Not only *what*, but also *when*: Understanding Brazilian political comments on legislative bills over time through Stance Detection and Topic Modeling

Matheus Cerqueira[a], *Nádia F. F. da* Silva[a,b], *Ellen* Souza[a,c], *Hidelberg O.* Albuquerque[c], *Márcio de S.* Dias[d], *André C. P. L. F. de* Carvalho[a]

[a]University of São Paulo, São Paulo, Brazil.
*matheuscerqueira@usp.br, andre@icmc.usp.br

[b]Federal University of Goiás, Goiás, Brazil.
nadia.felix@ufg.br

[c]Federal Rural University of Pernambuco, Pernambuco, Brazil.
{ellen.ramos,hidelberg.albuquerque}@ufrpe.br

[d]Federal University of Catalão, Goiás, Brazil.
marciodias@ufcat.edu.br

Submitted: 31 January 2025, Revised: 26 March 2025, Accepted: 21 April 2025, Published: 19 May 2025

Abstract. Legislative public spaces are important structures for participatory democracy, allowing citizens' voices to get engaged with politic decisions. As a consequence of the popularization of information and communication technologies, internet based tools have played an important role to improve public participation in political decisions, known as e-Democracy. These tools are usually composed of a set of functionalities or small services, named microservices. The better the microservices, the higher the citizen participation. This work investigates how to extract useful knowledge from citizen participation in the microservices of the public portal of the Brazilian Chamber of Deputies. For such, it analyzes public comments incorporating Natural Language Processing and Artificial Intelligence techniques in a platform named Ulysses. The tasks developed on this paper focus on a temporal analysis of comments on bills in the portal through Stance Detection and dynamic Topic Modeling tasks. For the first task, OxêSD, a BERTimbau-based model, was trained on two different corpora, one of them translated into Portuguese, and its predictive performance was evaluated using the F1 and ROC-AUC metrics, achieving 73% for both on our proposed Political-BRSD a mixed dataset containing both translated content from a bigger multilingual dataset (adapted from x-Stance) and bill-specific content (adapted from Ulysses-SD); for the second, BERTopic, a Topic Modeling framework, was used. Visualization tools to analyze how the proposed approach addressed the task were also used to explore the knowledge extracted. They allow the user to understand over time how the comments relate to each other and how the comments relate to a given legislative bill.

Keywords. Natural Language Processing, Legal Text Analysis, Stance Detection, Topic Modeling.
Research paper, DOI: https://doi.org/10.59490/dgo.2025.955

1. Introduction

Public portals, defined as inclusive and deliberative instances not monopolized by any entity, play a crucial role in democracies (Addor, 2018). As societies strengthen their relationships with the state, citizens evolve from passive to active participants, engaging in public discourse and decision-making processes. This evolution prompts a continuous redefinition of public spaces and democracy itself (Cardoso, 2007). Participatory democracy, often viewed as a combination of direct and indirect democracy, has undergone a contemporary reevaluation. It now encompasses universal participation, leveraging various tools to expand public spaces in political and public decision making (Macedo, 2008). However, universal participation raises concerns about potentially suppressing the voices of minorities, necessitating careful attention to the expression of diverse social groups (Lessa, 2010).

The advent of the internet and digital technologies has transformed interactions among citizens, politicians, and governmental institutions. This transformation has given rise to the concept of e-Democracy, wherein Information and Communications Technologies (ICTs) support democratic decision-making, enhancing both institutions and processes (Pereira et al., 2022). Early studies emphasized the need to analyze the quality of information and interactions on the internet, focusing on their relevance, scope, and consequences (Rothberg, 2008). This paper focus on e-Democracy, exploring how ICTs can improve democratic processes. E-democracy operates on three dimensions: establishing essential conditions for the use of information and communication technology (ICT) in government, emphasizing citizen participation (e-participation) in decision-making, and concentrating on the effective delivery of government services online (e-Government) (Kneuer, 2016; Ronchi, 2019).

Natural Language Processing (NLP), a subfield of Artificial Intelligence, has evolved in recent years. Fueled by increased computing power, abundant data, and advances in Machine Learning, NLP encompasses various tasks, including Word Processing, Information Retrieval, and Sentiment Analysis (Chowdhary, 2020; Hirschberg & Manning, 2015). Stance Detection, a subtask of sentiment analysis, involves categorizing the expressed position of a given text regarding a particular target. This classification task finds applications in diverse areas such as rumor stance classification and fake news detection (Kuçuk & Can, 2020). Another important NLP task is Topic Modeling, an unsupervised technique for extracting latent topics from textual data. It plays a significant role in uncovering the underlying themes within complex and large corpora across various domains (George & Birla, 2018; Yamunathangam et al., 2021). Therefore, considering the temporal dynamics of language use, Temporal Text Analysis allows understanding how discourses change over time. In the context of media storms, where a specific subject experiences a surge in news coverage within a short time on social media, temporal text analysis becomes relevant (Lukito et al., 2019).

Legal AI and Legal NLP have emerged as areas focused on augmenting, assisting, or potentially replacing legal actors in various tasks. These tasks include the analysis of legal documents, like laws, bills, legislation texts, case laws, and more, often scattered across different databases and formats (Avgerinos Loutsaris et al., 2021; Katz et al., 2023). In this context, the Brazilian Chamber of Deputies uses multiple NLP and AI tools in a platform named Ulysses Câmara dos Deputados, 2018, which processes legislative texts from its portal named *E-democracia*[1]. The goal of this research is to apply NLP methods in the temporal domain to analyze comments on bills written on the E-Democracia portal through the Ulysses platform in order to understand how public opinion towards a given subject evolves through the time. In this work, we (1) provide a deep and thorough review of textual analysis within the legal domain over time, emphasizing Stance Detection and Topic Modeling tasks; (2) present significant advances in the PT-BR Stance Detection task within the legal domain by introducing a new large corpus (Political-BRSD[2]) and a novel Stance Detection model (OxêSD[3]); and (3) through Stance Detection and Topic Modeling over time, develop tools that enable an understanding of how texts engage with a specific bill and with each other over time.

