When Tweets Get Viral - A Deep Learning Approach for Stance Analysis of Covid-19 Vaccines Tweets by Brazilian Political Elites

Lorena Guadalupe Barberia[1], **Pedro Henrique de Santana Schmalz**[1], **Norton Trevisan Roman**[2]

[1]Departamento de Ciência Política, Universidade de São Paulo

[2]Escola de Artes, Ciências e Humanidades, Universidade de São Paulo

{lorenabarberia,pedrosantanaschmalz,norton}@usp.br

Abstract. *Social media platforms are crucial for understanding public opinion about policy issues. In this regard, detecting stance in Twitter posts is a vital tool. In this study, we built a corpus of tweets from 2020 and 2021, annotated with stance towards COVID-19 vaccines and vaccination, and test BERTimbau as a way to automatically detect stance in such tweets. Our model reached 86% accuracy in 2020, 77% in 2021, and 79% in the combined 2020/2021 set. Our results also highlight the time-dependent nature of data distribution and, as a consequence, stance classification. Therefore, this research also contributes to the field by shedding some light on the existing methodological challenges in analyzing complex public policy debates over time.*

1. Introduction

Social media platforms, such as Twitter, are crucial for understanding public opinion on policy issues given their widespread use by the population. Its main drawback, however, lies in this popularity, in the sense that its users create enormous amounts of data every day and analyzing these data can become a cumbersome task. In this regard, automatically detecting stance in such media becomes a vital tool.

Automatic Stance Detection (SD) is a sub-field of Natural Language Processing (NLP) that seeks to automatically identify the stance of the author of some text towards a proposition or pre-chosen target [Mohammad et al. 2017]. Given its increased relevance in public debate and its widespread use by the public, Twitter has received increasing attention as a valuable source for monitoring public opinion (*e.g.* [Walker et al. 2012, Bar-Haim et al. 2017, Dey et al. 2017]).

Understanding the stance of tweet authors on a particular issue provides real-time user-generated information, which is valuable but is also challenging due to the highly nuanced and subjective nature of the language used by those expressing their opinions using this platform. To address these challenges, we introduce a *corpus* of tweets[1] by Brazilian Political Elites (*i.e.* candidates supported by their parties in local elections), collected during 2020 and 2021 and annotated with the position toward COVID-19 vaccines and

[1]Available at `https://github.com/PedroSchmalz/covid19-tweets-brazilian-mayoral-candidates` under CC BY-NC-SA 4.0.

vaccination, and propose an approach to automatically classify stance in such tweets[2].

The remainder of this article is organized as follows. Section 2 presents some related work on stance detection and the use of BERT [Devlin et al. 2018] and BERTimbau [Souza et al. 2020] in NLP tasks. Next, in Section 3 we describe the procedure we followed to collect and classify the *corpus* of tweets, along with the steps taken to train our automatic stance classifier. Section 4, in turn, presents and discusses our results. Finally, our conclusions and directions for future work are presented in Section 5.

2. Related Work

Stance detection has been extensively studied in the field of NLP, and is often used in opinion analysis (*e.g.* [Bar-Haim et al. 2017]), policy debates (*e.g.* [Somasundaran and Wiebe 2009, Addawood et al. 2017, Zhang et al. 2017, Augenstein et al. 2016, Walker et al. 2012]), social media monitoring (*e.g.* [Dey et al. 2017]) and fake news detection (*e.g.* [Lillie and Middelboe 2019]). We contribute to the stance detection literature by using a deep learning approach to identify political elites' stances on the subject of COVID-19 vaccines.

In this regard, one of the most successful deep learning models for NLP is the Bidirectional Encoder Representations from Transformers (BERT) [Devlin et al. 2018], which has achieved state-of-the-art results on a wide range of tasks. Minaee *et al.* [Minaee et al. 2021] analyze the increase in the usage of deep learning models for text classification. They review, amongst other things, the rise of Transformer-based PLMs (Pre-trained Language Models). As noted by the authors, when used for sentiment analysis on popular data sets (*e.g.* IMDB, SST-2, Amazon, etc.), BERT and its variants (BERT-large, RoBERTa, ALBERT, etc.) obtain promising results in classification tasks.

Although BERT has been successfully applied to tweet classification, challenges remain to be addressed. First, it should be noted that there are few *corpora* available in Brazilian Portuguese that could be used to this end. Still, some authors undertake the task of classifying text in Brazilian Portuguese (*e.g.* [Aguiar et al. 2021, Nascimento et al. 2015, Junqueira and Fernandes 2018, Torres et al. 2020, Brum and Nunes 2017, Silva et al. 2021]) with some of them using BERTimbau [Souza et al. 2020], a pre-trained BERT model for Brazilian Portuguese (*e.g.* [Martins 2022, Hammes and Freitas 2021, Silva and Freitas 2022]). We seek to contribute to this literature by providing an annotated data set in Portuguese for stance detection, proposing a keyword method for replicating and expanding the *corpora*, and applying BERTimbau to stance classification.

3. Material and Methods

To build our *corpus*, in January 2022 we retrieved tweets from 2020, a year with high uncertainty on the development of COVID-19 vaccines, and 2021, the beginning of the vaccination campaign that changed the debate on COVID-19 immunization from the previous year. Source profiles on Twitter were selected based on candidates registered and certified by the Brazilian Superior Electoral Court (TSE). Of the 300 candidates running

[2]Replication files available at `https://github.com/PedroSchmalz/when-tweets-get-viral-replication-files`.

for mayor at the 2020 elections in the 26 state capitals, we identified existing Twitter accounts for 243. Among them, 20 profiles were inactive[3] and 80 accounts did not publish content related to our research topic during the analyzed period and were therefore excluded from the analysis.

Our final sample consisted then of 143 mayoral candidates. We used Twitter's REST API[4]to collect all tweets from these candidates and filtered them in using a keyword selection. As noted by [Barbera et al. 2020], this method is preferred over other approaches, such as using subjective categories, because it allows for researcher control and can be replicated and even used in different media.

The set of keywords (Table 1) used in this study was developed in four test trials based on observations of spelling variations, term frequency, and usage. Orthographic and spelling issues were addressed after a preliminary analysis of common variations used by Twitter users. Moreover, we consider both lower and upper cases of terms. Subsequently, we classified tweets that could be excluded after determining that their content, while mentioning keyword terms, did not refer to COVID-19 vaccines or vaccination.

Table 1. Keywords used to retrieve tweets

Topic	Keywords
Vaccines and Vaccination (Gen.)	[Vv]acin; [Vv]assina; [Vv]acinacao [Vv]asina; [Ii]munização; [Ii]munisação; Vaccine Symbol (0)
Covid-19 Vaccines and Laboratories	CoronaVac - [Cc]orona[Vv]ac; [Cc]omuna[Vv]ac; [Ss]inovac AstraZeneca - [Aa]stra[Zz]eneca; [Aa]stra[Zz]enica; [Oo]xford; [Oo]xfort; [Oo]xfor; [Vv]axzvria Pfizer - [Pp]fizer; [Pp]eizer; [Pp]pfaizer; [Ff]aizer; [Ff]eizer; [Bb]iontech; [Bb]iontec; [Cc]omyrnaty; [Cc]omimaty; [Bb]iontech [Mm]anufacturing [Gg]mbh; Moderna - [Vv]acina da [Mm]oderna; [Vv]acina [Mm]oderna; [Ss]pikevax; [Mm]oderna [Bb]iontech; mRNA=1273; CX-024414 Sputnik - [Ss]putnik; [Ss]putinik; [Ss]putink; [Ss]putinic;[Ss]putinikV; [Gg]amaleya; [Gg]amaleia Janssen -[Jj]ansen; [Jj]anssen; [Jj]&[Jj]: [Jjohnson & [Jjohnson: [Jj]honson & [Jj]honson: [Jj]onson: [Jj] & [Jj]: [Jjohnson: [Jjjohnsons; [Jj]honson; [Jjancen;[Aa]d26.COV2S Covaxin - [Cc]ovaxin; [Cc]ovachin; [Bb]harat [Bb]iotech Novavax - [Nn]ovavax; [Cc]ovavax; [Nn]uvaxovid; NVX-CoV2373; TAK-019; SARS- CoV-2 rS with Matrix-M1 adjuvant; [Ss]erum [Ii]nstitute of [Ii]ndia; [Nn]ovavax; [Ff]ormulation Sinopharm - [Ss]inopharm; BIBP; [Ss]inofarm Other Laboratories - [Bb]utantan; [Bb]utanta; [Ff]iocruz; @fiocruz; [Ff]iocrus
Localities	[Vv]achina; [Vv]axina; [Vv]acina da [Cc]hina; [Vv]acina [Cc]hinesa; [Vv]acina [Bb]ritânica; [Vv]acina [Cc]ubana; [Vv]acina [Rr]ussa; [Vv]acina da [Rr]ussia
Additional Terms	[Dd]oriavac; [Cc]obaia; [Jj]acaré; [Jj]acare; [Dd]itadória; [Dd]itadoria; [Vv]a-china; [Vv]achinação; [Vv]a-chinacao; [Cc]olateral; [Ee]feito [Cc]olateral; [Oo]brigatoriedade; [Oo]brigar

[3]The activity status of these Twitter accounts was determined manually by coders to ensure these were professional candidate profiles with recent posts.

[4]Documentation available at: `https://developer.twitter.com/en`

As a result, of the 2,335 tweets retrieved in 2020, 1,589 referred to COVID-19 vaccines, which built our 2020 part of the *corpus*. During 2021, we managed to retrieve more than 17,000 tweets, of which we randomly selected a sample of 5,000 and, after the above-mentioned analysis, we kept 4,831 tweets related to Covid-19 vaccines. In total, the *corpus* comprises 6,420 tweets from 2020 and 2021.

Next, the *corpus* was manually annotated by five volunteers, which also reviewed each other's work. Our unit of analysis was an individual tweet and each post was classified either as *Favorable* to Covid-19 Vaccines, *Neutral*, or *Unfavorable*. Examples of the latter include calling Coronavac "*Vachina*" (vaccine + China), or expressions of doubt about vaccine effectiveness. Favorable tweets relate to those praising vaccines and their arrival. Finally, neutral tweets refer to factual posts, including news posts and those that do not communicate a clear stance about COVID-19 vaccines or vaccination in general.

Table 2 shows the class distribution for 2020, 2021, and the total for both years. During both years, the majority of tweets were classified as favorable. However, there is a higher percentage of unfavorable vaccine tweets in 2020 than in 2021 and a significant increase in neutral tweets in 2021 compared to 2020. In 2021, there was a larger volume of tweets related to Covid-19 vaccines, and most were favorable or neutral.

Table 2. Distribution of Classes

Class	2020	2021	Combined
Favorable	1,319	2,682	4,001
Neutral	82	2,085	2,167
Unfavorable	188	64	252
Total	1,589	4,831	6,420

With this *corpus* at hand, we set out to fine-tune a BERTimbau classifier [Souza et al. 2020] (a pre-trained language model for Brazilian Portuguese, based on BERT [Devlin et al. 2018]) to the task of automatic stance classification. To do so, we first randomly set (with stratified sampling) 10% of the *corpus* apart, as a final test set. We then preprocessed both sets separately, so as to avoid data leakage [Zheng and Casari 2018, Kuhn and Johnson 2019]. Preprocessing consisted of lowercasing, hyperlink removal, and other minor cleanings (retweets, emoticons, etc.).

Next, both *corpora* were tokenized with the BERT Tokenizer (WordPiece algorithm) [Wu et al. 2016, Song et al. 2020]. The models were then trained and validated at the training set (*i.e.* with 90% of the data). Training was done using five epochs, a 32 batch size, and a 2×10^{-5} learning rate for the ADAM optimizer. The model was validated through stratified 5-fold cross-validation [Sobhani 2017], so as to preserve a higher number of observations for the minority classes in all folds.

Three different models were then trained: one with observations only from 2020, another from 2021, and an additional one with the merged sample from both years. The metrics used for model evaluation were training and validation cross-entropy loss (which measures the difference between the output of the model and the expected class) [Sukhbaatar and Fergus 2014], Precision (the fraction of true positives among all positives pointed out by the classifier), Recall (the fraction of all positives the were correctly identified by the classifier) and F1-score (the harmonic mean between precision and recall) [Olson and Delen 2008].

4. Results and Discussion

Figure 1 shows the average loss in both training and validation sets, along with precision, recall and F1-score for each class, at each epoch and averaged across all folds in 2020. As it turns out, loss values seem to indicate an overfitting process after the second epoch. Figures for precision, recall and F1-score indicate that the model can successfully predict favorable tweets in 2020, peaking at the second or third epochs. At the remaining classes, performance drops, probably as a consequence of the small amount of available data. Overall, two or three epochs are enough for the model to reach stability with this data set.

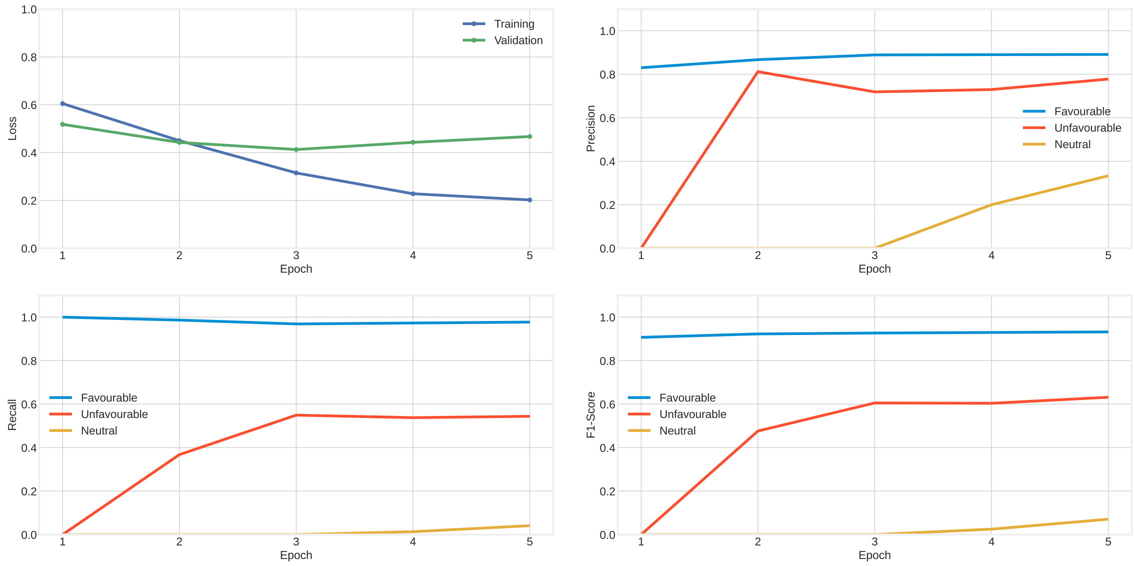

Figure 1. Loss, along with per class precision, recall and F1-score (2020)

During 2021, one sees a raise in the number of tweets related to COVID-19 vaccines and vaccination, possibly due to the authorization and start of vaccination in Brazil. In this year, although Favorable remained as the majority class, minority moved from Neutral to Unfavorable, with a pronounced increase in the Neutral class (see Table 2), thereby changing data distribution. Neutral tweets mostly corresponded to recently elected mayors advertising their municipal vaccination campaign progress.

Figure 2 shows training and validation average loss, along with precision, recall and F1-score for each class, at each epoch and averaged across all folds in 2021. As in 2020, loss values point to an overfitting after the second epoch, with precision, recall and F1-score being higher for the majority class, decreasing as the number of available examples drops, usually stabilizing at the second or third epoch. By comparing results from Figures 1 and 2, one sees how data drift impacts model performance, even for models trained at the specific data set, *i.e.* even after accounting for the new data distribution.

Considering the results obtained from 2020 and 2021, it comes as no surprise that the same behavior can be seen in the union of both sets, as illustrated in Figure 3. In this figure, one sees the model overfitting and all evaluation metrics stabilizing after the

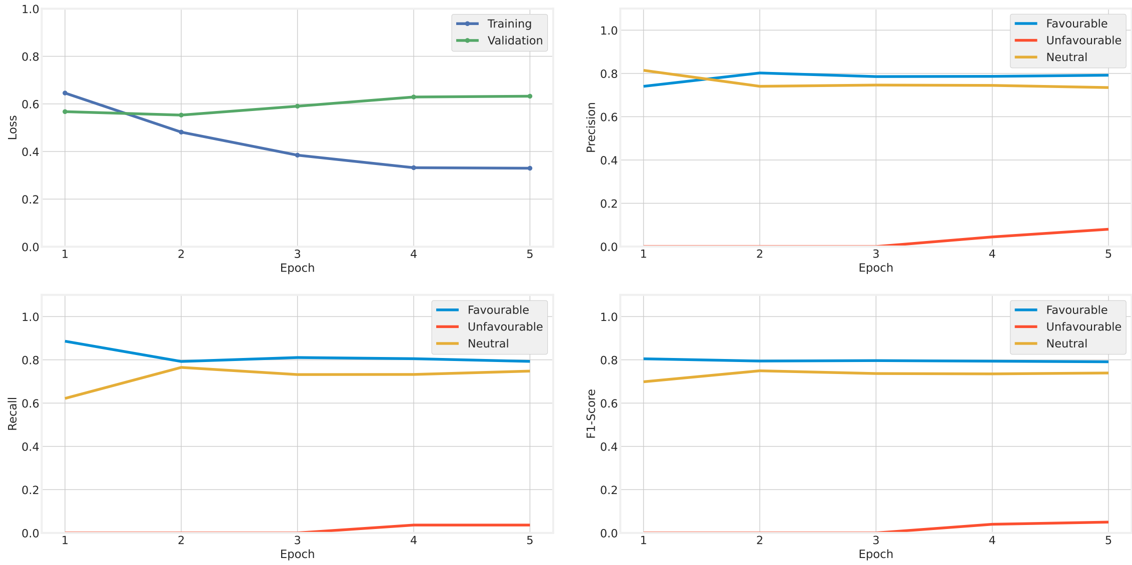

Figure 2. Loss, per class precision, recall and F1-score (2021)

second epoch. Despite this common behavior, overall performance was superior when using data from both years than from each individual year, in the sense that the model was more assertive in its classification of the minority class, even if it still lagged behind the other two classes.

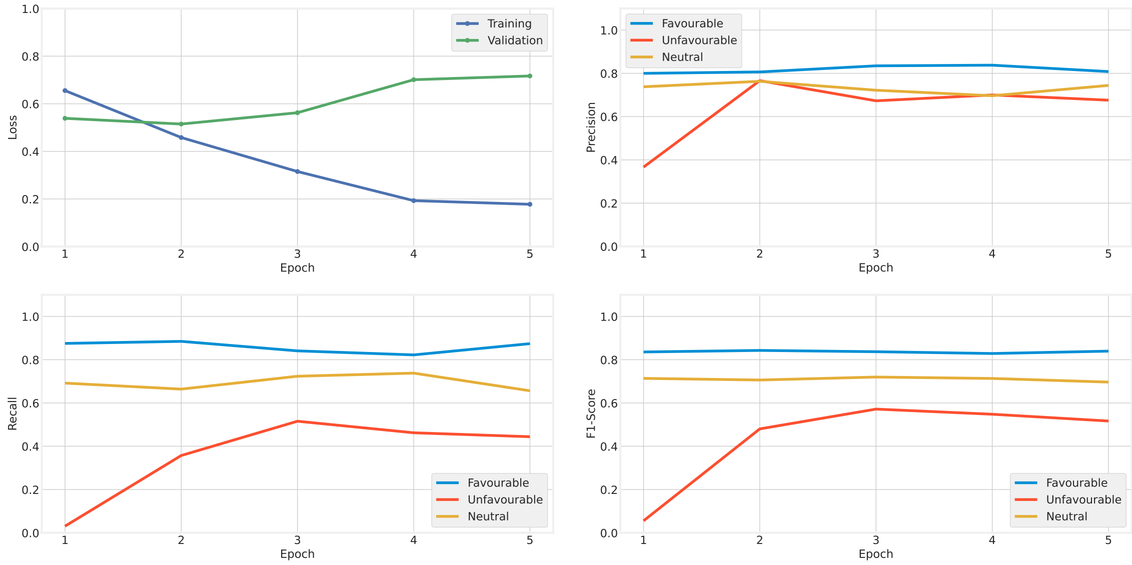

Figure 3. Loss, precision, recall and F1-score (2020 and 2021)

Average validation performance (across the five folds) for all data sets can be seen in Table 3, with a per-class breakdown being shown in Table 4. In these tables, we focus on the results at the second epoch only, so as to avoid overfitting and still reach model stability, and present both average and their associated 95% confidence interval (within parentheses). From these data, it becomes clear the difference in performance between 2020 and 2021 data. Interestingly, when dealing with 2020 and 2021 data together, there was a raise in macro-F1, reflecting the better performance at the minority classes.

Table 3. Average validation results (Epoch 2)

Data Set	Training Loss	Validation Loss	Validation Accuracy	Micro F1	Macro F1
2020	0.45 (0.43:0.46)	0.44 (0.40:0.49)	0.86 (0.84:0.88)	0.86 (0.84:0.88)	0.51 (0.44:0.57)
2021	0.48 (0.48:0.49)	0.55 (0.53:0.58)	0.77 (0.76:0.78)	0.77 (0.76:0.78)	0.67 (0.66:0.69)
2020 and 2021	0.46 (0.45:0.47)	0.52 (0.50:0.53)	0.79 (0.78:0.80)	0.79 (0.78:0.80)	0.67 (0.65:0.69)

Note: 95% Confidence Intervals in parentheses.

Table 4. Per-class average validation results (Epoch 2)

Year	Class	Precision	Recall	F1-Score	Support
2020	Favorable	0.87 (0.856:0.884)	0.986 (0.971:1.00)	0.92 (0.911:0.929)	~237
	Neutral	0 (0.00:0.00)	0 (0.00:0.00)	0 (0.00:0.00)	~15
	Unfavorable	0.81 (0.679:0.941)	0.37 (0.201:0.539)	0.476 (0.303:0.649)	~34
2021	Favorable	0.802 (0.764:0.840)	0.792 (0.735:0.849)	0.794 (0.777:0.811)	~483
	Neutral	0.740 (0.703:0.777)	0.766 (0.702:0.830)	0.748 (0.732:0.764)	~375
	Unfavorable	0 (0.00:0.00)	0 (0.00:0.00)	0 (0.00:0.00)	~12
2020 and 2021	Favorable	0.808 (0.785:0.831)	0.882 (0.85:0.914)	0.840 (0.834:0.846)	720
	Neutral	0.762 (0.716:0.808)	0.664 (0.605:0.723)	0.708 (0.692:0.724)	390
	Unfavorable	0.768 (0.701:0.835)	0.358 (0.297:0.419)	0.480 (0.425:0.535)	~45

Note: 95% Confidence Intervals in brackets.

One possible reason for this difference in class distribution may lie in the observed structural change in the COVID-19 vaccine debate. During 2020, COVID-19 vaccine availability was uncertain, and so candidates tweeted less often and only a few expressed their hope that clinical trials would turn into vaccines. When vaccines were authorized by Anvisa, Brazil's regulatory agency, and began to be offered to some adult groups in January 2021 with a slow roll-out to additional adults, there were significant rise in tweets about COVID-19 vaccines and vaccination. Interestingly, unfavorable tweets were rare in both years, especially in 2021, revealing that these individuals were not likely to voice unfavorable positions on COVID-19 vaccination.

Finally, from the models shown in Tables 3 and 4, the best performance model

Table 5. Results for the Test Set

Year	Class	Precision	Recall	F1-Score	Support
2020	Favorable	0.89	0.98	0.94	132
	Neutral	0.50	0.12	0.20	8
	Unfavorable	0.82	0.47	0.60	19
2021	Favorable	0.79	0.88	0.83	269
	Neutral	0.81	0.71	0.76	209
	Unfavorable	0.0	0.0	0.0	6
2020 and 2021	Favorable	0.79	0.87	0.83	400
	Neutral	0.73	0.65	0.69	217
	Unfavorable	0.73	0.32	0.44	25

among all folds, in terms of macro F1, was run in the test set, whose results can be seen in Table 5. Results were aligned to those obtained during validation in that the model run on 2020 data performs poorly on the neutral class and the results for unfavorable tweets were worse in 2021. Overall, the results improve over all classes for the combined 2020 and 2021 data set, but there are still problems related to the unfavorable class, with poor recall and F1-scores. In future steps, measures to deal with class imbalance may be explored.

5. Conclusions and Future Work

The present work introduced a *corpus* collected from Twitter with posts from Brazilian political elites about COVID-19 vaccines. The *corpus* was used to detect stance using a deep learning model (BERTimbau). The results indicate that class imbalance increased the difficulty of the task, given the . Poor results were obtained for the minority class in each set. This highlights that even in a short period the model can be affected by data drift, where data distribution changes, influencing the performance of the learner.

This study contributed to the understanding of political discourse during a crisis period (COVID-19 pandemic), which may generalize to other domains or moments. Although BERTimbau is recognized as state of the art for Portuguese, we are still confined to its text representation and the use of a linear layer of neurons as a classifier. In this sense, using its language model with other classifiers could yield better results. This will be a subject of future investigation. Furthermore, the five-fold approach restricts the potential generalization of the model, although it was necessary to maintain class distribution across sets. In future work, we plan to collect additional data to overcome this problem. Finally, we plan to compare our results to those by other commonly used models, following more traditional machine learning techniques, such as support vector machines (SVMs) or random forests (RFs), or other deep learning methods (RNNs, CNNs).

Acknowledgments

This work was carried out at the Center for Artificial Intelligence of the University of São Paulo (C4AI - http://c4ai.inova.usp.br/), with support by the São Paulo Research Foundation (FAPESP grant #2019/07665-4) and by the IBM Corporation. The project was also supported by the Ministry of Science, Technology and Innovation, with resources of Law N. 8.248, of October 23, 1991, within the scope of PPI-SOFTEX, coordinated by Softex and published as Residence in TIC 13, DOU 01245.010222/2022-44.

References

Addawood, A., Schneider, J., and Bashir, M. (2017). Stance classification of twitter debates. In *Proceedings of the 8th International Conference on Social Media & Society - #SMSociety17*. ACM Press.

Aguiar, A., Silveira, R., Pinheiro, V., Furtado, V., and Neto, J. A. (2021). Text classification in legal documents extracted from lawsuits in brazilian courts. In *Intelligent Systems*, pages 586–600. Springer International Publishing.

Augenstein, I., Rocktäschel, T., Vlachos, A., and Bontcheva, K. (2016). Stance detection with bidirectional conditional encoding. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 876–885, Austin, Texas. Association for Computational Linguistics.

Bar-Haim, R., Bhattacharya, I., Dinuzzo, F., Saha, A., and Slonim, N. (2017). Stance classification of context-dependent claims. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 251–261, Valencia, Spain. Association for Computational Linguistics.

Barbera, P., Boydstun, A., Linn, S., McMahon, R., and Nagler, J. (2020). Automated text classification of news articles: A practical guide. *Political Analysis*, 29(1):19–42.

Brum, H. B. and Nunes, M. d. G. V. (2017). Building a sentiment corpus of tweets in brazilian portuguese (version 1). In *11th International Conference on Language Resources and Evaluation (LREC 2018).*

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding (version 2). *arXiv*.

Dey, K., Shrivastava, R., and Kaushik, S. (2017). Twitter stance detection — a subjectivity and sentiment polarity inspired two-phase approach. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 365–372, New Orleans, LA, USA.

Hammes, L. O. A. and Freitas, L. A. d. (2021). Utilizando bertimbau para a classificação de emoções em português. In *Anais do XIII Simpósio Brasileiro de Tecnologia da Informação e da Linguagem Humana*. Simpósio Brasileiro de Tecnologia da Informação e da Linguagem Humana, Sociedade Brasileira de Computação.

Junqueira, K. T. C. and Fernandes, A. M. R. (2018). Análise de sentimento em redes sociais no idioma português com base em mensagens do twitter. In *IX Computer on the Beach*, Brazil.

Kuhn, M. and Johnson, K. (2019). *Feature Engineering and Selection: A Practical Approach for Predictive Models*. Chapman and Hall/CRC.

Lillie, A. E. and Middelboe, E. R. (2019). Fake news detection using stance classification: A survey (version 1). *arXiv*.

Martins, G. F. (2022). Um estudo utilizando-se de análise de sentimentos e aprendizado de máquina para a classificação de tweets sobre a vacinação no brasil. Bachelor's thesis, Universidade Federal do Rio Grande do Sul. In Portuguese.

Minaee, S., Kalchbrenner, N., Cambria, E., Nikzad, N., Chenaghlu, M., and Gao, J. (2021). Deep learning-based text classification. *ACM Computing Surveys*, 54(3):1–40.

Mohammad, S. M., Sobhani, P., and Kiritchenko, S. (2017). Stance and sentiment in tweets. *ACM Transactions on Internet Technology*, 17(3):1–23.

Nascimento, P., Osiek, B., and Xexéo, G. (2015). AnÁlise de sentimento de tweets com foco em notÍcias. *Revista Eletrônica de Sistemas de Informação*, 14(2):2.

Olson, D. L. and Delen, D. (2008). *Advanced Data Mining Techniques*. Springer, 1st edition.

Silva, F. and Freitas, L. (2022). Brazilian portuguese hate speech classification using bertimbau. In *International FLAIRS Conference Proceedings*, volume 35. University of Florida George A Smathers Libraries.

Silva, N. F. F. d., Silva, M. C. R., Pereira, F. S. F., Tarrega, J. P. M., Beinotti, J. V. P., Fonseca, M., Andrade, F. E. d., and de Carvalho, A. C. P. d. L. F. (2021). Evaluating topic models in portuguese political comments about bills from brazil's chamber of deputies. In *Intelligent Systems*, pages 104–120. Springer International Publishing.

Sobhani, P. (2017). *Stance Detection and Analysis in Social Media*. PhD thesis, University of Ottawa, Ottawa, Canada. Thesis submitted to the Faculty of Graduate and Postdoctoral Studies in partial fulfillment of the requirements for the Ph.D. degree in Computer Science, School of Electrical Engineering and Computer Science, Faculty of Engineering.

Somasundaran, S. and Wiebe, J. (2009). Recognizing stances in online debates. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 226–234, Suntec, Singapore. Association for Computational Linguistics.

Song, X., Salcianu, A., Song, Y., Dopson, D., and Zhou, D. (2020). Fast wordpiece tokenization. *arXiv*.

Souza, F., Nogueira, R., and Lotufo, R. (2020). Bertimbau: Pretrained bert models for brazilian portuguese. In *Intelligent Systems*, pages 403–417. Springer International Publishing.

Sukhbaatar, S. and Fergus, R. (2014). Learning from noisy labels with deep neural networks. *arXiv preprint*, arXiv:1406.2080(2(3)):4.

Torres, B. A., Moraes, L. C., and Pimenta, D. A. C. (2020). