimate thousands of documents. Only 3.6% of LSTM's predictions contain the words *algemas* or *algemado*, and 1.4% for BERT. A manual inspection shows that these models tend to detect documents merely containing the words *uso* or *nulidade*, independently of the surrounding context. For instance, among the 148,923 documents containing the word *uso*, the models LSTM and BERT predict, respectively, 6903 and 13,866 documents, accounting for 51,6% and 36,1% of predicted documents. This suggests that these models are biased toward detecting this word, regardless of its role in the conjunction *uso de algemas* (use of handcuffs).

To analyze this phenomenon further, let us consider the TF-IDF's most relevant words, obtained in Table 3: *algemado, algemas, audiencia, fuga, justificada, nulidade,* and *uso* (handcuffed, handcuffs, audience, runaway, justified, nullity, use). We give in Fig. 4 the correlations between these seven words, when restricting the dataset to the documents explicitly citing BP 11, as well as the documents estimated by the models. As explained in Section 4.1, a well-fitted model is expected to show similar correlations as the ground truth. The first image shows a high correlation between *uso* and *algemas*, indicating that they mainly belong to the set *uso de algema*. Besides, only the TF-IDF-based models reflect this behavior. This observation accounts for the fact that LSTM and BERT, by not taking into account the correlation between these words, show poor results when used on the larger dataset.
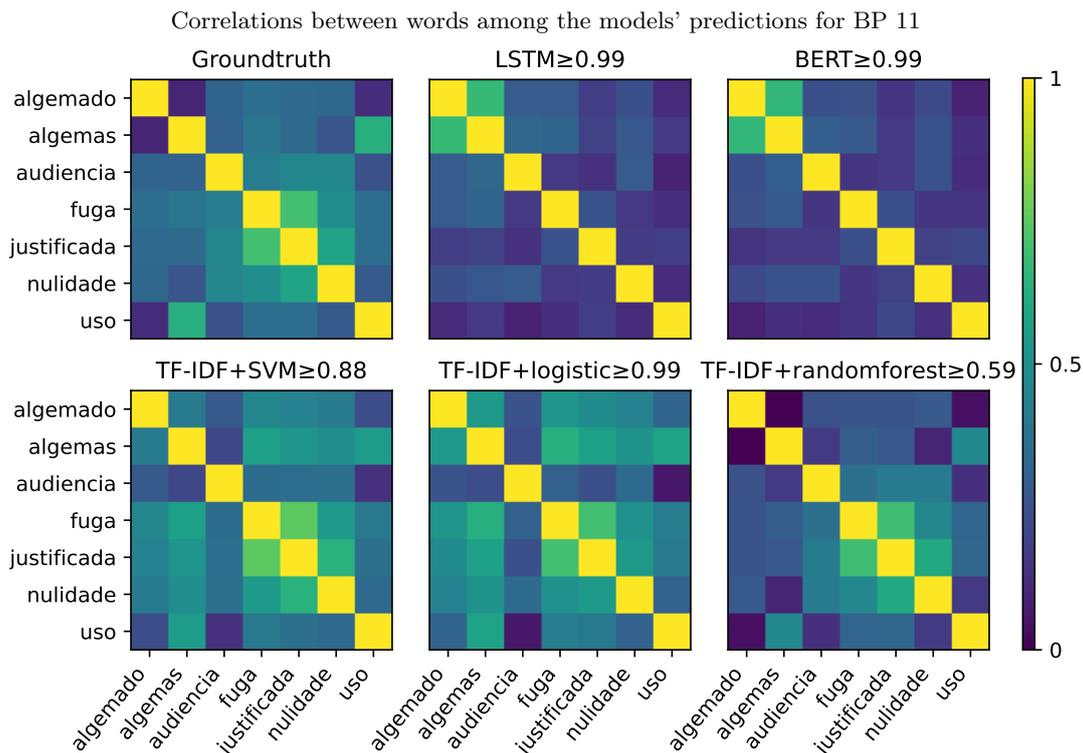
Figure 4: Correlations between the selected relevant words, when restricted to documents citing BP 11 ("Groundtruth") or documents predicted by one of our models. Please note that, as discussed in Section 4.1, correlations are calculated by viewing the words as binary random variables, hence correlation values lie between 0 and 1.

Juridical discussion

As a consequence of the analysis above, we take TF-IDF models to be the most reliable for the retrieval of documents similar to those citing BP 11. We shall see in the next sections that this conclusion holds for all the precedents considered. The case of BP 11, however, is unfortunate: no document is predicted by the models before the creation of the precedent, hence no interesting behavior can be observed in Fig. 3. This restricts the application of our methodology described in Section 4.1. Nevertheless, two juridical insights can be drawn. First, the regex search shows that the term "handcuffs" has come to be used in STF's cases recently, thanks to the edition of BP 11. It should be noted that a regex search for *uso de força* (use of force), a term commonly used as a synonym of handcuffs, only finds 153 documents, 74.8% of which also cite the precedent. On top of that, the fact that the TF-IDF models do not detect any similar documents before 2008 suggests that BP 11 not only changed the way cases are formulated, but really introduced this matter in the high court. This observation would fall under the *New theme* hypothesis, evoked in Section 4.1 as a way of explaining the inefficiency of the precedent in reducing the number of cases.

We see in this precedent that a BP does not solely serve to pacify jurisprudence, but it might also be used as a tool of judicial activism by STF, to resolve a situation deemed contrary to the Constitution. Among these situations is the widespread use of handcuffs, addressed in the precedent, which, despite being widely accepted by the public, the Justices deemed to be incorrect and to violate the basic rights of the defendants. In

this case, the BP did not originate from a recurrent legal controversy, but from the realization, by the court, that common social practices were in sharp contradiction with the Law.

Following our methodology, we represent in Fig. 5 the type of juridical processes used in conjunction with BP 11. As observed, most of these cases are *reclamações* (complaints/appeals, abbreviated Rcl), a specific process filed to contest a decision or address a grievance, typically related to matters falling within the jurisdiction of the STF. To understand the purpose of these claims, we read some of the 632 documents citing BP 11. It turns out that many decisions of the STF are merely "procedural", in the sense that they discuss the unsuitability of the *reclamação* as the proper avenue for the case. All these cases are deferred, based on internal procedural regulation. This seems to indicate that the *reclamação* is used by lawyers as a substitute for an appeal, corresponding to our hypothesis *New avenue to the court and procedural issues*.

Additional metadata for documents predicted by TF-IDF+randomforest for BP 11
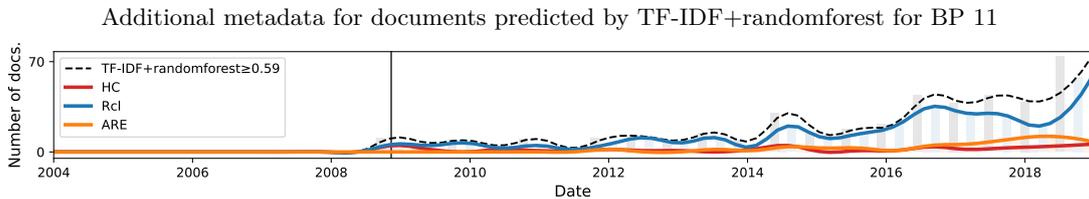
Figure 5: Type of legal process used in the documents predicted by TF-IDF+randomforest for BP 11 in Dataset #2, among ARE (*Recurso Extraordinário com Agravo*), Rcl (*Reclamação*) and HC (*Habeas Corpus*).

Besides, we point out that, among the *reclamações* that actually address the question of the use of handcuffs — i.e., non-procedural cases — the majority found their use adequately justified. However, these decisions often refrain from delving into a detailed discussion of the justification. This issue is reinforced by the legal precedent of *preclusão* (preclusion), which states that if there were no objections at the time of handcuff use, and if an authority figure was present, then the *reclamação* cannot be used. The often sparse written justification of handcuffs use, as well as the use of *preclusão*, call for a clarification of BP 11's statement, as already mentioned in legal literature [54, 17].

4.3 BP 14

Juridical context

Binding Precedent 14 derived directly from the Attorney Statute of 1994[6], which predicts the defense attorney's rights "to examine, in any institution responsible for conducting investigations, even without authorization, records of flagrant crimes and investigations of any kind, completed or ongoing, even if already concluded for the authority, being able to copy documents and take notes, in physical or digital format".

This right, however, is not free from dispute, especially in situations when classified investigations take place. The *habeas corpus* 88.190[7] exemplifies these further complications: in 2005, the Brazilian newspaper *O Globo* brought light to the existence of an investigation aiming at one of the owners of an important exportation company, a fact that was also new to the investigation's target, given its classified status. After

---

[6] Article XIV `https://www.oab.org.br/Content/pdf/LegislacaoOab/Lei-8906-94-site.pdf`
[7] HC 88.190 `https://redir.stf.jus.br/paginadorpub/paginador.jsp?docTP=AC&docID=382091`

consistent refusals of investigation disclosure, lawyers decided to file an *habeas corpus* to the Federal Court of the Second Region (*Tribunal Regional Federal da 2ª Região*, TRF-2) against the Prosecutor's Office, asking for access to the investigation proceedings. After failing to obtain a favorable decision at TRF-2 and the Supreme Court of Justice[8] (*Superior Tribunal de Justiça*), they filed an *habeas corpus* to STF pleading the recognition of unconstitutionality of the denial of access to the ongoing investigation documents, based on the guarantees of the right of sustaining contradictory positions through the whole criminal litigation. After debates, the STF Justices decided to allow the responsible attorneys to access the classified materials regarding their clients.

Cases similar to the *habeas corpus* described in the previous paragraph convinced Justices to approve Binding Precedent 14 in March 2009, which states the following:

> **Binding Precedent 14.** "It is the right of the defender, in the interest of the represented party, to have broad access to the evidence elements that are already documented in an investigative procedure conducted by an authority with jurisdiction in criminal investigation, and that pertain to the exercise of the right to defense." (STF, 09/2009)

> **Súmula Vinculante 14.** "É direito do defensor, no interesse do representado, ter acesso amplo aos elementos de prova que, já documentados em procedimento investigatório realizado por órgão com competência de polícia judiciária, digam respeito ao exercício do direito de defesa." (STF, 09/2009)

The time series at the top of Fig. 1 indicates that, for the first few years after its approval, BP 14 was effective in reducing STF cases regarding its topic, although this is followed by a steep increase in references to the object. This calls for a throughout study of the reasons behind this trend.

Mathematical results

A quick analysis of the predicted documents in the bottom part of Fig. 6 indicates that, among the five models considered, only TF-IDF+random forest achieved satisfactory results, by faithfully following the groundtruth curve. This is also the case for the regex search of the expression *acesso aos elementos/autos/documentos* (access to elements, records, or documents).

To investigate why the other models estimate many documents, we will not consider the correlations, as we did for BP 11 (in Fig. 4), but simply the frequencies of the most relevant words, represented in Fig. 7. Almost all groundtruth documents cite *acesso aos autos* (access to records). Similarly, 96.3% of TF-IDF+randomforest's predictions mention *autos*, of which 81.8% also mention *acesso*. Besides, among TF-IDF+SVM and TF-IDF+logistic's predictions, respectively, 93.5% and 93.6% mention *autos*, while only 44.6% and 41.8% of them conjointly contain *acesso*. This suggests an underfitting of both models.

On the other hand, LSTM and BERT predict a quantity of documents not even mentioning the words *acesso* and *autos*. As it turns out, the decisions detected before the creation of BP 14 seemed of little relevance regarding its topic. Most of them[9]

---

[8]Different from the STF, this institution is the highest court of appeal in cases where no direct constitutional rights seem to be trespassed.

[9]such as MS 26.024 `https://redir.stf.jus.br/paginadorpub/paginador.jsp?docTP=AC&docID=4` `35164` and HC 83.703 `https://portal.stf.jus.br/processos/detalhe.asp?incidente=2180019`
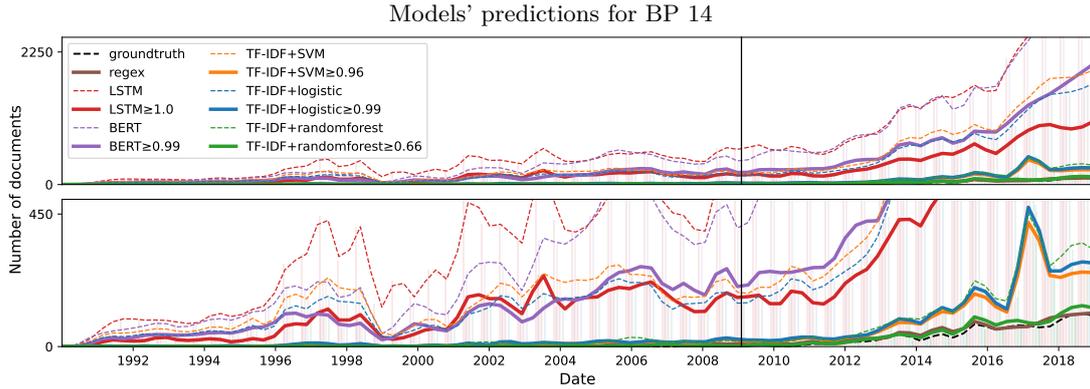
Figure 6: Number of documents predicted by each model for BP 14 in Dataset #2, represented as a histogram (window length of six months) and interpolated via quadratic spline. We give the predictions for the thresholds adapted to Dataset #1 (dashed) and #2 (solid), as described in Table 5. Two views are given, the second one zooming in on the ordinate axis.
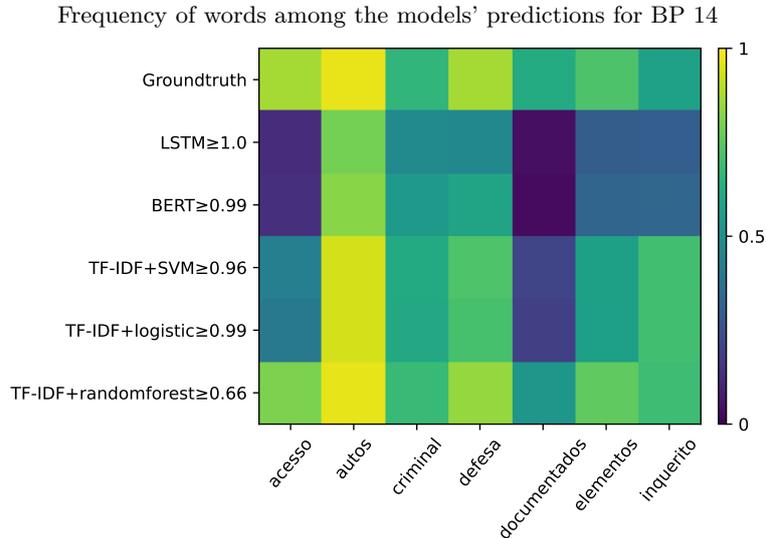

Figure 7: Frequency of the selected relevant words, when restricted to documents citing BP 14 ("Groundtruth"), or documents predicted by one of our models.

are related to "investigations", although not addressing directly the attorneys' right of access to the investigation material. We consider these predictions to be unreliable.

Interestingly, both the predictions of TF-IDF+SVM and TF-IDF+logistic peaked in 2017. At first sight, a manual inspection of a sample of the detected documents, excluding elements of the groundtruth class, raises questions about the qualities of these predictions. Although certain decisions[10] are on topics regarding investigations (specifically, on the initiation of an official investigation and the transcription of some telephonic tapping record), neither of them concerns topics related to *access* of investigation material by the investigated part, which is the subject of the BP. The addition of metadata,

---

[10]Such as Inq 4.260 `https://portal.stf.jus.br/processos/downloadPeca.asp?id=313917544&ext=.pdf` and RHC 133.298 `https://portal.stf.jus.br/processos/downloadPeca.asp?id=315014507&ext=.pdf`

as in Fig. 8, shows that this peak mainly consists of cases from *Distrito Federal* that also mention the word *sigilo* (secrecy/confidentiality). This suggests that these cases are complaints filed by attorneys representing individuals investigated or condemned for white-collar crimes, most relevantly, by attorneys representing politicians or individuals of political interest. It is, therefore, unsurprising that this increase in usage of the BP after 2012 partly matches the interval when *Operação Lava Jato* took place, a joint cooperative investigation aiming at politicians accused of corruption-related crimes. We continue this discussion, from a legal point of view, in the next paragraph.
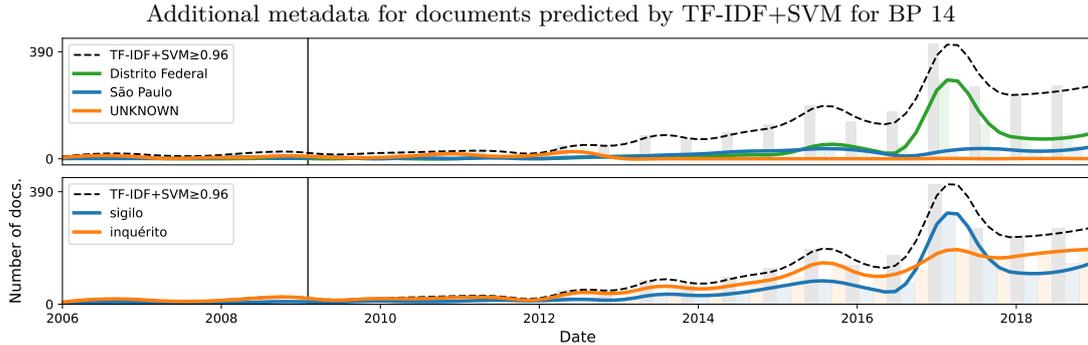
Figure 8: **Top:** State to which the cases predicted by TF-IDF+SVM for BP 14 belong, with label "UNKNOWN" when the metadata is unavailable. **Bottom:** Regex search for the presence of the words *sigilo* (secrecy/confidentiality) and *inquérito* (inquiry/investigation).

Juridical discussion

First, we shall analyze the predictions of models TF-IDF+SVM and TF-IDF+logistic, which, although not following the ground truth curve, evoke a political event worth pointing out. A reading of the documents shows that the peak of Fig. 8 can be dated to April 4th, 2017, the day which Judge Edson Fachin decided on 255 cases to proceed with investigations of criminal cases of corruption against holders of federal political positions, such as deputies and senators, based on denunciation of participation in the *Lava Jato* scheme. These decisions were pushed by the public prosecutors' office's accusation, based on information retrieved from leniency deals involving those others investigated through a special criminal institution known as *delação premiada* (turns state's evidence). However, in the law establishing this form of testimony (12,850/2013[11]), secrecy is imposed on the shared information and bargain terms, aiming to preserve both the collaborator's and the investigation's integrity. The cases of April 4th, on the other hand, had important public and media impact, together with the many distinct requests from the public prosecutors' office involving the establishment and continuation of investigation on individuals, petitions for publicity of the content of delations, and the list of those investigated. Not only did Judge Fachin decide to make the information public, but justified so, in all 255 documents, using very similar language. It is likely that such a "copy-and-paste" extract of text, in which the decider argues that the publicity of the investigation information does not detriment "the defense's right, after collecting the accusatory piece, and with the means and financial resources to the adversary, the

---

[11]https://www.planalto.gov.br/ccivil_03/_ato2011-2014/2013/lei/l12850.htm

possibility of appealing against a complaint" — which *is* the topic of BP 14, although the decision, as a whole, is only marginally related to this issue — caused the models to detect some of these instances as potential applications of the BP.

We now turn to the predictions of TF-IDF+random forest, which can be considered accurate, in view of their agreement with the groundtruth curve in Fig. 6. We gather additional metadata in Fig. 9, showing that BP 14 was not enough to settle all debates about lawyers' access to investigation material, as already predicted right after its publication by [52]. By manually inspecting the decisions citing BP 14, we identified several complaints in which the defense of individuals, although not directly investigated, requested access to classified investigation material, due to some potential relation of their clients with the investigated individuals. Although unaddressed by the statement of the BP, the occurrence of this sort of complaint is frequent enough that the STF itself considered it as a subtopic of the BP[12], with decisions going in both directions. The fact that BP 14 is cited in many cases not directly related to its content fits the hypothesis *New avenue to STF or procedural issues*.

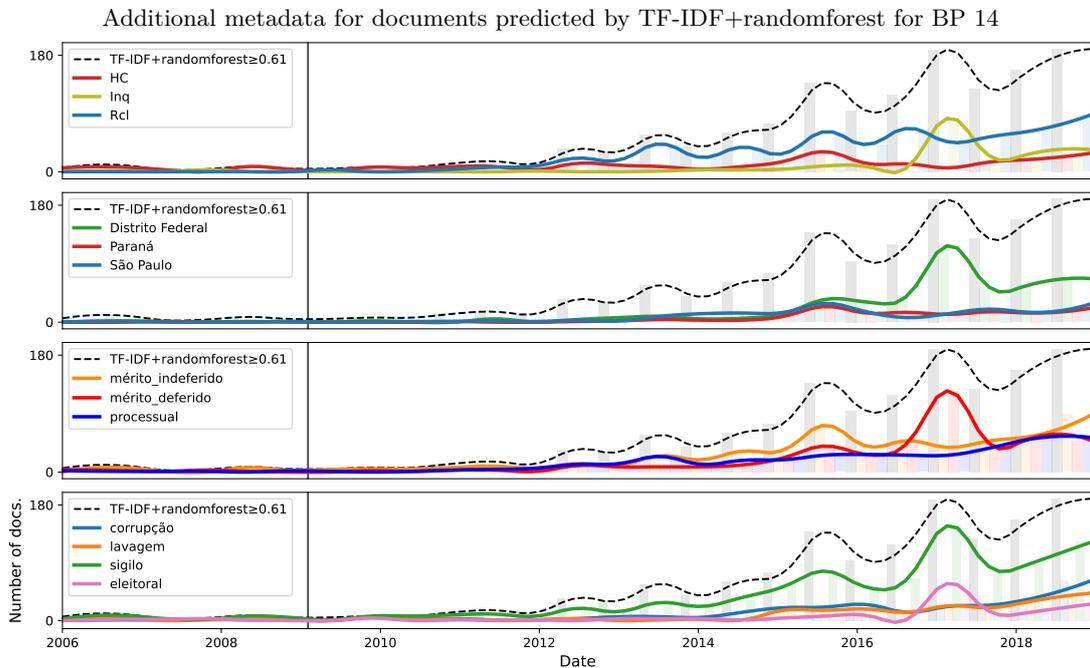Additional metadata for documents predicted by TF-IDF+randomforest for BP 14

Figure 9: **First:** Type of legal process used in the documents predicted by TF-IDF+randomforest for BP 14, among HC (*Habeas Corpus*), Inq (*Inquérito*) and Rcl (*Reclamação*). **Second:** State to which these cases belong, with label "UNKNOWN" when the metadata is unavailable. **Third:** Type of decision taken at trial: *processual* if deferred for procedural matters, or *deferido* (accepted) and *indeferido* (rejected). **Fourth:** Regex search for the presence of the words *corrupção* (corruption), *lavagem* (money laundering), *sigilo* (secrecy/confidentiality) and *eleitoral* (electoral).

In addition, and as mentioned above, it seems reasonable to associate the increasing trend of uses of BP 14 with the *lavajatism* process of the end of the 2010s, in line with our hypothesis *External conjecture not directly related to the precedent*, defined in Section 4.1. As a matter of fact, most complaints concern the investigation of corruption

---

[12]See `https://portal.stf.jus.br/jurisprudencia/sumariosumulas.asp?base=26&sumula=1230`.

cases aimed at politicians and individuals in power. In this case, not only the access to the records, or the accords of *delação premiada*, is the source of disagreement, but also the fact that the very status of the investigated individuals poses a threat to the security of informants and the process of evidence gathering[13]. Further corroboration to the hypothesis that the cause of increasing uses of BP 14 is tied to potential white-collar crimes involving politics might be gathered from the comparison of the available time series of citations with data after 2019, the end of the period of analysis of this article, because of the ever-changing policies regarding the combat against corruption in Brazil.

This showcases how the realities of the Judiciary Power may be less predictable than most juridical theories on BPs seem to suggest. Although BPs are created to address specific legal problems, their subsequent use, in actuality, might be greatly influenced by new dynamics and unforeseeable trends in Case Law. Thus, as a greater number of BPs are created over time, the complexity of the legal system tends to increase and more opportunities arise for external factors to influence the usage of this legal instrument.

4.4 BP 17

Juridical context

In Brazil, the mechanism by which the public administration (the Union, states, municipalities, and autarchies) pays debts resulting from final judicial sentences is called the *precatórios* system (court order payment). This is a privileged payment system designed to protect the public treasures, regulated by Article 100 of the Constitution[14]. More precisely, *precatórios* are payment requisitions issued by the Judiciary in favor of individuals or legal entities that have won lawsuits against the public administration. These debts can be of an alimentary nature (such as salaries, pensions, retirements, indemnities for death or disability) or non-alimentary (other debts, such as expropriations and contracts).

Article 100 establishes a set of rules on how *precatórios* should be paid. The first one concerns the order of payments: *precatórios* must be paid in the chronological order of presentation, respecting the division between alimentary and non-alimentary debts. There is a specific queue for each indebted public entity, meaning each entity has its own list of *precatórios*. However, alimentary *precatórios* have priority over non-alimentary ones. Elderly individuals (over 60 years old), those with serious illnesses, and people with disabilities also have priority in receiving amounts, up to a stipulated limit.

Another important aspect of this system relates to the functioning of public budgets in Brazil. Public entities must forecast the necessary amounts for the payment of *precatórios* in their annual budgets, under penalty of federal, state, or municipal intervention, as the case may be. Thus, Paragraph 5 of Article 100 requires public entities to include in their annual budgets the necessary amounts for the payment of judicial *precatórios* presented by April 2 of each year[15]. These *precatórios* must be paid by the end of the following fiscal year. This rule ensures that debts arising from final judicial sentences are properly budgeted for and settled within a reasonable timeframe, preserving the real value owed to the creditor.

---

[13]These recurrent situations were also perceived by the tribunal, as again indicated in `https://portal.stf.jus.br/jurisprudencia/sumariosumulas.asp?base=26&sumula=1230`.

[14]`https://www.planalto.gov.br/ccivil_03/constituicao/Constituicao.htm#art100`.

[15]This provision was originally located in Paragraph 1 of article 100 and the limit date for presenting *precatórios* was July 1st. These changes were introduced by Constitutional Amendments 62 and 114, as is further discussed ahead

Both issues — that of the order of payments and that of the processing of payments through the budget — become relevant because the *precatórios* system faces several challenges, such as delays in payments, especially in states and municipalities with a large volume of debts, and the constant need for constitutional amendments and adaptations to address default. The Binding Precedent 17 was created in November 2009 to clarify in which circumstances the Brazilian Public Treasuries needed to pay late payment interests (*juros de mora*). Specifically, the issue was whether this type of interest should apply to the period designated in the current Paragraph 5 (Paragraph 1, at the time of the BP's creation) of article 100 of the Constitution. It was debated whether late payment interest would apply to the period between the deadline for presenting the *precatórios* and the actual date in which the payment is finally concluded by the Public Treasury. The BP states that, in this period, late payments are not applicable.

> **Binding Precedent 17.** "No late payment interest is charged on court orders paid during the period stipulated in Paragraph 1 of Article 100 of the Constitution" (STF, 11/2009)
>
> **Súmula Vinculante 17.** "Durante o período previsto no parágrafo 1º do artigo 100 da Constituição, não incidem juros de mora sobre os precatórios que nele sejam pagos." (STF, 11/2009)

The legal interpretation adopted by STF in BP 17 was based on the fact that the Public Treasuries were forced to process the payments through the budget. Thus, according to the Court, the time between the provision of the payment and the actual payment could not be considered a "delay" and should not be subjected to late payment interests.

We draw the reader's attention to the fact that only one month after the creation of the BP, Constitutional Amendment 62[16] included a Paragraph 12 to Article 100, which establishes how late payment interests should be calculated, regardless of the period. Consequently, certain Justices understand this paragraph as overcoming the BP completely[17]. However, it does not seem to be the prevalent understanding of the Court, as expressed by Justice Rosa Weber in the *recurso extraordinário* (extraordinary appeal) RE 577.465/2016. Many Justices still hold that late payment interests referred to in Paragraph 12 are only applicable if the payment does not occur by the end of the next fiscal year, subsequent to the provision. Then there would be an actual "delay" in the payment, not attributable to normal processing of payments through the budget.

Mathematical results

Fig. 10 presents the predictions of the five models for BP 17, together with a regex search for *precatório* and *juros de mora* (court orders, late payment interest). Some common trends stand out: a peak between 1995-1998 and between 2002-2005, as well as an increasing tendency starting from 2014, predicted by almost all the models. Besides, in addition to 1995-1998, regex detects two other peaks: in 2007 and 2015.

The difference between these predictions is directly explained by the words correlation in Fig. 11: except TF-IDF+randomforest, none of the models accurately captured the correlation of *precatório* (court order) with the others words (*mora*, late payment). In particular, the most important peak, between 2002-2005, is principally composed of

---

[16]https://www.planalto.gov.br/ccivil_03/constituicao/Emendas/Emc/emc62.htm#art1

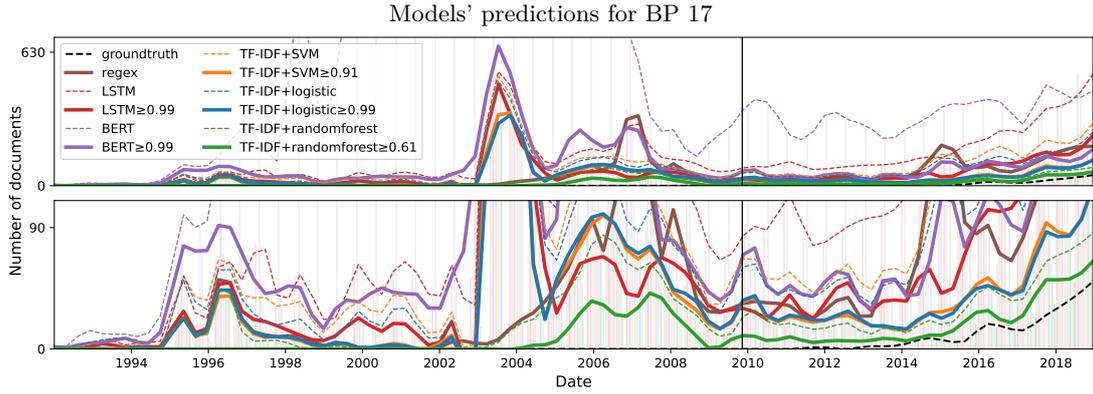[17]https://www.youtube.com/watch?v=7Rig_iyM8mY

Figure 10: Number of documents predicted by each model for BP 17 in Dataset #2, represented as a histogram (window length of six months) and interpolated via quadratic spline. We give the predictions for the thresholds adapted to Dataset #1 (dashed) and #2 (solid), as described in Table 5. Two views are given, the second one zooming in on the ordinate axis.

documents mentioning "court order" but not "late payment". This explains the absence of regex's predictions in this period. This peak, even though not referring to cases directly linked to BP 17, contains relevant information, as studied below.
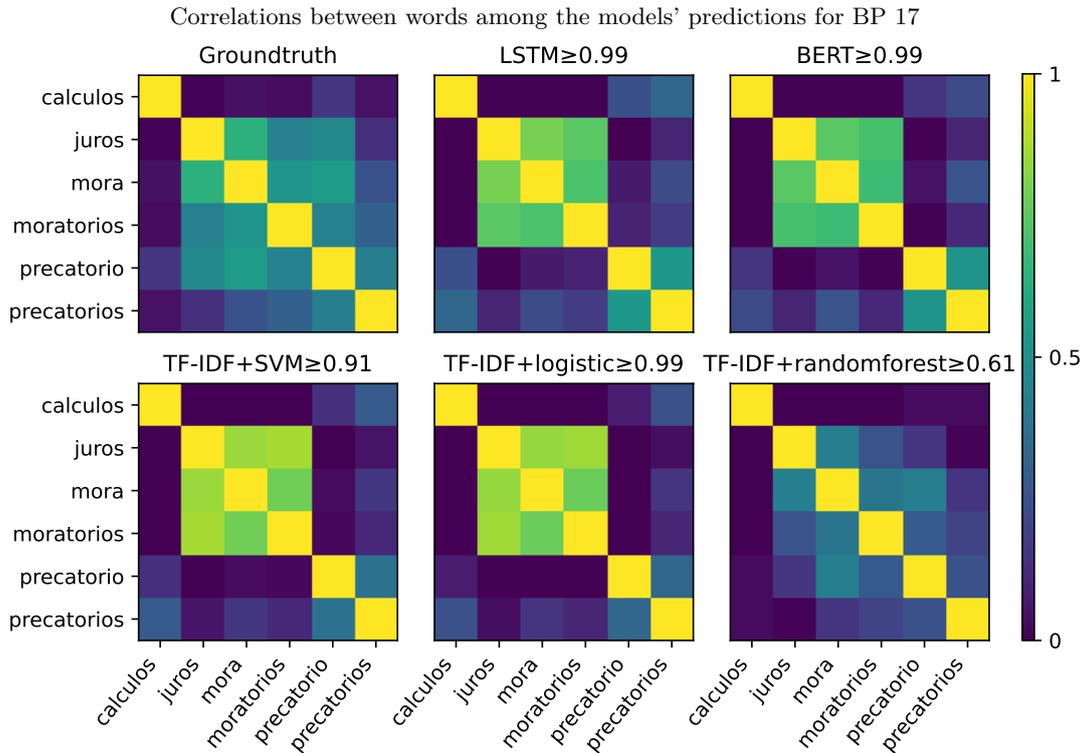

Figure 11: Correlations between the selected relevant words, when restricted to documents citing BP 17 ("Groundtruth"), or documents predicted by one of our models.

Another interesting observation is the significant correlation, in most of the models, between the terms *moratório* and *juros/mora*. The term *moratório* (moratorium) can,

indeed, be used as a synonym for *mora* (late). However, it seems that the models have imposed that predicted documents when they contain the latter, must contain the former. This would explain why regex's peak in 2007, mainly composed of documents citing *juros de mora* but not *moratório*, has not been detected by the models.

As with all the other binding precedents, BP 17 had random forest for the best model, in the sense that it predicts better the groundtruth curve, after the publication of the precedent. An additional search of words, given in Fig. 12, shows that during 2005-2009, almost no documents predicted by TF-IDF+randomforest satisfied the initial regex search. Indeed, in these documents, instead of the expression *juros de mora* are used the equivalent words *juros* and *moratório*. These are documents[18] for which BP 17 directly applies, highlighting the limitations of a naive regex search.
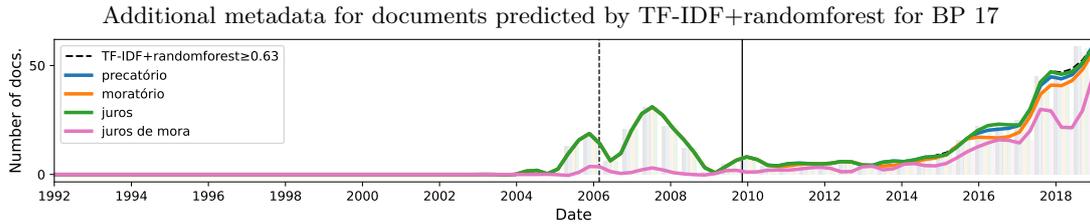

Figure 12: Presence, in the documents predicted by the model TF-IDF+randomforest for BP 17, of the words *precatório* (court order), *moratório* (moratorium), *juros* (interests) and *juros de mora* (late payment interests).

Juridical discussion

The analysis above shows that the history of BP 17's similar documents should be studied from both the models' predictions — that may use synonyms instead of exact words from the precedent — and the regex-detected documents — that do not suffer from the correlations learned by the models. We note that they both share a small peak in 1993 (see Fig. 10), representing when the matter started to be taken to the court[19].

Starting with the models' prediction, we give in Fig. 13 some additional information regarding TF-IDF+SVM. As it stands out, the 2002-2005 peak is almost exclusively composed of cases with origin in the state of São Paulo and of type *Intervenção Federal* ("Federal Intervention"). This is an exceptional measure of interference by the Union in the States, temporarily suppressing the autonomy of these entities. This event corresponds to the well-known refusal of São Paulo to pay its debts to alimentary/food creditors[20], prioritizing non-alimentary debts instead, in the early 2000s. A significant change occurred after a pivotal Supreme Court trial deciding on a Federal Intervention request in 2002, led by Attorney Antônio Roberto Sandoval Filho, and brought to trial by Justice Marco Aurélio. Though at the time there were questions about the claims brought before the court in this case — which, in fact, would not be successful —, this decision marked a turning point, drawing attention to the unsustainable nature of the

---

[18]for instance, AI 525.183 `https://redir.stf.jus.br/paginadorpub/paginador.jsp?docTP=AC&docID=620350` and AI 546.938-7 `https://redir.stf.jus.br/paginadorpub/paginador.jsp?docTP=AC&docID=501102`

[19]with a highlight to RE 149.466, in which Justice Sepúlveda Pertence concluded that *juros de mora envolvem inadimplência*, i.e., "late payment interests entail inability (to meet a financial obligation)".

[20]`https://www.sandovalfilho.com.br/ha-20-anos-pedido-de-intervencao-federal-no-estado-de-sao-paulo-comecou-a-mudar-a-historia-dos-precatorios-judiciais/`

court-ordered debt situation in Brazil. Even though we still observe a high value on the curve from 2016 (Figure 13), the payment of debts in São Paulo was considered maximum priority by the Court of Justice of the State of São Paulo[21], with more than 19 billion reais (4 billion dollars) being released for court-ordered payments in 2023[22].
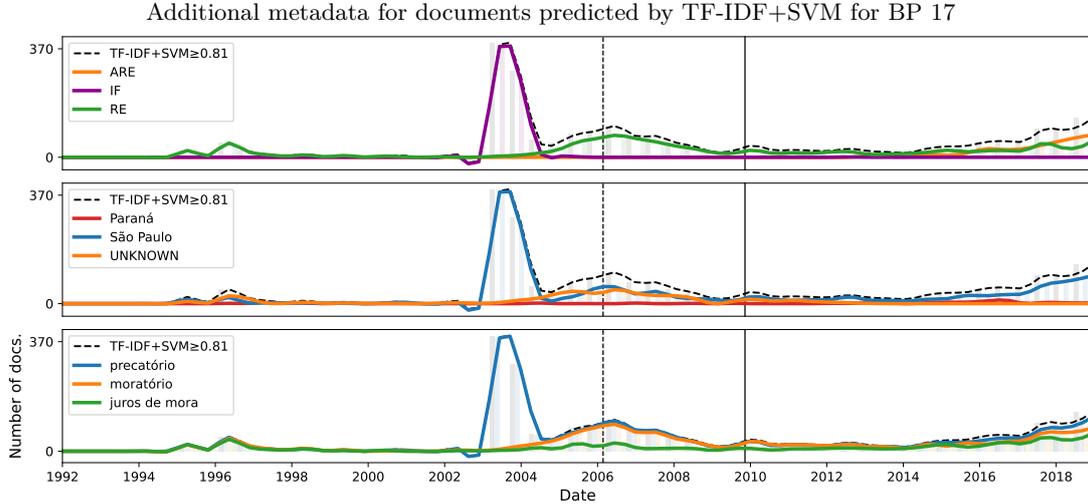

Figure 13: **First:** Type of legal process used in the documents predicted by the model TF-IDF+SVM for BP 17, among RE (*Recurso Extraordinário*), ARE (*Recurso Extraordinário com Agravo*) and IF (*Intervenção Federal*). **Second:** State to which these cases belong, with label "UNKNOWN" when the metadata is unavailable. **Third:** Presence of the words *precatório* (court order), *moratório* (moratorium) and *juros* (interests).

We now turn to regex's predictions, with Fig. 14 representing additional metadata. We notice a 2006-2007 peak regarding *militares* ("military personnel") in the State of Rio de Janeiro. A search shows that this peak mainly consists of documents employing the word *militares* which may be traced to some issues with wages[23]. This peak seems to coincide with a problem that Rio de Janeiro was dealing with in 2004, regarding the adjustment of military wage. At that time, civil public servants received a wage adjustment to overcome inflation, and Justice decided that military public servants would not be included. Further analyses are necessary to assess whether this is the case for the predicted documents, however, our methodology is useful in identifying these patterns and justifying these analyses.

In 2015, the peak of documents satisfying the regex query can be explained by ADIs 4.357[24] and 4.425[25] (*Ação Direta de Inconstitucionalidade*, Direct Action for Unconstitutionality). This relates to a very specific, technical legal discussion, regarding a transition rule introduced by Constitutional Amendment 62. The Court found Constitutional Amendment 62 to be partially unconstitutional, due to some of the rules it introduced for the order of payments — specially some of the provisions concerning preference due to old age — and, as a result, there was uncertainty regarding how the interest rate for late payments should be calculated during the transition period.

---

[21]https://www.tjsp.jus.br/Imprensa/Campanhas/Precatorios

[22]https://www.tjsp.jus.br/Noticias/Noticia?codigoNoticia=95996&pagina=1

[23]https://www.conjur.com.br/2004-jan-28/militares_rio_nao_receber_reajuste_1307/

[24]https://portal.stf.jus.br/processos/detalhe.asp?incidente=3813700

[25]https://redir.stf.jus.br/paginadorpub/paginador.jsp?docTP=TP&docID=5067184

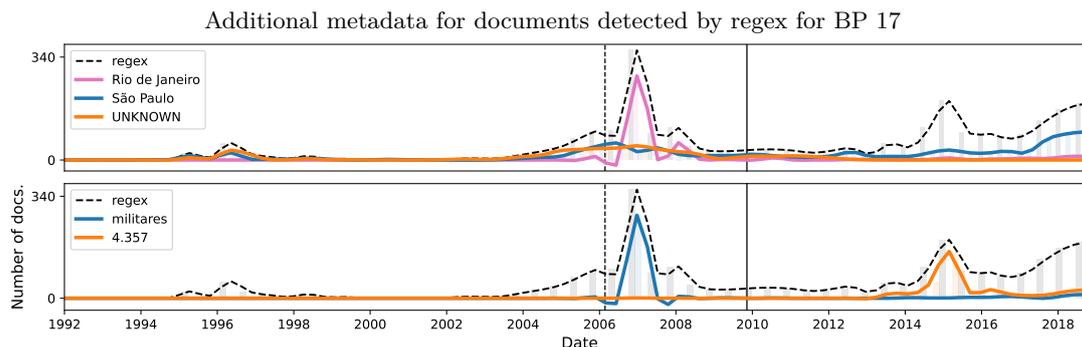Additional metadata for documents detected by regex for BP 17

Figure 14: **Top:** State to which the cases satisfying the regex query belong, with the label "UNKNOWN" when the metadata is unavailable. **Bottom:** Presence of the words *militares* (militaries) and 4.357.

The data suggests that BP 17 was partially ineffective in reducing the number of cases brought to STF. The BP was created after the major peek of cases concerning its main subject had already subsided. Soon after its creation, major changes to the underlying constitutional text were introduced by Constitutional Amendments 62 and 114, raising questions if the BP was still standing. Furthermore, by the end of the period of analysis, we may see again a tendency towards an increase in the number of cases concerning the subject of the BP. The only peek identified by the regex query after creation of the BP, in 2015, can be attributed to an external event — matters related to the effects of the declarations of unconstitutionality in ADIs 4.357 and 4.425. This corresponds to our hypothesis *External conjecture not directly related to the precedent.*

But the main lesson that can be drawn from this case seems to be that the dynamic nature of the legal system sometimes will bring changes that render a BP ineffective, or that raise questions about its validity soon after it is issued. It might be argued that, in a Civil Law system, the preferred solution to a problem of legal interpretation or implementation comes from a change in the underlying legal texts — and not through Case Law or a BP. Nevertheless, it is frequently necessary to rely on judicial solutions as they provide a faster solution to urgent problems and address concrete problems more directly. In the case of BP 17, however, we see how the dynamics of the legal system, how changes of the legal text and disputes over related topics can limit the effectiveness of this judicial tool to address legal problems and reduce litigation.

4.5   BP 26

Juridical context

In December 2009, Binding Precedent 26 was introduced to address two key legal issues, regarding individuals convicted of heinous crimes (*crimes hediondos e equiparados*), a category that encompasses offenses like torture, homicide, drug dealing, or rape. As specified in the 1990's Law n°8.072[26], a more severe treatment was imposed on these prisoners, in particular prohibiting progress through the three prison regimes, closed, semi-open, and open (Article 2, Paragraph 1). In HC 82.959/SP[27], the court found this

---

[26]Law 8.072 https://www.planalto.gov.br/ccivil_03/leis/l8072.htm

[27]HC 82.959/SP https://redir.stf.jus.br/paginadorpub/paginador.jsp?docTP=ac&docID=79206&pgI=156&pgF=160

provision to be unconstitutional. This sparked two subsequent discussions related to the determination of regime progression that would be answered in BP 26. Firstly, BP 26 clarified whether the law 11.464/07[28], allowing sentence progression for heinous crimes, applied retroactively. Secondly, it determined whether judges could require a criminological examination (*exame criminológico*) for sentence progression. The following wording was established, which would address both issues raised:

> **Binding Precedent 26.** "For the purpose of sentence progression in the case of the imprisonment for a heinous crime or an equivalent offense, the execution judge shall consider the unconstitutionality of art. 2 of Law 8.072, dated July 25, 1990, without prejudice to assessing whether the convicted individual meets the objective and subjective requirements for the benefit, and may, for this purpose, order, with reasoned justification, the performance of a criminological examination." (STF, 12/2009)

> **Súmula Vinculante 26.** "Para efeito de progressão de regime no cumprimento de pena por crime hediondo, ou equiparado, o juízo da execução observará a inconstitucionalidade do art. 2º da Lei 8.072, de 25 de julho de 1990, sem prejuízo de avaliar se o condenado preenche, ou não, os requisitos objetivos e subjetivos do benefício, podendo determinar, para tal fim, de modo fundamentado, a realização de exame criminológico." (STF, 12/2009)

However, as shown in Fig. 1, the implementation of BP 26 has been followed by certain growth of cases related to it. As reported by [3], these cases are mainly appeals (*reclamação*) aimed to challenge decisions from lower courts that had not fully embraced the new legal theories. We expect to reveal, via our methodology, additional explanations for this issue.

It is worth mentioning that the debate that led to the voting of BP 26 reveals a high level of criticism among public defenders. The dispute, especially led by the Public Defender's Office of the State of São Paulo, represented by Dr. Rafael Ramia Muneratti, focused on the second part of the binding precedent, regarding the criminological examination. It was argued that this examination, extensively rejected by psychologists and other specialists[29], would contradict the *princípio da individualização da pena* (principle of individualization of punishment). The next year, this disagreement was explicitly written by the Public Defender's Office under the name "Thesis 53"[30].

Mathematical results

We show in Fig. 15 the predictions of the five models, as well as a regex search for the words *exame criminológico* and *progressão de regime* (criminological examination, regime progression). Except for TF-IDF+randomforest, they all present a peak in 2006, as well as an increasing tendency starting from 2010. Although exhibiting similar behaviors, the five models differ a lot regarding the number of documents estimated: for 2018, LSTM estimates approximately five times as many documents as TF-IDF+SVM or TF-IDF+logistic.

---

[28] Law 11.464 `https://www.planalto.gov.br/ccivil_03/_Ato2007-2010/2007/Lei/L11464.htm`

[29] The contradiction between the conduct of criminological examination and the Psychologists' Code of Ethics is discussed in `https://www.crpsp.org/noticia/view/1846/nota-tecnica-sobre-a-suspensao-da-resolucao-cfp-0122011---atuacao-dao-psicologao-no-ambito-do-sistema-prisional`.

[30] Tese 53 da defensoria Pública do Estado de São Paulo `https://www2.defensoria.sp.def.br/dpesp/Conteudos/Materia/MateriaMostra.aspx?idItem=61257&idModulo=9706`.
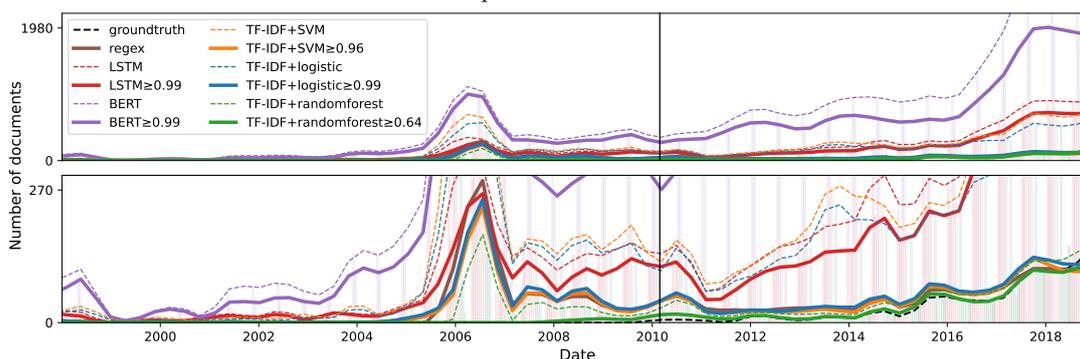
Figure 15: Number of documents predicted by each model for BP 26 in Dataset #2, represented as a histogram (window length of six months) and interpolated via quadratic spline. We give the predictions for the thresholds adapted to Dataset #1 (dashed) and #2 (solid), as described in Table 5. Two views are given, the second one zooming in on the ordinate axis.

Moreover, a manual inspection shows that the documents estimated by LSTM after 2010 are of poor quality, in the sense that they do not fit in the context of BP 26. To explain this discrepancy, we illustrate in Fig. 16 the correlations between the top TF-IDF words obtained in Table 3: *criminológico, execuções, observará, progressão, realizacão* and *regime* (criminological, executions, will observe, progression, realization, regime). It is worth noting the significant correlation between *criminológico* and *progressão* in the groundtruth image. All the models show a similar correlation. To complement this data, we give in Fig. 17 the frequency of citations of the words. As observed, LSTM tends to predict many documents solely based on the word *regime*. In fact, among the 324,756 documents of Dataset #2 that employ this word, the model predicts 5328 documents. In comparison, TF-IDF+SVM only predicts 1574. We could say that LSTM tends to detect positive cases that present the word *regime*, taking little account of its conjunction with the context of BP 26. This inevitably results in a very large number of documents, given the importance of the theme of regime progression for prisoners.

Another interesting observation is the quantity of documents detected by all the models, except TF-IDF+randomforest, during the peak between 2006 and 2008. Similarly, the regex search finds 659 documents in this period. As it turns out, a manual inspection of these documents shows that they are indeed related to the theme of BP 26 and, hence are considered correct predictions of the models. A further inspection shows that, among the most important features of the random forest (see Table 3), the words *realizacão* (realization) and *subjetivos* (subjective) are present in respectively 83.7% and 80.7% of the 241 predicted documents. In opposition, they are only present in 60.8% and 60.4% of the 659 documents detected by regex between 2006 and 2008. This suggests that this TF-IDF model has learned unnecessary restrictions, imposing the presence of words *realizacão* and *subjetivos*, resulting in too few predicted documents. In what follows, we will only consider the documents given by the models TF-IDF+SVM and TF-IDF+logistic, as these seem to be the most reliable estimations.
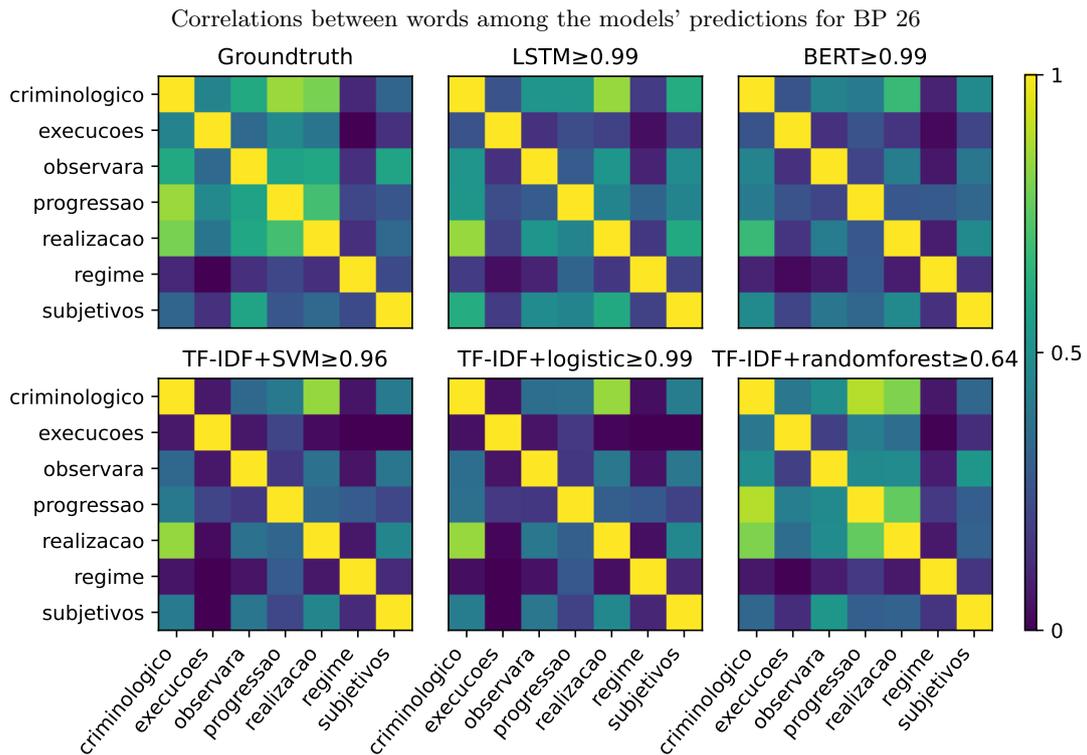
Figure 16: Correlations between the selected relevant words, when restricted to documents citing BP 26 ("Groundtruth"), or documents predicted by one of our models.
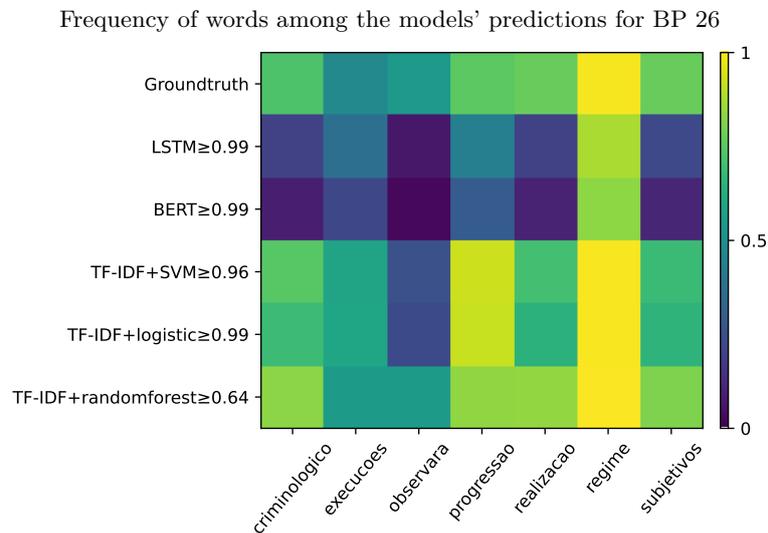

Figure 17: Frequency of the selected relevant words, when restricted to documents citing BP 26 ("Groundtruth"), or documents predicted by one of our models.

Juridical discussion

The trends of the TF-IDF+SVM or TF-IDF+logistic curves, in Fig. 15, can be mapped to exact juridical events. First, the peak around 02/23/2006 corresponds to the declaration of unconstitutionality in HC 82.959/SP, the case that launched the discussion

leading up to BP 26. The almost immediate reaction of the Brazilian Congress to address the unconstitutionality declared by this judgment, with the presentation of a law project as early as 03/23/2006, would explain the subsequent decline.

Besides, the trend of increased citations, evident after the second half of 2015, can be understood further with Fig. 18: most of the processes are appeals, emitted in the state of São Paulo. A reading of the texts suggests that from this date, a large number of complaints regarding the rationale behind the *request* for criminological examinations, which, according to the wording of the precedent, should be ordered "with reasoned justification" (*de modo fundamentado*). As evoked earlier, this matter had already been raised by the Public Defender's Office of São Paulo at the time of BP 26's debate, arguing that a precise justification for the request for criminological examination is essential. It is reasonable to hypothesize that São Paulo's defenders follow this thesis and express themselves through appeals to the STF. This phenomenon falls under the hypothesis *Resistance by a group of litigants or regional specificity* defined in Section 4.1.
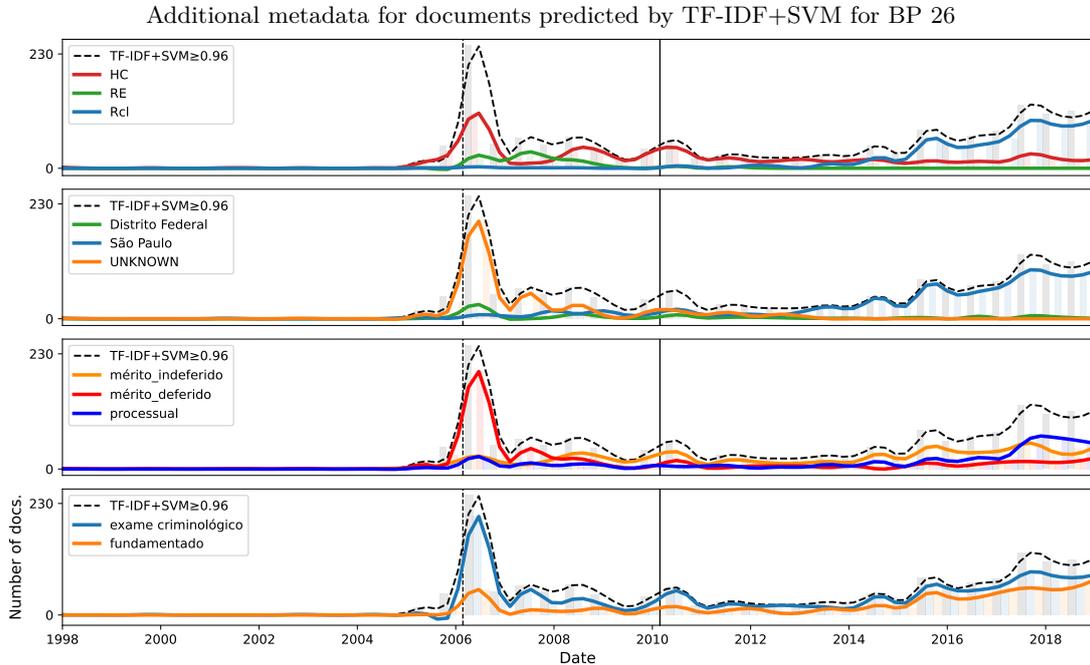

Figure 18: **First:** Type of legal process used in the documents predicted by the model TF-IDF+SVM for BP 26, among RE (*Recurso Extraordinário*), Rcl (*Reclamação*) and HC (*Habeas Corpus*). **Second:** State to which these cases belong, with label "UNKNOWN" when the metadata is unavailable. **Third:** Type of decision taken at trial: *processual* if deferred for procedural matters, or *deferido* (accepted) and *indeferido* (rejected). **Fourth:** Presence of the words *exame criminológico* (criminal examination) and *fundamentado* (grounded/founded).

Fig. 18 reveals other important information: starting from 2017, the cases tend to be deferred by the High Court based on procedural matters. More precisely, we have observed that it is common for cases to arrive in front of the judges having already lost their purpose (for instance because the convicted, in the meantime, obtained regime progression through other means). But more frequently cases are deferred because the STF considered that *reclamação* was not the correct procedural instrument. Even though the chance of an appeal being accepted might be low, we could reasonably think

that defenders continue submitting them, gambling on this chance. This constitutes a typical case of the *New avenue to the court and procedural issues* hypothesis.

In conclusion, our analysis shows that the expected effect of unifying the jurisprudence surrounding BP 26 — regime progression for convicted of heinous crimes — was not entirely achieved. The *possibility* of the progression is taken for granted, but it has simultaneously provoked a new wave of discussions, centered around the use of criminological examination, and more precisely the justification of its request. In this respect, the vague wording of BP 26 is to blame. However, we should note that vague wording is a common trait of legal provisions. Under certain conditions, it may be viewed not as a problem inherent to the production of Law, but as a quality of it. Vagueness in legal text allows for flexibility of interpretation and may be intentionally used when relevant context-specific dimensions of a legal problem are not known *ex ante.* Thus, when a general or vague term is used in Law, aspects of the solution of the problem are, in practice, deferred to be solved at a posterior time, when more information about the concrete situation is available to the parties involved. Vagueness in the legal text is both unavoidable, to a certain extent, and desirable, as precision can generate arbitrary solutions that are contradictory to the main purpose of a norm under specific or unforeseen circumstances. As Endicott [22, p. 14] puts it: "Far from being repugnant to the idea of making a norm, vagueness is of central importance to law-makers (...) It is a central technique of normative texts."

When elaborating a binding precedent, STF is in the unique position of being both the creator of a general normative text and the ultimate interpreter of that same text. This induces legal tensions that our study of BP 26 has brought to light: if the vague wording of BP 26 has allowed for more flexibility in the determination of possible uses of criminological examination, it has also caused many cases to be brought to STF for clarification, rendering one of the main purposes of a BP — that of reducing legal disputes and legal uncertainty — ineffective.

4.6   BP 37

Juridical context

Many legal cases awaiting judgment by the Justices of STF have been recognized as having general repercussion. In August 2014, the Justices ruled on one of these cases (the extraordinary appeal RE 592.317), which would release 1100 ordinary cases for judgment in lower judiciary branches. In this appeal, the STF revisited and applied one of its precedents to overrule a decision from a Brazilian State court. This (non-binding) precedent, *Súmula 339* (Precedent 339), establishes that the judiciary cannot increase the salaries of public servants on the grounds of isonomy. During the debates on the mentioned appeal, Justice Gilmar Mendes suggested converting Precedent 339 into a binding precedent, given its consistent application in STF's decisions. This way, lower judiciary branches would have to rule according to what was established by the Supreme Court, and hopefully, that would reduce the number of cases related to it that reach STF. In October 2014, BP 37 was created with the same wording as Precedent 339:

> **Binding Precedent 37.** "It is not up to the Judiciary, which has no legislative function, to increase the salaries of public servants on the grounds of isonomy." (STF, 10/2014)

Súmula Vinculante 37. "Não cabe ao Poder Judiciário, que não tem função legislativa, aumentar vencimentos de servidores públicos sob o fundamento de isonomia." (STF, 10/2014)

Long before the creation of this binding precedent, a law from 2003 (Law 10,698) was enacted to grant an increase of R$ 59.87 to all federal public servants. This resulted in 30,000 legal cases over 15 years arguing that the adjustment should not be a fixed amount but rather proportional to each servant's salary (specifically 13.23%)[31]. Different judges issued different rulings, so legal uncertainty took place. Many cases reached the STF, and the Justices, opposed to the adjustment, ended up citing Precedent 339 and later Binding Precedent 37 in their decisions. While BP 37 and repeated rulings from the STF demonstrated the unconstitutionality of such an adjustment, the controversy surrounding the issue continued to create some degree of legal uncertainty. As a result, the number of decisions citing this BP throughout time remained high.

Mathematical results

We show in Fig. 19 the predictions of the five models, as well as a regular expression search for the words *isonomia*, *vencimentos* and *servidores públicos* (named "regex" in the legend), and a search for *Súmula 339* (Precedent 339). For BP 37, we tuned the models' thresholds as those such that at least 80% are predicted, and not 90%, as is the case for BPs 11, 14 and 26 (see Table 5). In fact, BP 37 is the precedent that has the largest number of documents citing it; choosing this smaller threshold means that fewer documents are detected, and we conducted a better analysis with this value. As demonstrated in the bottom part of the figure, the TF-IDF models presented the best results, along with regex, since they faithfully follow the groundtruth curve after the BP's publication.
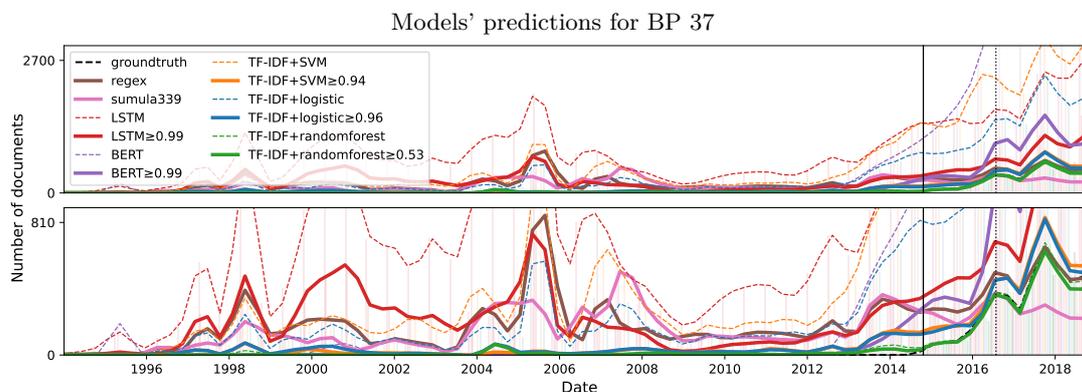

Figure 19: Number of documents predicted by each model for BP 37 in Dataset #2, represented as a histogram (window length of six months) and interpolated via quadratic spline. We give the predictions for the thresholds adapted to Dataset #1 (dashed) and #2 (solid), as described in Table 5. Two views are given, the second one zooming in on the ordinate axis. In addition, we represent a regex query for *Súmula 339*.

One observes a few common trends in the figure. Before the publication of the BP, the model LSTM, as well as the regex search, shows a clear peak in 2005. These documents

---

[31]https://www.conjur.com.br/2017-out-23/lei-deu-aumento-59-servidores-produziu-30-mil-processos

indeed contain words that characterize BP 37 and related decisions. A manual inspection shows that these documents involve Law 8,622/1993, granting an increase of 28.86% in the salary of some categories of public servants. The documents are essentially appeals, raised by servants who, although belonging to categories not foreseen by this law, were requesting the same increase. As these servants were not subject to Precedent 339, the precedent did not apply, and the High Court granted them a salary increase. In summary, these predicted documents seem somehow related to the BP's topic and cases that cite it, but do not refer to the application of Precedent 339. Ideally, they should not have been detected. The TF-IDF models did not predict them, reinforcing their good performance in the task. We point out that all the documents predicted by the TF-IDF models in 2004 actually mention Precedent 339. As this precedent and the BP have the same wording, these documents were correctly predicted.

Besides, after the publication of BP 37, two peaks in the number of predictions can be observed (in 2016 and 2017), reflected by all the models. As it turns out, these peaks can be mapped to specific juridical events, as we will study further in the juridical discussion below (see Fig. 21).

In terms of quantity, LSTM is the model with the highest number of predictions before the BP publication date, switching places with BERT afterward. Reading a sample of these documents shows that LSTM and BERT's predictions are only vaguely related to the theme of BP 37. Curiously, the TF-IDF models predicted almost no document before 2012. One could say that TF-IDF is overfitted (that is, only estimate documents citing the BP), while the other models are underfitted (that is, estimate non-relevant documents). To analyze the reasons for this discrepancy, we give in Fig. 20 the correlations between a few relevant words in the BP's wording: *aumentar, cabe, isonomia, legislativa, públicos, servidores* and *vencimentos* (to increase, fits/suits, isonomy, legislative, public, servants, salaries). As expected, the documents citing BP 37 show a significant correlation between *servidores* and *públicos* (as in "public servants", a common expression in Portuguese), but also, more interestingly, *aumentar* and *legislativa*. As observed in the figure, all the models were able to catch these correlations. However, presenting these correlations is not enough for a model to give correct predictions: for instance, only 47.6% of LSTM's predictions contain the expression "public servants", compared to 78.7%, 83.5% and 86.9% for the TF-IDF models. The large number of predictions from LSTM and BERT after 2014, following an increasing trend, simply suggests a poor generalization of these models — we recall that when validated on Dataset #1, they obtained good results (see Table 2).

To understand the low number of predictions of TF-IDF before 2014, compared to the documents citing Precedent 339, one can study the important words of TF-IDF, collected in Table 3. In particular, *isonomia* has the greatest importance, hence we expect that a large portion of the predicted documents contain this word. It is indeed the case, at a rate of 79.1%, 83.6%, and 88.9% for SVM, logistic, and randomforest. Comparatively, only 63.9% of the documents citing Precedent 339 mention this word. This suggests that the model is biased towards estimating principally documents employing *isonomia*, which is not representative of all the documents similar to BP 37, explaining the small number of documents predicted.

Juridical discussion

As a consequence of the above analysis, the predictions of TF-IDF+randomforest are only trustworthy after the publication of BP 37. Fortunately, similar documents be-
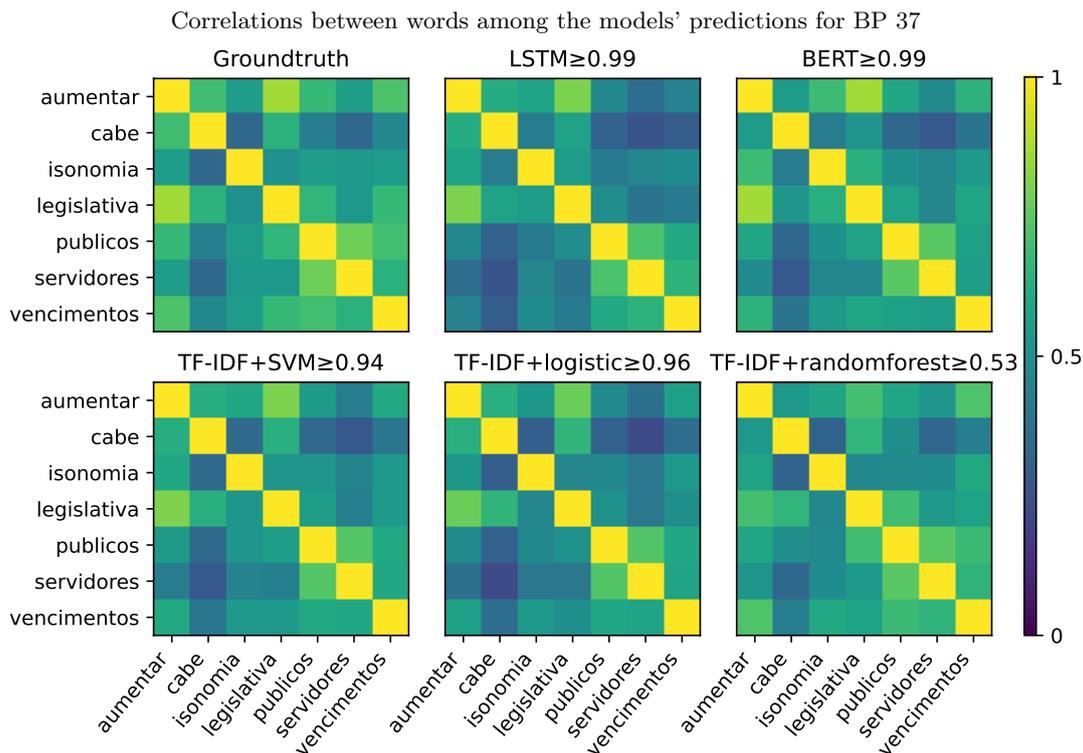
Figure 20: Correlations between the selected relevant words, when restricted to documents citing BP 37 ("Groundtruth") or documents predicted by one of our models.

fore the publication can be studied through citations of Precedent 339. We will therefore divide our discussion into two parts, first studying the predictions of TF-IDF+randomforest. We note that it would have been equivalent, as far as post-publication documents are concerned, to consider the groundtruth curve (actual citations of BP 37, see Fig. 19).

Fig. 21 gathers relevant metadata regarding TF-IDF+randomforest's predictions: the type of process, the state of emission, and the mention of the words *1987*, *10.698*, *Mogi Guaçu* and *magistratura*, which we will explain throughout the text. The first peak of the curve, in 2016, can be related to real-world juridical events as follows. As depicted in the figure, most predicted documents in that year were Extraordinary Appeals (RE) that came from the State of Rio de Janeiro (RJ). A manual inspection of some of these documents shows the application of BP 37 in the context of the RJ State Law 1,206/1987. This is reflected by the citation curve for "1987". Given that this law was enacted to increase the salaries of public servants from the Executive and Legislative branches, servants from the Judiciary questioned their exclusion by taking legal action. After citing BP 37 in several decisions, STF issued a final ruling dismissing the servants' request in October 2016[32].

The second peak of decisions, from 2017 onwards, can be understood as a combination of several initiatives. The first reason is Law 13,317/2016 which altered the career and salary structure of public servants. This led to a significant increase in cases claiming a 13.23% salary adjustment, based on Law 10.698/2003, as shown by the citation curve

---

[32]General repercussion thesis ARE 909,437 `https://redir.stf.jus.br/paginadorpub/paginador.jsp?docTP=TP&docID=11828219`

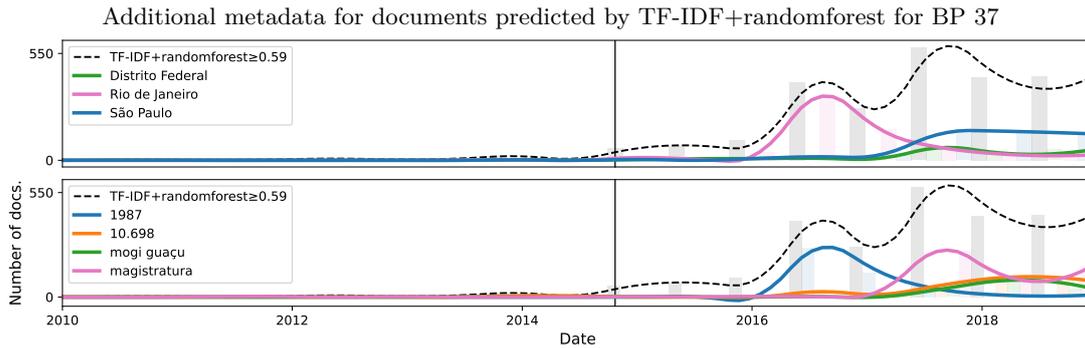Additional metadata for documents predicted by TF-IDF+randomforest for BP 37

Figure 21: **Top:** State to which the legal process used in the documents predicted by TF-IDF+randomforest for BP 37 belong, among the top three. **Bottom:** Documents containing the words *1987*, *10.698*, *Mogi Guaçu* or *magistratura*.

in Fig. 21 (see citations of "10.698"). As a response, Justice Gilmar Mendes proposed in 2017 the creation of a new and specific binding precedent that would explicitly prohibit the granting of the mentioned adjustment[33]. After further analysis of the proposition, which considered the Prosecutor General's manifestation[34] and others, the creation of this new BP was rejected.

In addition, the many mentions of "Mogi Guaçu", a municipality in São Paulo, can be attributed to a collective agreement signed between the municipality and workers' syndicate (*Lei Complementar Municipal 1.121/2011*[35]), determining the incorporation of bonuses into the salaries of all municipal employees and civil servants. A somewhat similar movement can be observed through the documents citing *magistratura* (members of the judiciary, such as judges), which relates to a controversy concerning the *licença prêmio*, a bonus awarded every five years to public servants (federal, municipal or state). Certain members of the judiciary invoked the principle of isonomy to obtain this bonus, which was rejected, according to the terms of BP 37[36].

In conclusion, the four trends observed on the citations of BP 37 — salary increases for public servants from Law 1,206/1987 or Law 13,317/2016, the Mogi-Guaçu syndicate and *licença prêmio* for judiciary — are protests led by specific groups, as formulated by our hypothesis *Resistance by a group of litigants or regional specificity.*

Next, we turn to the analysis of BP 37's similar documents before its publication, through those that mention Precedent 339. We remind the reader that the latter is a direct descendant of the former, hence they share similar documents. Fig. 22 contains metadata relative to these documents.

A first trend can be observed between 1997 and 2002: most of the documents mention the state "Santa Catarina", as visualized on the blue curve. A manual inspection revealed that most documents refer to two general cases that discuss permanent increases in the salaries of state servants of Santa Catarina who have temporarily held commissioned positions (Complementary State Law 43/1992 and State Law 9,847/1995). Thus, citing Precedent 339 in its arguments, the State of Santa Catarina filed extraordinary

---

[33]https://www.conjur.com.br/dl/psv-128-reajuste-1323.pdf

[34]https://portal.stf.jus.br/processos/downloadPeca.asp?id=312747008&ext=.pdf

[35]https://sistema.camaramogiguacu.sp.gov.br/consultas/norma_juridica/norma_juridica_mostrar_proc?cod_norma=5779

[36]https://portal.stf.jus.br/jurisprudenciaRepercussao/verPronunciamento.asp?pronunciamento=7188706
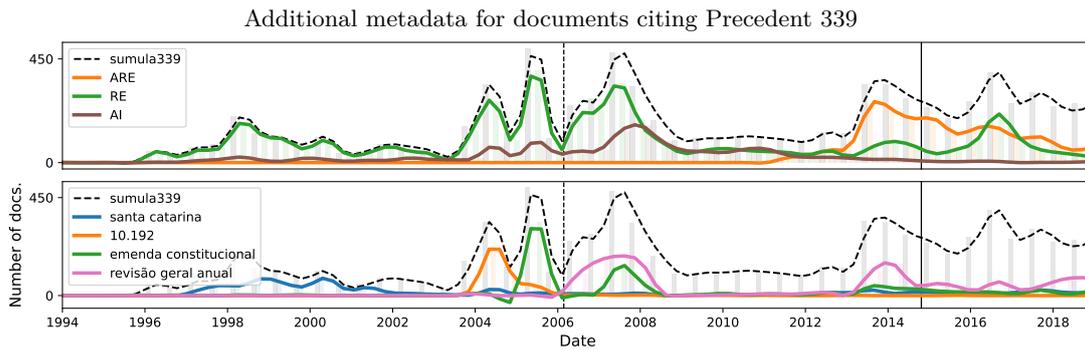
Figure 22: **Top:** Type of legal process used in the documents citing Precedent 339, among RE (*Recurso Extraordinário*), ARE (*Recurso Extraordinário com Agravo*) and AI (*Agravo de Instrumento*). **Bottom:** Search for the presence of the words *Santa Catarina* (the state), *10.192*, *emenda consitucional* (constitutional amendment) or *revisão geral anual* (annual general revision).

appeals to the STF against lower court decisions that guaranteed to the servants (i) the right to the permanent increase or (ii) equal salary adjustment to those servants receiving such an increase due to past commissioned positions and those who are currently in equivalent positions.

To continue, one observes a significant rise in citations between 2003 and 2010, made of three peaks. They are all linked to specific event, represented by the presence of the words *10.192*, *emenda consitucional* (constitutional amendment) or *revisão geral anual* (annual general review). Regarding the first one, a manual inspection revealed they refer to the application of Precedent 339 in the context of Law 10,192/2001, a law that increased the salaries of non-public servants (10,87%) and which was used by public servants to require a similar increase.

As far as the 2005 peak is concerned, it must be traced back to 1993. Indeed, that was the year the Laws 8,622 (already mentioned) and 8,627 came into effect, adjusting the salaries of higher ranks in the armed forces by 28.86%. STF Justices then recognized the same right to lower-ranking military personnel, understanding it was a case of general revision of salaries, therefore not applying Precedent 339 but rather Article 37, Section X of the Constitution. The Union filed extraordinary appeals[37] (*Recursos extraordinários*), and one of the arguments was that the Constitution allowed for different adjustments for different ranks (Article 142), but the Justices upheld their decision based on Article 37. Article 37 was amended by Constitutional Amendment 19/1998[38], which established the necessity of a specific law for any form of salary setting or change, thus exempting the requirement for adjustments to consistently follow the same percentage. In light of this, the Federal Government requested, and the STF accepted, that the 28.86% adjustment should be valid only until June 1998, when the Constitutional Amendment came into effect.

We now turn to the last peak. In addition to the need for a specific law for any form of setting or change in the salary of public servants, Constitutional Amendment 19/1998 also established that the general revision of salaries must occur annually and based on annual laws proposed by the Head of the Executive branch (President of the

---

[37]such as AI 551.744 https://portal.stf.jus.br/processos/detalhe.asp?incidente=2308620

[38]https://www.planalto.gov.br/ccivil_03/constituicao/Emendas/Emc/emc19.htm#art3

Republic or state governors). A manual inspection of documents between 2006 and 2008 found extraordinary appeals[39] requesting compensation for public servants due to omissions in issuing such a law. STF denied these appeals based on Precedent 339. After many years of judgments on this matter, STF stated in 2019 that not submitting a law's proposition regarding the annual revision of public servants' salaries does not create a subjective right to compensation; however, the Executive branch must provide justification for why the revision was not proposed (STF *Tema* 19[40]) [56].

In conclusion, our analysis shows that the tumultuous history of Precedent 339, marked by various controversies led by public servants (seen in Fig. 22), continues in a similar way through its new avatar, BP 37 (as seen in Fig. 21), corresponding to the hypothesis *Resistance by a group of litigants or regional specificity.*

It has also created, due to the persistence of this resistance, an influx of cases to the higher courts, not only to STF, but also to STJ (*Superior Tribunal de Justiça*), corresponding to the hypothesis *New avenue to the court.* As the controversy around this issue ensues, new cases continue to be brought before STF, as new attempts to adjust the salary of public servants occur in the absence of formal laws. The contradiction of this never-ending administrative practice to the content of the BP, in itself, creates a clear avenue for the issue to be brought up time and time again before the court. In this case, however, it is interesting to note that this new avenue is not a consequence of lawyers trying to bend formal restrictions or take advantage of indirect legal connections to gain access to the court. Rather, this avenue is a result of direct confrontation between persevering administrative practices and what has been determined as the correct legal interpretation of the Law by STF.

4.7 Final legal considerations

Empirical insights into the legal debate on binding precedents

As discussed in Section 2.1, a BP is understood as a legal tool for standardizing judicial decisions [62]. The previously mentioned requirements for the creation of BPs, as set forth in Paragraph 1 of Article 103-A of the Federal Constitution, reinforce the understanding that this instrument aims to unify jurisprudence, addressing recurring controversies that cause uncertainty or legal insecurity, thus reducing the volume of litigation. Indeed, it is expected that the BP, by granting binding effects to the consolidated understanding of the court on controversial issues, will produce desirable effects for the legal system [37]. Among such effects, the following can be highlighted: (i) the reduction of litigation arising from doubts about the correct application of the law, thereby accelerating the resolution of conflicts and increasing the efficiency of judicial activity; (ii) clearer guidance for social agents on how to conduct their legal relations, which in turn would discourage behaviors incompatible with the prevailing legal order; and (iii) the promotion of equity and formal equality, by restricting the decisions rendered by lower courts (which, in turn, reduces the likelihood of similar legal problems receiving divergent judicial solutions) [64].

Moreover, the BP can also enhance the effectiveness of constitutional norms and strengthen the role of the STF as the court responsible for constitutional interpretation [63]. In this regard, the instrument represents a natural course of evolution in Brazil, given the increasing importance acquired by Constitutional Law since 1988, both for the

---

[39]such as RE 524.145 `https://portal.stf.jus.br/processos/detalhe.asp?incidente=2461438`

[40]`https://portal.stf.jus.br/jurisprudenciaRepercussao/tema.asp?num=19`

interpretation and application of infraconstitutional norms across various branches of law and for the realization of fundamental rights provided for in the Federal Constitution itself [1, 63].

Much of the academic and doctrinal discussion that accompanied the introduction of the BP in Brazil has focused on understanding how this instrument affects the role of the STF — and the judiciary more broadly — in the country [57, 18, 35, 1]. There was debate over whether the BP, a legal instrument typically associated with the common law system, would be compatible with Brazilian law, which is rooted in a civil law tradition. Additionally, questions were raised about the extent to which the BP differs from the legislative function and how the introduction of such a prerogative for the STF would impact the Brazilian separation of powers [49, 15, 9]. It should be noted that these are essentially normative issues, aimed at assessing whether the BP is a desirable instrument and determining what its contours, implementation limits, and main characteristics should be.

However, the main justifications for the introduction of the BP in the Brazilian constitutional system are predicated on empirical claims. After all, the instrument results from an effort to standardize constitutional jurisprudence and contain the demands that reach the STF, given the high volume of litigation the court must handle every year. This effort, in fact, not only led to the introduction of the BP but also to other instruments aimed at reducing the number of cases judged by the STF. Despite this fact, there are few empirical studies that seek to verify whether the BPs produce the expected effects — particularly studies that aim to ascertain whether there is indeed a standardization of jurisprudence and a reduction in litigation concerning the subject of the BP. The reasons for this are partly methodological, as discussed earlier in this text.

Thus, this research, by proposing a methodology that allows measuring the volume of litigation pertaining to the subject of a BP from before its creation, aimed only to confirm the results typically predicted by legal doctrine — that is, to empirically demonstrate that the introduction of a binding precedent indeed reduces the number of cases judged by the STF on the subject in question. However, the research results revealed a counter-intuitive scenario that has not received appropriate attention in the legal field. The investigated cases showed an increase in the number of decisions addressing the theme or issue covered by the BP after its introduction. We proposed some hypotheses to explain this result, whose legal significance points to interesting issues in constitutional theory that still deserve to be the subject of applied studies in the future.

The first and most frequent explanatory hypothesis for the increase in litigation following the creation of an BP is what we refer to as the creation of a *New avenue to STF or procedural issues*. This hypothesis is relevant for explaining almost all the cases studied[41]. This arises from the fact that the judicial system involves actors who act strategically. It is often a mistake to assume, when seeking a social outcome through a legal change, that the recipients of that change will continue to behave the same way they did before its introduction. Generally, it is necessary to understand how a change in the law will alter the incentives of the actors involved. The claim that a judicial decision is not in accordance with an BP allows the case to reach the STF, creating an interest among litigants to address this type of issue precisely to gain access to the court. In other words, by establishing an BP on a particular subject, the STF increases the constitutional relevance of the issue, signaling to litigants that it will hear cases

---

[41]It is less present only in BP 17, but was explicitly identified as a relevant cause of litigation in BPs 11, 14, 26, and 37.

involving violations of that BP.

Therefore, the establishment of an BP always creates incentives for parties to invoke the issue of the BP to access the court, even if their cases' connection to the subject of the BP is weak or indirect. The same process, in fact, can be observed with the creation of legal rules through the usual legislative process: it is a well-known legal phenomenon that the multiplication of normative acts establishing specific norms tends to increase the complexity of the existing legal system, generating more legal uncertainty and litigation, instead of providing clarity to the law applicable to each specific case. The language of the BP itself can introduce new complexities if the way the BP addresses its subject contains ambiguities or vague terms that give rise to new judicial discussions.

In our analysis, we found that the hypothesis we called *Vague wording* was clearly related only to BP 26, but this case highlighted some peculiarities regarding the problem of vagueness in BPs. As previously noted, the use of vague terms plays an important role in the production of legal norms, allowing different social situations to be addressed by the norm, even when the author of the normative text does not know the specifics of each particular case to which the norm applies. However, in the case of the STF, there is a rather peculiar situation: the body that creates the vague normative text will also be the one that has the final word on how the text should be interpreted and applied. The contemporary state has other bodies with normative and judicial functions, but the decisions issued by these bodies can always be reviewed judicially. In the case of BPs, it will always be up to the STF, ultimately, to resolve controversies regarding its own binding precedents.

The semantic openness of normative statements is also what allows groups of litigants to resist the application of an BP or for divergent interpretations to arise in specific lower courts, creating regional specificities in how the BP is addressed. As a rule, these points of resistance need to mobilize some ambiguity or legal gap left by the BP. The hypothesis we called *Resistance by a group of litigants or regional specificity* proved relevant in two cases: BP 26 and BP 37. In these cases, we see that the creation of an BP will not always have the effect of pacifying social relations and reducing conflicts; it can also incite more litigation and opposition from those affected, especially when the content of the BP does not seem to adequately address a particular case and the agents seek to differentiate it from the original scope of application of the BP (the well-known *distinguishing* strategy).

External circumstances to the judicial dynamics can also create problems for the application of the BP. An unpredictable social or economic event that stimulates new conflicts or affects the conditions for the application of the BP — or even changes to the constitutional text itself — can lead to new peaks of litigation on the BP's subject. Two cases were also identified where the hypothesis we called *External conjecture not directly related to the precedent* proved relevant. These are BP 14 and BP 17. In these cases, we observe how the influx of litigation can be related to factors beyond the control of the STF. The BP is intended to govern a complex, dynamic, and somewhat unpredictable socio-economic reality [62]. When deciding to create a new BP, the STF relies on the cases that have reached the court up to that moment, which are also the result of changing social dynamics. Therefore, the court's perspective on BP issues is subject to biases.

Finally, there was a case (BP 11) where the issue had not been recurrently debated before the introduction of the BP, a hypothesis we call *New theme*. In this case, the STF determined that a practice, which had not been questioned before the court, constituted a significant constitutional violation. The BP was used, therefore, as a way to intervene in social reality to ensure compliance with the Constitution. This is a particularly

relevant use of the BP because it relates to a broader debate about the increasing role of the judiciary. The methodology developed by this research allows for measuring the cases in which the use of the BP does not directly result from recurrent societal demands, but rather from a more assertive action by the court to ensure the effectiveness of the Constitution.

All five hypotheses listed in this study suggest that the establishment of an BP can, in certain cases — or possibly in most cases — result in an increase in the number of litigations, raising new controversies and legal issues that demand judicial solutions. This problem has already been discussed theoretically by legal doctrine, but it has not yet been measured, despite being based on claims that are, as stated, essentially empirical. We hope that the methodology proposed here will provide important contributions to a better understanding of the impacts of BPs and the possible causes of increased litigation after the issuance of an BP.

Limitations and potential improvements

It is important to note that our entire analysis, which claims to uncover the mechanisms underlying the repetitive use of a binding precedent, is significantly biased by the fact that we perform Similar Case Retrieval only on the documents brought up to trial at the Supreme Court. In particular, our method is vulnerable to a survivorship bias, unaware of themes that are confined to the lower courts, or not taken to court at all (e.g., theses from advocacy groups, reports from various agencies, political events).

This problem can be mitigated in at least three ways. The first, already mentioned above, would be to acquire data from courts below the STF. In the case of STJ (Superior Court of Justice), data is available on the official website[42]. This way, the authors of [41] extracted 129,602 cases, enabling an analysis of corporate law precedents. For other courts, however, a search engine is not available, and we are not aware of any public database — the data acquisition must be done by asking each court.

Another important avenue would be a field survey, involving interviews with legal professionals (e.g., lawyers, public defenders, prosecutor, or judges). In this regard, we had the opportunity to speak with members of the Public Defender's Office and Public Prosecutor's Office of the State of São Paulo, who kindly presented to us, from their point of view, the important issues at stake in the creation of the BPs (more specifically BP 26). However, the development of a sociological analysis, which would take into account the various political and legal positions with regard to Supreme Court precedents, would require the collection of many other points of view.

Finally, our analysis would benefit from the incorporation of expert annotations in the training set, in order to isolate key elements of the documents. Indeed, our models are trained in such a way that all parts of the text are considered equally; this potentially introduces an important bias, by encouraging the integration of elements that are too specific to certain cases, or by giving equal importance to cases judged solely on procedural grounds. Such an approach is expected to refine our models, which suffer, as our work has shown, from a considerable generalization error. In this regard, and given the considerable number of documents studied, it would be interesting to produce automatic annotations using Large Language Models.

---

[42]`https://www.stj.jus.br/sites/portalp/Processos/Consulta-Processual`

5 Conclusion

The Brazilian Federal Supreme Court routinely deals with an overwhelming number of cases with repetitive demands. While there is a juridical instrument specially created to reduce the number of repeated demands through the establishment of a normative jurisprudence — the binding precedent —, this instrument often fails in this regard.

In this article, we have performed a series of mathematical and juridical analyses that allowed us to explain why BPs 11, 14, 17, 26, and 37 (some of the most cited in the STF's decisions) fail in reducing STF's workload. First, we have applied and compared different Similar Case Retrieval methods (TF-IDF-based models, LSTM, BERT, and regular expressions) on a database composed of STF's decisions. By studying time series of similar cases, our methodology enabled us to assess the impact of laws on jurisprudence, and to perform an empirical study of the juridical mechanisms behind these BPs' inefficiency. We finally described five main hypotheses that explain the large number of cases reaching STF.

From the mathematical point of view, the TF-IDF models used in our analyses performed better than LSTM and BERT. This was observed both on the annotated dataset (Dataset #1), where the scores of the first models are higher, and on the larger dataset (Dataset #2), where the deep learning models face a considerable generalization error. Although we have used a Portuguese-specific BERT model, fine-tuned on a large collection of legal documents, we plan to incorporate further improvements: circumventing the truncation of input text to the first 510/512 tokens (e.g., via pooling [76] or Big Bird [74]), and fine-tuning to our specific dataset. In addition, we intend to investigate whether the incorporation of expert-annotated key parts would lead to better results or new insights.

From a legal perspective, our study uncovers a counter-intuitive effect, that has been overlooked in purely doctrinal works: the very nature of the binding precedent brings with it reasons for increasing repetitive demands. This includes the fact that, when a BP is published, it is declared to be a new topic of relevance that the Supreme Court is committed to dealing with; that it can lead to the formation of diverse groups of litigants; that it introduces new complexities and potentially ambiguities; that it exposes and legitimizes a new discussion topic; and that it is naturally subject, in an indirect way, to other economic, political or social conjectures. In order to refine our study, which is based solely on the analysis of STF decisions, we plan to integrate somehow lower court decisions, as well as conducting a field survey to gather the views of various legal professionals.

References

[1] Walber de Moura Agra. A reconstrução da legitimidade do Supremo Tribunal Federal: densificação da jurisdição constitucional brasileira. Forense, Rio de Janeiro, 2005.

[2] Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preoţiuc-Pietro, and Vasileios Lampos. Predicting judicial decisions of the european court of human rights: A natural language processing perspective. PeerJ Computer Science, 2:e93, 2016.

[3] Thiago Bottino do Amaral. Habeas corpus nos tribunais superiores: uma análise e proposta de reflexão, 2016.

[4] Vithor Gomes Ferreira Bertalan and Evandro Eduardo Seron Ruiz. Predicting judicial outcomes in the brazilian legal system using textual features. In DHandNLP@ PROPOR, pages 22–32, 2020.

[5] Vithor Gomes Ferreira Bertalan and Evandro Eduardo Seron Ruiz. Using attention methods to predict judicial outcomes. Artificial Intelligence and Law, pages 1–29, 2022.

[6] Paheli Bhattacharya, Kripabandhu Ghosh, Arindam Pal, and Saptarshi Ghosh. Legal case document similarity: You need both network and text. Information Processing and Management, 59(6):103069, 2022.

[7] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. Journal of machine Learning research, 3(Jan):993–1022, 2003.

[8] Alexandre Freitas Câmara. Manual de direito processual civil. São Paulo: Atlas, 2022.

[9] Mauro Cappelletti. Juízes legisladores? Sergio A. Fabris, Editor, Porto Alegre, 1993.

[10] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. Universal sentence encoder. arXiv preprint arXiv:1803.11175, 2018.

[11] Yahui Chen. Convolutional neural network for sentence classification. Master's thesis, University of Waterloo, 2015.

[12] Kenneth Ward Church. Word2vec. Natural Language Engineering, 23(1):155–162, 2017.

[13] Fernando A. Correia, Alexandre A.A. Almeida, José Luiz Nunes, Kaline G. Santos, Ivar A. Hartmann, Felipe A. Silva, and Hélio Lopes. Fine-grained legal entity annotation: A case study on the brazilian supreme court. Information Processing and Management, 59(1):102794, 2022.

[14] Fernando A Correia, José Luiz Nunes, Guilherme FCF de Almeida, Alexandre AA Almeida, and Hélio Lopes. An exploratory analysis of precedent relevance in the brazilian supreme court rulings. In Proceedings of the ACM Symposium on Document Engineering 2019, pages 1–4, 2019.

[15] Sérgio Sérvulo da Cunha. O efeito vinculante e os poderes do juiz. Saraiva, São Paulo, 1999.

[16] Flávia Pires Da Silva. A influência política na edição da súmula vinculante nº 11–algemas. Centro Universitário do Distrito Federal, 2009.

[17] Diego Oliveira da Silveira and Gilberto Schäfer. Análise crítica da súmula vinculante do uso das algemas. Revista Brasileira de Direitos Fundamentais & Justiça, 8(28):229–249, 2014.

[18] Dalmo de Abreu Dallari. O poder dos juízes. Saraiva, São Paulo, 1996.

[19] Jingpei Dan, Lanlin Xu, Weixuan Hu, Yuming Wang, and Yingfei Wang. Enhancing similar case matching with event-context detection in legal intelligence. In Fei Liu, Nan Duan, Qingting Xu, and Yu Hong, editors, Natural Language Processing and Chinese Computing, pages 549–561, Cham, 2023. Springer Nature Switzerland.

[20] Pedro Henrique Luz De Araujo, Teófilo Emídio de Campos, Fabricio Ataides Braz, and Nilton Correia da Silva. Victor: a dataset for brazilian legal documents classification. In Proceedings of the Twelfth Language Resources and Evaluation Conference, pages 1449–1458, 2020.

[21] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805, 2018.

[22] Timothy A. O. Endicott. The value of vagueness, pages 14–30. Oxford University Press, Oxford, New York, 2011.

[23] Joaquim Falcão, Pablo de Camargo Cerdeira, and Diego Werneck Arguelhes. I relatório supremo em números-o múltiplo supremo. Escola de Direito do Rio de Janeiro da Fundação Getulio Vargas, 2012.

[24] William Paulo Ducca Fernandes, Luiz José Schirmer Silva, Isabella Zalcberg Frajhof, Guilherme da Franca Couto Fernandes de Almeida, Carlos Nelson Konder, Rafael Barbosa Nasser, Gustavo Robichez de Carvalho, Simone Diniz Junqueira Barbosa, and Hélio Côrtes Vieira Lopes. Appellate court modifications extraction for portuguese. Artificial Intelligence and Law, 28(3):327–360, 2020.

[25] James H. Fowler and Sangick Jeon. The authority of supreme court precedent. Social Networks, 30(1):16–30, 2008.

[26] Filippo Galgani, Paul Compton, and Achim Hoffmann. Citation based summarisation of legal texts. In PRICAI 2012: Trends in Artificial Intelligence: 12th Pacific Rim International Conference on Artificial Intelligence, Kuching, Malaysia, September 3-7, 2012. Proceedings 12, pages 40–52. Springer, 2012.

[27] Aurélien Géron. Hands-on machine learning with scikit-learn, keras, and tensorflow, 2022.

[28] Ben Hachey and Claire Grover. Extractive summarisation of legal texts. Artificial Intelligence and Law, 14:305–345, 2006.

[29] Hiroaki Hayashi, Zecong Hu, Chenyan Xiong, and Graham Neubig. Latent relation language models. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 34-05, pages 7911–7918, 2020.

[30] Zhilong Hong, Qifei Zhou, Rong Zhang, Weiping Li, and Tong Mo. Legal feature enhanced semantic matching network for similar case matching. In 2020 International Joint Conference on Neural Networks (IJCNN), pages 1–8, 2020.

[31] Jon M. Kleinberg. Authoritative sources in a hyperlinked environment. J. ACM, 46(5):604–632, sep 1999.

[32] Sushanta Kumar, P Krishna Reddy, V Balakista Reddy, and Aditya Singh. Similarity analysis of legal judgments. In Proceedings of the fourth annual ACM Bangalore conference, pages 1–4, 2011.

[33] Quoc Le and Tomas Mikolov. Distributed representations of sentences and documents. In International conference on machine learning, pages 1188–1196. PMLR, 2014.

[34] Fernando Leal and Ana Paula Barllos. Vii relatório supremo em números: os precedentes no supremo, 2020.

[35] Victor Nunes Leal. Passado e futuro da súmula do stf. Revista de Direito Administrativo, 145:1–20, 1981.

[36] Greg Leibon, Michael Livermore, Reed Harder, Allen Riddell, and Dan Rockmore. Bending the law: geometric tools for quantifying influence in the multinetwork of legal opinions. Artificial Intelligence and Law, 26, 06 2018.

[37] Rodolfo de Camargo Mancuso. Divergência jurisprudencial e súmula vinculante. Revista dos Tribunais, São Paulo, 1999.

[38] Rachele Mignone, Ivan Spada, Chiara Bonfanti, Michele Colombino, Giorgia Iacobellis, Marianna Molinari Laurentiu Jr Marius Zaharia, Susanna Marta, Ilaria Angela Amantea, Emilio Sulis, Luigi Di Caro, et al. Augmented reading and similar case matching: from legal domain experts' modus operandi to a computational pipeline. CEUR Workshop Proceedings (CEUR-WS, 2021.

[39] Alexandre de Moraes. Direito constitucional. Atlas, São Paulo, 2022.

[40] Thom Neale. Citation analysis of canadian case law. J. Open Access L., 1:1, 2013.

[41] José Nunes and Ivar Hartmann. A quantitative approach to ranking corporate law precedents in the brazilian superior court of justice. Artificial Intelligence and Law, 30, 03 2022.

[42] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), pages 1532–1543, 2014.

[43] Thomaz Pereira, Diego Werneck Arguelhes, and Guilherme da Franca Couto Fernandes de Almeida. Viii relatório supremo em números: Quem decide no supremo?: tipos de decisão colegiada no tribunal, 2020.

[44] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers), pages 2227–2237, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.

[45] Matthew E Peters, Mark Neumann, Robert L Logan IV, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A Smith. Knowledge enhanced contextual word representations. arXiv preprint arXiv:1909.04164, 2019.

[46] K Raghav, P Krishna Reddy, and V Balakista Reddy. Analyzing the extraction of relevant legal judgments using paragraph-level and citation information. AI4JCArtificial Intelligence for Justice, page 30, 2016.

[47] Lucas Emanuel Resck, Jean R Ponciano, Luis Gustavo Nonato, and Jorge Poco. Legalvis: Exploring and inferring precedent citations in legal documents. IEEE Transactions on Visualization and Computer Graphics, 2022.

[48] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16, pages 1135–1144, New York, NY, USA, August 2016. Association for Computing Machinery.

[49] Walter Claudius Rothenburg. A dialética da democracia: entre legisladores e jurisdição constitucional, 2007.

[50] Haşim Sak, Andrew Senior, and Françoise Beaufays. Long short-term memory based recurrent neural network architectures for large vocabulary speech recognition, 2014.

[51] Luis Felipe Salomão. Artificial inteligence: technology applied to conflict management within the brazilian judiciary, 2022.

[52] Cláudio José Palma Sanches and Sergio Fedato Batalha. Reflexos aplicativos entre a sumula vinculante n. 14 e a nova redação do artigo 155 do cpp. ETIC-ENCONTRO DE INICIAÇÃO CIENTÍFICA-ISSN 21-76-8498, 5(5), 2009.

[53] Ingo Wolfgang Sarlet and Jayme Weingartner Neto. Súmula vinculante nº11 (stf), imposição das algemas e dignidade da pessoa humana. Ius Gentium, 7(1):341–359, 2016.

[54] Rogerio Barros Sganzerla. O controle da justificativa do uso de algemas após a edição da súmula vinculante 11. Revista da Faculadade de Direito da UFMG, 61:473–498, 2012.

[55] Thomas A C Smith. The web of law. SSRN Electron. J., 2005.

[56] Renato Nonato Xavier Sobrinho. O direito à revisão geral anual dos servidores públicos em tempos de calamidade pública. Revista Controle-Doutrina e Artigos, 19(2):421–443, 2021.

[57] Lenio Luiz Streck. Súmulas no direito brasileiro: eficácia, poder e função. Livraria do Advogado, Porto Alegre, 1998.

[58] Keet Sugathadasa, Buddhi Ayesha, Nisansa de Silva, Amal Shehan Perera, Vindula Jayawardana, Dimuthu Lakmal, and Madhavi Perera. Legal document retrieval using document vector embeddings and deep learning. In Kohei Arai, Supriya Kapoor, and Rahul Bhatia, editors, Intelligent Computing, pages 160–175, Cham, 2019. Springer International Publishing.

[59] Zhongxiang Sun, Jun Xu, Xiao Zhang, Zhenhua Dong, and Ji-Rong Wen. Law article-enhanced legal case matching: A causal learning approach. In Proceedings

of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23, page 1549–1558, New York, NY, USA, 2023. Association for Computing Machinery.

[60] Fabien Tarissan and Raphaëlle Nollez-Goldbach. Temporal properties of legal decision networks: a case study from the International Criminal Court. In 28th International Conference on Legal Knowledge and Information Systems (JURIX'2015), Braga, Portugal, December 2015.

[61] Fabien Tarissan, Yannis Panagis, and Urška Šadl. Selecting the cases that defined europe: Complementary metrics for a network analysis. In 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pages 661–668, 2016.

[62] André Ramos Tavares. Tribunal e jurisdição constitucional. Instituto Brasileiro de Direito Constitucional, São Paulo, 1998.

[63] André Ramos Tavares. Teoria da Justiça Constitucional. Saraiva, São Paulo, 2005.

[64] André Ramos Tavares. Paradigmas do judicialismo constitucional. Saraiva, São Paulo, 2012.

[65] Howard Turtle. Text retrieval in the legal world. Artificial Intelligence and Law, 3:5–54, 1995.

[66] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is All You Need. In Proceedings of the 31st International Conference on Neural Information Processing Systems, pages 6000–6010, Long Beach, California, USA, 2017. Curran Associates Inc.

[67] Yen Thi-Hai Vuong, Quan Minh Bui, Ha-Thanh Nguyen, Thi-Thu-Trang Nguyen, Vu Tran, Xuan-Hieu Phan, Ken Satoh, and Le-Minh Nguyen. Sm-bert-cr: a deep learning approach for case law retrieval with supporting model. Artificial Intelligence and Law, pages 1–28, 2022.

[68] Ruili Wang and Yiming Zeng. Nonlinear nearest-neighbour matching and its application in legal precedent retrieval. In Third International Conference on Information Technology and Applications (ICITA'05), volume 1, pages 341–346 vol.1, 2005.

[69] Ryan Whalen. Modeling Annual Supreme Court Influence: The Role of Citation Practices and Judicial Tenure in Determining Precedent Network Growth, pages 169–176. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[70] Radboud Winkels, Jelle Ruyter, and Henryk Kroese. Determining authority of dutch case law. Legal knowledge and information systems, 235:103–112, 2011.

[71] Chaojun Xiao, Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Zhiyuan Liu, Maosong Sun, Tianyang Zhang, Xianpei Han, Zhen Hu, Heng Wang, and Jianfeng Xu. Cail2019-scm: A dataset of similar case matching in legal domain, 2019.

[72] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. arXiv preprint arXiv:1412.6575, 2014.

[73] Weijie Yu, Zhongxiang Sun, Jun Xu, Zhenhua Dong, Xu Chen, Hongteng Xu, and Ji-Rong Wen. Explainable legal case matching via inverse optimal transport-based rationale extraction, July 2022.

[74] Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. Big Bird: Transformers for longer sequences. CoRR, abs/2007.14062, 2020.

[75] Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. Ernie: Enhanced language representation with informative entities. arXiv preprint arXiv:1905.07129, 2019.

[76] Peng Zhou, Zhenyu Qi, Suncong Zheng, Jiaming Xu, Hongyun Bao, and Bo Xu. Text classification improved by integrating bidirectional lstm with two-dimensional max pooling. arXiv preprint arXiv:1611.06639, 2016.