2. Related Work

In this section, we introduce related works in NLP, specifically in Stance Detection and Topic Modeling, considering Legal Domain and Temporal applications.

2.1. Legal Domain Applications

Avgerinos Loutsaris et al., 2021 address the challenge posed by the enormous volume of legal information in each European Union country due to the integration of Digital Governance across all levels of society. They

[1] Available in https://www.camara.leg.br/participe
[2] Available in https://huggingface.co/datasets/cerqueiramatheus/Political-BRSD
[3] Available in https://huggingface.co/cerqueiramatheus/OxeSD

have developed an open and automated legal system that handles various tasks, including information extraction (such as law title, number, number of pages, and creation date), Named Entity Recognition (NER), law decomposition, and opinion extraction. On a different note, Ikram and Chakir, 2019 highlight the difficulty of archiving legal documents and accessing relevant information in Morocco. To assist legal professionals, they propose text classification methods employing multiple Machine Learning (ML) induced models. In the context of the Portuguese language, there are also several studies that deal with different legal text analysis problems. Carvalho and Barbosa, 2018 implement a legal visualization based software using tokenization, tagging, and lemmatization to enhance the understanding and analysis of documents such as the Portuguese Constitution. The Lima et al., 2021 study focuses on comparing different feature extraction models in Brazilian Portuguese texts from the Tribunal de Justiça do Rio Grande do Norte (TJRN), a clustering technique that can be useful for lawyers and even to finding controversial issues in legal texts. Also in Brazilian Portuguese, Medina et al., 2022 propose a model with different tasks, such as summarization, feature extraction, and classification, focused on legal proceedings texts.

Nababan et al., 2022 study is focused on understanding public positions and sentiments about the Job Creation Bill in Indonesia, through a Twitter data analysis; they collected and annotated several tweets to create a new corpus and also conducted experiments with different ML algorithms, such as Naïve Bayes, Support Vector Machine, and Logistic Regression. These algorithms were chosen because of their different inductive biases. Bergam et al., 2022 work delves into the analysis of US Supreme Court documents through Stance Detection, introducing and comparing new metrics that allowed them to correlate the responsiveness to public opinion to ideology expression. They also proposed *SC-stance*, a new dataset focused on legal Stance Detection, matching written opinions to legal questions. Usually, studies adopt a multilingual approach. Lai, Cignarella, et al., 2020 work select topics about politics in multiple languages, investigate the portability of tools, and then evaluate MultiTACOS algorithm performance. Egli et al., 2023 proposal is centered on the analysis of official voter information booklets; their experiments achieved best results using a fine-tuned multilingual BERT model Devlin et al., 2019. Finally, within the Brazilian Chamber of Deputies context, Maia et al., 2022 use digital media comments on political bills to understand the stances of the citizens and group the polarity of public opinions about new discussion topics. To achieve this, they developed a new dataset, *Ulysses-SD* and implemented a mBERT model. It is the nearest application in domain and approach compared to our goal; however, none of them leverages the temporal information.

Sarne et al., 2019 work is focused in large-scale corpora of privacy policies using unsupervised learning techniques. They developed a framework for privacy policy Topic Modeling and applied it over a corpus of privacy policies of mobile applications crawled from the Google Play Store. To build their framework, they used GraphLab, a machine learning platform that supports Latent Dirichlet allocation (LDA) Blei et al., 2003 Topic Modeling. Similarly, LawNet-Viz is a web-based system developed by La Cava et al., 2022 for modeling, analyzing and visualizing law reference networks extracted from law corpus. Their work was demonstrated using the Italian Civil Code and implemented the LDA topic model Blei et al., 2003 available on the Gensim Python library for the language modeling process. LawNet-Viz can benefit since lawyers and jurists to citizens. Joshi and Saha, 2020 developed a framework to automatically extract knowledge from Code of Federal Regulations from the United States government, such as key terms, rules, relationships and topic summaries. They used the Gensim library and the LDA Blei et al., 2003 topic model for the extraction of topic from summarized subsections task. Their work automated the analysis from CFR regulations. For the Brazilian Portuguese context, Silva et al., 2021 designed a specific sentence embedding technique to build topic models and achieved good results compared to other multilingual sentence embedding techniques. Their approach applied BERTopic Grootendorst, 2022 to extract topics from political comments on Brazilian bills and produced three different datasets. Again, on the same domain, the temporal information was not explored, one of our main goals in this work.

2.2. Temporal Applications

Usually, the temporal applications are focused in time information extraction. Viani et al., 2020 developed a system to identify relevant electronic health records and extract time expressions from them; this system was applied to two electronic health records datasets and resulted in the creation of a new corpus annotated with normalized values for time expressions. However, within the context of this work, our main task is to understand how texts can be grouped over time, that is, the analysis of documents in document networks Zhang and Lauw, 2022, specifically with Stance Detection and Topic Modeling. Although there are multiple applications using temporal information, our work differs by combining two tasks, Stance Detection and Topic Modeling, for understanding texts in the legal domain.

Alturayeif et al., 2023 review highlighted that, despite that opinions change through time and the importance to record and evaluate temporal data, diachronic evolution is not a common explored topic within Stance Detection field; it also presented three different studies. The first, from Lai et al., 2019, introduces how political debate can be understood in a diachronic perspective, focusing on the aggregation of empirical data in a given granularity (days, weeks, months). Their main goal was the analysis of the Twitter debate about the 2016 referendum of Italian Constitution; they developed a new corpus leveraging the diachronic perspective to understand the citizens' opinion evolution through time. The second, proposed by Lai, Patti, et al., 2020, introduces a novel approach and annotation schema for Stance Detection to investigate the impact of social network community features and the evolution of stance over time within the Brexit referendum context. Their analysis highlighted the importance of temporal features. Finally, Alkhalifa et al., 2021 work produced two large-scale stance datasets using temporal information to mitigate the performance deterioration of classifiers over time, that is, to reduce impacts of the temporal gap between training and test data. Their approach adapted word embeddings used in training of the stance classifier.

One of the earliest approaches to dealing with the Topic Modeling task over time was proposed by Blei and Lafferty, 2006. Their work developed the Dynamic Topic Models, that employed probabilistic time series models has been developed to analyze the temporal evolution of topics within extensive document collections. Since then, several different proposals have been developed. He et al., 2014 emphasize the increasing interest in this area and mention, beyond the Dynamic Topic Models, its newer variations, like continuous-time and online multi-scale proposals. Their work also proposes a new method joining sentiment and topic model over time. Sharma et al., 2022 applied dynamic Topic Modeling on vaccine-related discussions on Twitter from 2020 to 2021, aiming to track public perception of vaccines over time and identify the factors influencing these evolving opinions. Ghosh et al., 2017 study explores the use of supervised temporal topic models to analyze large corpora of news articles and identify trends; they evaluated their method through multiple infectious disease outbreaks data from the United States of America, China and India. Finally, Grootendorst, 2022 study proposes BERTopic, a Topic Modeling framework that implements dynamic Topic Modeling.

3. Method

In this section, we describe our method, focusing on the steps, input and output of our task, the models used, and the evaluation metrics. Briefly, our approach follows the presented on Figure 1.

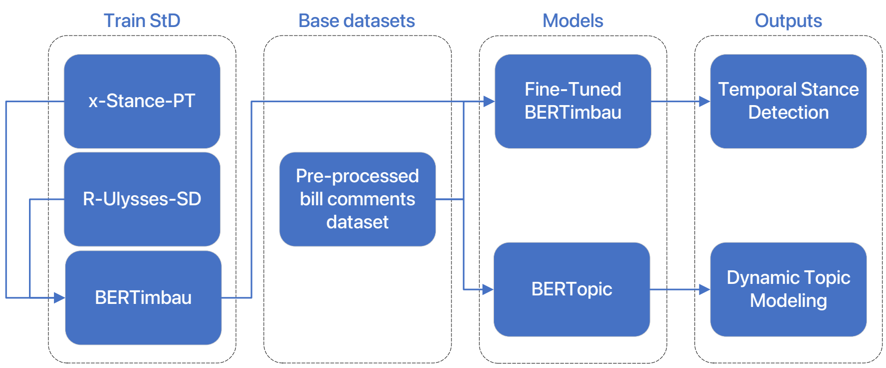

Fig. 1 – Methodological approach

3.1. Comment Analysis

Our main task is focused on the temporal analysis of Brazilian political comments through Stance Detection and Topic Modeling. First, we fine-tuned a model for the Stance Detection task using two different datasets (train datasets). For the main task, our input is composed by comments written about a given bill on the E-Democracia portal (base datasets). We employ a pre-processing step for the base datasets. Then, we execute models for Stance Detection and Topic Modeling given the comments, what will be discussed in the next section. Finally, we provide graphs containing the output of each task (polarities over time and topics) for the given comments set.

3.2. Models

First, for the Stance Detection task, we fine-tuned a pre-trained BERT model, specifically BERTimbau. BERT models employ a bidirectional context approach, creating representations of words considering their contexts within a given sentence. BERTimbau, specialized for Brazilian Portuguese, outperforms generic BERT on Portuguese tasks. For the temporal Topic Modeling task, we utilized BERTopic (Grootendorst, 2022), a model leveraging class-based TF-IDF with pre-trained S-BERT models (Reimers & Gurevych, 2019) containing multiple Topic Modeling approaches, including Dynamic Topic Modeling.

3.3. Evaluation Metrics

To assess the performance of our models, we used the Hugging Face Evaluate library (HuggingFace, 2023), leveraging both the weighted F1 Score and Receiver Operating Characteristic Area Under the Curve (ROC AUC) metrics. The weighted F1 score is the harmonic mean of precision and recall, computed for each class independently and then averaged, with the contribution of each class weighted by its prevalence in the dataset. Simultaneously, the ROC AUC metric was employed to evaluate the model ability to discriminate between positive and negative instances, offering insights into their overall predictive power.

4. Experimental Setup

In this section, we describe our base dataset used as input of the tasks, the datasets used to training the Stance Detection Model, its setup and implementation.

4.1. Base Datasets

To generate the outputs for the Stance Detection and Topic Modeling tasks, we extract comments associated with specific bills from the E-democracia portal, each producing a unique dataset. Similarly to the works of (Maia et al., 2022; Silva et al., 2021), our extracted datasets are composed by comments written in Brazilian Portuguese from CSV files, containing columns such as comment ID, author, timestamp, number of likes, and manual labels. Finally, we preprocess the comments, filtering for a minimum length (default set at 15 characters) in order to keep only real comments and produce better results.

4.2. Train Datasets

Given the lack of large datasets for Brazilian Portuguese Stance Detection, we adopted a combination of the x-Stance dataset (Vamvas & Sennrich, 2020) and the Ulysses-SD dataset (Maia et al., 2022). To address language disparities, we translated the x-Stance dataset to Portuguese using the Google Translator class from the Deep Translator package (Baccouri, 2023), resulting in the x-Stance-PT dataset[4]. We also reduced the Ulysses-SD to two classes (`FAVOR` and `AGAINST`) in R-Ulysses-SD[5]. The Political-BRSD dataset, a combination of x-Stance-PT and R-Ulysses-SD, was created to join information from both datasets, as illustrated in Table 1. Each of them followed a 70/15/15 division for train, test and validation.

4.3. Setup and Implementation

Our Stance Detection experiments consisted in four approaches trained on a NVIDIA Tesla P100 GPU with 16GB memory, employing the same hyperparameters (`learning_rate` = 2^{-5} and `epochs` = 4). The approaches

[4]Available in https://huggingface.co/datasets/cerqueiramatheus/x-Stance-PT
[5]Available in https://huggingface.co/datasets/cerqueiramatheus/reduced-Ulysses-SD

Tab. 1 – **Number of comments for the train datasets** – Number of comments in each label for the R-Ulysses-SD, x-Stance-PT and Political-BRSD datasets.

Dataset	Favor	Against	Total
R-Ulysses-SD	524	685	1209
x-Stance-PT	34001	33270	67271
Political-BRSD	34525	33955	68480

included models fine-tuned only on Ulysses-SD, only on x-Stance-PT, sequentially on x-Stance-PT and Ulysses-SD, and finally on the merged Political-BRSD dataset. Each model received as input a combination of topic information, the comment itself from each dataset and a given label (FAVOR or AGAINST). For temporal Topic Modeling, we used Dynamic Topic Modeling within the existing BERTopic implementation on Ulysses services, which automatically selects 10 topics, automatically named.

5. Results

To exemplify our results, we selected two different bills from the E-democracy portal to show the outputs of each task. We selected the following bills because they have a high number of comments and are recent, both created in 2020: 1) Public Administration – **PL 2630/2020** (Câmara dos Deputados, 2020) (to combat fake news); and Health – **PL 2564/2020** (Câmara dos Deputados, 2021) (to establish salary floor for nursing professionals).

5.1. Stance Detection Model Training

Tab. 2 – **F1 and AUC measures** – F1 and AUC measure results for the proposed models 1, 2, 3 and 4 on x-Stance-PT, R-Ulysses-SD and Political-BRSD datasets.

Dataset	Model							
	1		2		3		4 (best)	
	F1	AUC	F1	AUC	F1	AUC	F1	AUC
x-Stance-PT	0.40	0.51	0.71	0.71	0.58	0.61	**0.73**	**0.72**
R-Ulysses-SD	0.86	0.85	0.54	0.56	**0.87**	**0.86**	0.85	0.84
Political-BRSD	0.58	0.61	0.71	0.71	0.59	0.62	**0.73**	**0.73**

We follow the model names from the previous section: 1) BERTimbau trained on R-Ulysses-SD; 2) BERTimbau trained on x-Stance-PT; 3) BERTimbau trained on R-Ulysses-SD and then on x-Stance-PT; and 4) BERTimbau trained on Political-BRSD. As Table 2 shows, 1 and 2, as expected, produced the worst outputs; they specialized only on the context of the dataset that they were trained. Model 3 appears to be forgetting its first training on x-Stance-PT, but this is a known issue since the recommended number of epochs for this task, at maximum, 4 (Devlin et al., 2019) and we performed two training steps of 4 epochs each. Finally, model 4 produces the best results, leveraging both datasets.

Our best model, BERTimbau trained on Political-BRSD, outperforms results from the original x-Stance work (F1=0.69). To the best of our knowledge, our produced results are the best for cross-topic Stance Detection in the PT-BR political domain. Then, we evaluated its confusion matrix produced for the test datasets. For R-Ulysses-SD, its performance is better at class AGAINST, with the same number of false negatives and false positives. However, for x-Stance-PT and Political-BRSD, the model produces better results for the FAVOR class; the number of false positives and false negatives is still similar. This indicates that our model learned to better classify the FAVOR stances, even with the use of well-balanced data sets, as previously shown.

5.2. Tasks Outputs

Usually, the number of comments on the bills is not constant, and there are months with much more interaction and this is due to external factors, like news or social networks; some bills are also more controversial than others. However, these factors are not explored in this work, as we propose a first step to this analysis, focused on identifying the reactions and interactions within the comments of a given bill.

5.2.1. Stance Detection

As shown in Figure 2, there is a prevalence of the `FAVOR` class for the PL 2630/2020, increasing from December 2021 to January 2022, what indicates that the bill is probably less controversial; its purpose (salary floor for nursing professional) can be a reason to it. However, PL 2564/2020 (Figure 3) seems to be more controversial, with small peaks before 2023 and two huge peaks between April 2023 to June 2023. Despite its majority class is `AGAINST`, there is a significant number of `FAVOR` comments.

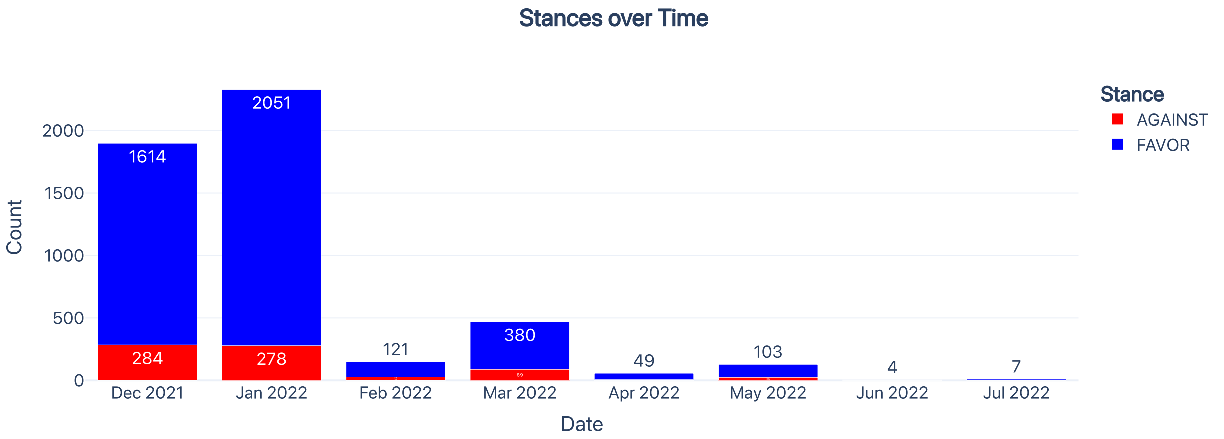

Fig. 2 – Stance Detection output for PL 2630/2020

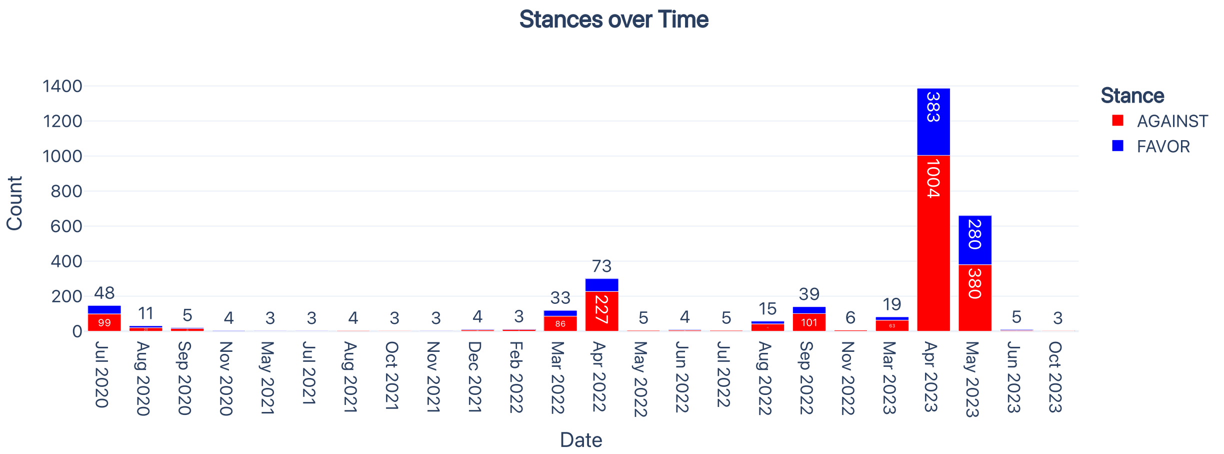

Fig. 3 – Stance Detection output for PL 2564/2020

In our experiments with multiple bills, there was a prevalence of comments distribution over time on bills with this same behavior, that is, concentrated in two or three months, with isolated peaks over time.

5.2.2. Topic Modeling

Figure 4 shows the comment groups for PL 2630/2020, matching the frequency peaks from the Stance Detection task. Topic 0 (PT-BR: `paciente_respeito_aprovar_trabalho`, EN-US: `patient_respect_approve_work`) groups most of the comments. For PL 2564/2020, as PL 2564/2020, there is again a concentration in a single topic, with other minor topics. Topic 0 (PT-BR: `rede_informação_fake_internet`, EN-US: `network_information_fake_internet`), as shown in Figure 5, groups most of the comments over time.

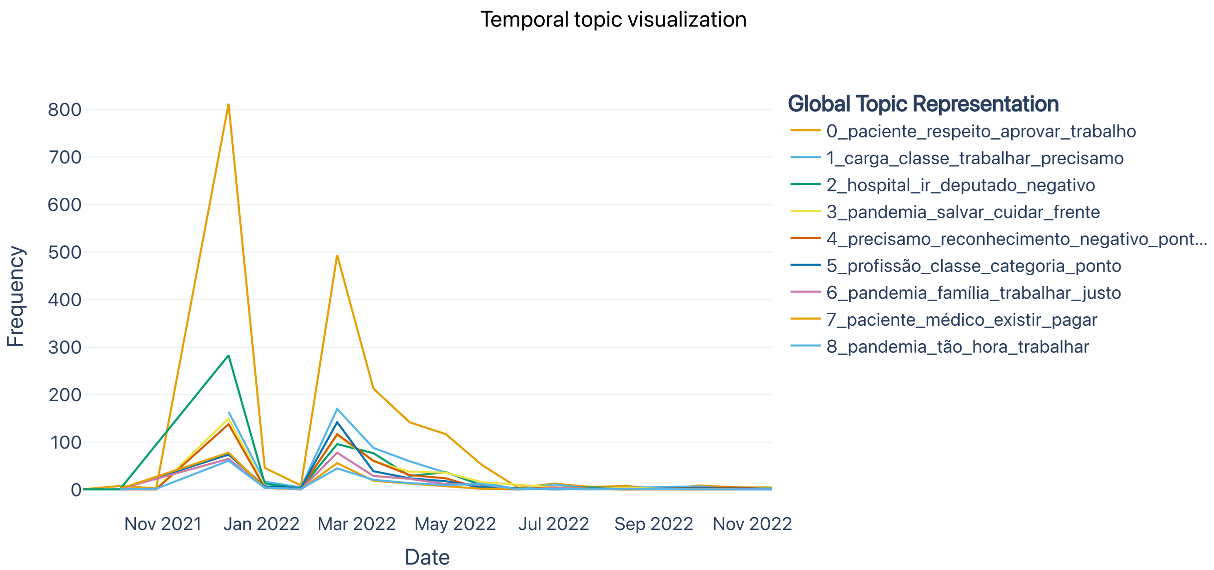

Fig. 4 – Topic Modeling output for PL 2630/2020

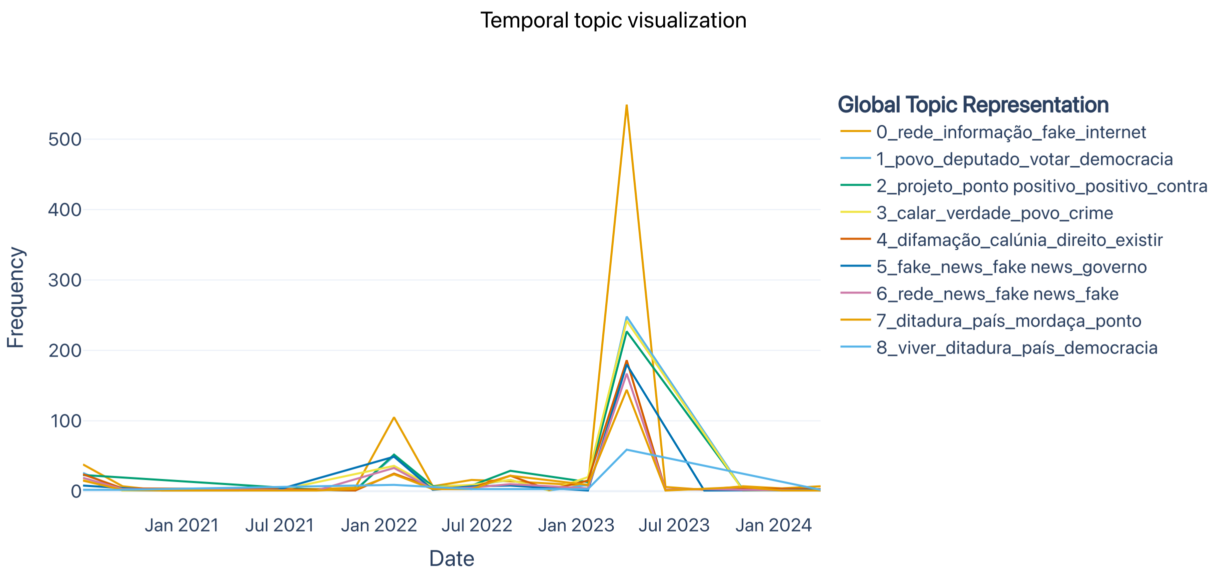

Fig. 5 – Topic Modeling output for PL 2564/2020

With multiple bills, there was always a main topic, as in the previous examples. As all the comments are about the same bill, this is expected behavior. However, the other topics seem to be relevant and can reveal important insights.

5.3. Discussion

Even though the Stance Detection and Topic Modeling tasks have been designed to be globally available, that is, to work on all available bills, their performance depends on the comments distribution over time (as shown in the first example), the total amount of available comments (usually bills take months to have a few comments), external factors (as the influence of external discussion and the bill relevance). Our selected examples were manually picked, since the Brazilian Chamber of Deputies services do not provide filters to select the most commented bills in a complete period, only in the last six months. In the examples, the Stance Detection visualization provides good insights about citizens' responses to bills; although the Topic Modeling focuses on one topic, it groups other relevant secondary topics, which can be better seen through dynamic filtering in the graphs, which are interactive.

5.4. Limitations

Due to the lack of well-established evaluation metrics for dynamic topic models, we were unable to quantitatively assess the performance of our proposed model. However, we acknowledge the recent work of (James et al., 2023), that proposed a novel evaluation measure for dynamic topic models that analyzes the changes in the quality of each topic over time; they propose Temporal Topic Coherence, Temporal Topic Smoothness, Temporal Topic Quality and Dynamic Topic Quality. It is important to note that the non-temporal evaluation, that is, the traditional Topic Modeling evaluation, is already developed by other projects within Ulysses services, which is beyond the scope of this work.

Although significant progress has been made in developing cross-topic Stance Detection models, the evaluation of their performance remains a challenge. As highlighted in (Reuver et al., 2021), the generalizability of cross-topic Stance Detection models across diverse topics and socio-cultural contexts is a critical aspect that requires further investigation and demonstrates that topic, class, and their interaction can significantly impact model performance. This underscores the need for continuous evaluation of cross-topic Stance Detection models across a wide range of contexts and topics over different datasets to ensure their robustness and generalizability. Our research was limited to the context of Political-BRSD.

6. Conclusions

Our best-performing Stance Detection model (OxêSD, a BERTimbau-based model fine-tuned on Political-BRSD) surpassed results from existing works, demonstrating the effectiveness of our approach in cross-topic scenarios. The dynamic Topic Modeling framework provided a detailed temporal analysis, contributing to a deeper understanding of the evolving discourse around legislative proposals. Despite promising results, our research has limitations. The lack of well-established evaluation metrics for dynamic topic models poses challenges in quantitatively assessing their performance. Future work could focus on refining evaluation measures for temporal analysis tasks. Furthermore, expanding the scope to include more diverse topics and datasets would enhance the generalizability of our models. In this work, we addressed the challenges of cross-topic Stance Detection and dynamic Topic Modeling in the context of Brazilian political discussions. Our experiments with Stance Detection models revealed that training on a combination of datasets (Political-BRSD) significantly improves performance compared to models trained on individual datasets (x-Stance-PT and R-Ulysses-SD). The best performing model (trained on Political-BRSD) achieved superior F1 and ROC AUC scores, showcasing its effectiveness in capturing cross-topic stances in the PT-BR political domain.

The temporal analysis of stance over time provided valuable insights into citizens' responses to specific bills. By visualizing stances and topics, we observed patterns of interaction and identified potentially controversial bills. Our dynamic Topic Modeling framework successfully captured the evolution of topics over time, offering a nuanced understanding of the discussions surrounding legislative proposals. In conclusion, our research contributes to advancing the field of cross-topic Stance Detection and temporal analysis in the context of political discussions. The combination of effective models, dynamic visualizations, and an interactive demo platform positions our work as a valuable resource for researchers, policymakers, and the public interested in understanding and analyzing political discourse.

7. Acknowledgements

We sincerely appreciate the support provided by the research funding agencies CAPES, FAPESP, and CNPq, whose contributions have been essential to the development of this work. Furthermore, we extend our gratitude to the Brazilian Chamber of Deputies, the Center for Applied Research in Artificial Intelligence Recreating Environments (CPA IARA), and the Brazilian National Institute of Artificial Intelligence (IAIA) for their significant support in advancing this research.

References

Addor, F. (2018). Reflections on local participatory democracy in latin america. *Revista de Administração Pública*, *52*(6), 1108–1124. DOI: https://doi.org/10.1590/0034-761220170131.

Alkhalifa, R., Kochkina, E., & Zubiaga, A. (2021). Opinions are made to be changed: Temporally adaptive stance classification. *Proceedings of the 2021 Workshop on Open Challenges in Online Social Networks*, 27–32. DOI: https://doi.org/10.1145/3472720.3483620.

Alturayeif, N., Luqman, H., & Ahmed, M. (2023). A systematic review of machine learning techniques for stance detection and its applications. *Neural Computing and Applications*, *35*(7), 5113–5144. DOI: https://doi.org/10.1007/s00521-023-08285-7.

Avgerinos Loutsaris, M., Lachana, Z., Alexopoulos, C., & Charalabidis, Y. (2021). Legal text processing: Combing two legal ontological approaches through text mining. *DG.O2021: The 22nd Annual International Conference on Digital Government Research*, 522–532. DOI: https://doi.org/10.1145/3463677.3463730.

Baccouri, N. (2023). Deep translator.

Bergam, N., Allaway, E., & Mckeown, K. (2022). Legal and political stance detection of SCOTUS language. *Proceedings of the Natural Legal Language Processing Workshop 2022*, 265–275. DOI: https://doi.org/10.18653/v1/2022.nllp-1.25.

Blei, D. M., & Lafferty, J. D. (2006). Dynamic topic models. *Proceedings of the 23rd International Conference on Machine Learning*, 113–120. DOI: https://doi.org/10.1145/1143844.1143859.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, *3*(Jan), 993–1022.

Câmara dos Deputados. (2018, November). Câmara lança ulysses, robô digital que articula dados legislativos. *Portal da Câmara dos Deputados*. Retrieved November 27, 2023, from https://www.camara.leg.br/noticias/548730-camara-lanca-ulysses-robo-digital-que-articula-dados-legislativos/o

Câmara dos Deputados. (2020, July). Enquete do PL 2630/2020. *Portal da Câmara dos Deputados*. Retrieved November 27, 2023, from https://www.camara.leg.br/enquetes/2256735/resultados

Câmara dos Deputados. (2021, December). Enquete do PL 2564/2020. *Portal da Câmara dos Deputados*. Retrieved November 27, 2023, from https://www.camara.leg.br/enquetes/2309349/resultados

Cardoso, F. H. (2007). A democracia no centro da agenda. In B. Sorj & M. D. d. Oliveira (Eds.), *Sociedade civil e democracia na américa latina: Crise e reinvenção da política* (pp. 7–10). SciELO Books - Centro Edelstein.

Carvalho, N. R., & Barbosa, L. S. (2018). Transforming legal documents for visualization and analysis. *Proceedings of the 11th International Conference on Theory and Practice of Electronic Governance*, 23–26. DOI: https://doi.org/10.1145/3209415.3209424.

Chowdhary, K. R. (2020). Natural language processing. In *Fundamentals of artificial intelligence* (pp. 603–649). Springer India. DOI: https://doi.org/10.1007/978-81-322-3972-7_19.

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. DOI: https://doi.org/10.18653/v1/N19-1423.

Egli, E., Mamie, N., Dolev, E. L., & Muller, M. (2023). Voting booklet bias: Stance detection in swiss federal communication. *arXiv preprint arXiv:2306.08999*. https://arxiv.org/abs/2306.08999

George, L. E., & Birla, L. (2018). A study of topic modeling methods. *2018 Second International Conference on Intelligent Computing and Control Systems (ICICCS)*, 109–113. DOI: https://doi.org/10.1109/ICCONS.2018.8663152.

Ghosh, S., Chakraborty, P., Nsoesie, E. O., Cohn, E., Mekaru, S. R., Brownstein, J. S., & Ramakrishnan, N. (2017). Temporal topic modeling to assess associations between news trends and infectious disease outbreaks. *Scientific Reports*, *7*(1), 40841. DOI: https://doi.org/10.1038/srep40841.

Grootendorst, M. (2022). Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*. https://arxiv.org/abs/2203.05794

He, Y., Lin, C., Gao, W., & Wong, K.-F. (2014). Dynamic joint sentiment-topic model. *ACM Trans. Intell. Syst. Technol.*, *5*(1). DOI: https://doi.org/10.1145/2542182.2542188.

Hirschberg, J., & Manning, C. D. (2015). Advances in natural language processing. *Science*, *349*(6245), 261–266. DOI: https://doi.org/10.1126/science.aaa8685.

HuggingFace. (2023). Evaluate. https://github.com/huggingface/evaluate

Ikram, A. Y., & Chakir, L. (2019). Arabic text classification in the legal domain. *2019 Third International Conference on Intelligent Computing in Data Sciences (ICDS)*, 1–6. DOI: https://doi.org/10.1109/ICDS47004.2019.8942343.

James, C., Nagda, M., Haji Ghassemi, N., Kloft, M., & Fellenz, S. (2023, September). *Evaluating dynamic topic models*. DOI: https://doi.org/10.48550/arXiv.2309.08627.

Joshi, K. P., & Saha, S. (2020). A semantically rich framework for knowledge representation of code of federal regulations. *Digit. Gov.: Res. Pract.*, *1*(3). DOI: https://doi.org/10.1145/3425192.

Katz, D. M., Hartung, D., Gerlach, L., Jana, A., & Bommarito, I., Michael J. (2023, February). *Natural language processing in the legal domain* [Provided by the SAO/NASA Astrophysics Data System]. DOI: https://doi.org/10.48550/arXiv.2302.12039.

Kneuer, M. (2016). E-democracy: A new challenge for measuring democracy. *International Political Science Review / Revue internationale de science politique*, *37*(5), 666–678. Retrieved September 18, 2023, from https://www.jstor.org/stable/26556880

Kuçuk, D., & Can, F. (2020). Stance detection: A survey. *ACM Comput. Surv.*, *53*(1). DOI: https://doi.org/10.1145/3369026.

La Cava, L., Simeri, A., & Tagarelli, A. (2022). Lawnet-viz: A web-based system to visually explore networks of law article references. *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 3300–3305. DOI: https://doi.org/10.1145/3477495.3531668.

Lai, M., Cignarella, A. T., Hernández Farías, D. I., Bosco, C., Patti, V., & Rosso, P. (2020). Multilingual stance detection in social media political debates. *Computer Speech & Language*, *63*, 101075. DOI: https://doi.org/https://doi.org/10.1016/j.csl.2020.101075.

Lai, M., Patti, V., Ruffo, G., & Rosso, P. (2020). #brexit: Leave or remain? the role of user's community and diachronic evolution on stance detection (D. Pinto, V. Singh, & F. Perez, Eds.). *Journal of Intelligent &amp Fuzzy Systems*, *39*(2), 2341–2352. DOI: https://doi.org/10.3233/jifs-179895.

Lai, M., Tambuscio, M., Patti, V., Ruffo, G., & Rosso, P. (2019). Stance polarity in political debates: A diachronic perspective of network homophily and conversations on twitter. *Data & Knowledge Engineering*, *124*, 101738. DOI: https://doi.org/https://doi.org/10.1016/j.datak.2019.101738.

Lessa, R. (2010). Democracia, representação e desenvolvimento. In F. d. S. Silva, F. G. Lopez, R. R. C. Pires, et al. (Eds.), *Estado, instituições e democracia: Democracia* (pp. 47–89, Vol. 2). Instituto de Pesquisa Econômica Aplicada (Ipea).

Lima, J. P., Alfredo Costa, J., & Araujo, D. C. (2021). Comparison of feature extraction methods for brazilian legal documents clustering. *2021 IEEE Latin American Conference on Computational Intelligence (LA-CCI)*, 1–5. DOI: https://doi.org/10.1109/LA-CCI48322.2021.9769839.

Lukito, J., K Sarma, P., Foley, J., & Abhishek, A. (2019). Using time series and natural language processing to identify viral moments in the 2016 U.S. presidential debate. *Proceedings of the Third Workshop on Natural Language Processing and Computational Social Science*, 54–64. DOI: https://doi.org/10.18653/v1/W19-2107.

Macedo, P. S. N. d. (2008). Democracia participativa na constituição brasileira. *Revista de informação legislativa*, *45*(178), 181–193. http://www2.senado.leg.br/bdsf/handle/id/176529

Maia, D. F., Silva, N. F. F., Souza, E. P. R., Nunes, A. S., Procópio, L. C., Sampaio, G. d. S., Dias, M. d. S., Alves, A. O., Maia, D. F., Ribeiro, I. A., Pereira, F. S. F., & de Carvalho, A. P. d. L. F. (2022). Ulyssessd-br: Stance detection in brazilian political polls. In G. Marreiros, B. Martins, A. Paiva, B. Ribeiro, & A. Sardinha (Eds.), *Progress in artificial intelligence* (pp. 85–95). Springer International Publishing.

Medina, M. C. C., Da Silva Oliveira, L. M., Ferreira, J. F. C., Silva, L. H. D. S., Rodrigues, C. M. O., De Oliveira, J. F. L., Sobral, P. C., Souza, B., Feitosa, D., & Fernandes, B. J. T. (2022). Classification of legal documents in portuguese language based on summarization. *2022 IEEE Latin American Conference on Computational Intelligence (LA-CCI)*, 1–6. DOI: https://doi.org/10.1109/LA-CCI54402.2022.9981852.

Nababan, A. H., Mahendra, R., & Budi, I. (2022). Twitter stance detection towards job creation bill [Sixth Information Systems International Conference (ISICO 2021)]. *Procedia Computer Science*, *197*, 76–81. DOI: https://doi.org/https://doi.org/10.1016/j.procs.2021.12.120.

Pereira, M. A., Bernardes, C. B., & Vale, M. L. (2022). O papel da moderação nas audiências públicas interativas do e-democracia: Entre a promoção do debate e o receio da censura. *Revista Brasileira de Ciência Política*, (37), e249679. DOI: https://doi.org/10.1590/0103-3352.2022.37.249679.

Reimers, N., & Gurevych, I. (2019, November). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In K. Inui, J. Jiang, V. Ng, & X. Wan (Eds.), *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (emnlp-ijcnlp)* (pp. 3982–3992). Association for Computational Linguistics. DOI: https://doi.org/10.18653/v1/D19-1410.

Reuver, M., Verberne, S., Morante, R., & Fokkens, A. (2021, November). Is stance detection topic-independent and cross-topic generalizable? - a reproduction study. In K. Al-Khatib, Y. Hou, & M. Stede (Eds.), *Pro-*

ceedings of the 8th workshop on argument mining (pp. 46–56). Association for Computational Linguistics. DOI: https://doi.org/10.18653/v1/2021.argmining-1.5.
Ronchi, A. M. (2019, February). *E-democracy* (1st ed.). Springer International Publishing.
Rothberg, D. (2008). Por uma agenda de pesquisa em democracia eletrônica. *Opinião Pública*, *14*(1), 149–172. DOI: https://doi.org/10.1590/S0104-62762008000100006.
Sarne, D., Schler, J., Singer, A., Sela, A., & Bar Siman Tov, I. (2019). Unsupervised topic extraction from privacy policies. *Companion Proceedings of The 2019 World Wide Web Conference*, 563–568. DOI: https://doi.org/10.1145/3308560.3317585.
Sharma, A., Susan, S., Bansal, A., & Choudhry, A. (2022). Dynamic topic modeling of covid-19 vaccine-related tweets. *2022 the 5th International Conference on Data Storage and Data Engineering*, 79–84. DOI: https://doi.org/10.1145/3528114.3528127.
Silva, N., Silva, M., Pereira, F., Tarrega, J., Beinotti, J., Fonseca, M., Andrade, F., & Carvalho, A. (2021). Evaluating topic models in portuguese political comments about bills from brazil's chamber of deputies. *Anais da X Brazilian Conference on Intelligent Systems*. https://sol.sbc.org.br/index.php/bracis/article/view/19061
Vamvas, J., & Sennrich, R. (2020). X-stance: A multilingual multi-target dataset for stance detection [5th SwissText & 16th KONVENS Joint Conference 2020, SwissText and KONVENS 2020 ; Conference date: 23-06-2020 Through 25-06-2020]. *Proceedings of the 5th Swiss Text Analytics Conference (SwissText) & 16th Conference on Natural Language Processing (KONVENS)*. https://swisstext-and-konvens-2020.org/
Viani, N., Kam, J., Yin, L., Bittar, A., Dutta, R., Patel, R., Stewart, R., & Velupillai, S. (2020). Temporal information extraction from mental health records to identify duration of untreated psychosis. *Journal of Biomedical Semantics*, *11*(1), 2. DOI: https://doi.org/10.1186/s13326-020-00220-2.
Yamunathangam, D., Priya, C., Shobana, G., & Latha, L. (2021). An overview of topic representation and topic modelling methods for short texts and long corpus. *2021 International Conference on Advancements in Electrical, Electronics, Communication, Computing and Automation (ICAECA)*, 1–6. DOI: https://doi.org/10.1109/ICAECA52838.2021.9675579.
Zhang, D. C., & Lauw, H. (2022, July). Dynamic topic models for temporal document networks. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, & S. Sabato (Eds.), *Proceedings of the 39th international conference on machine learning* (pp. 26281–26292, Vol. 162). PMLR. https://proceedings.mlr.press/v162/zhang22n.html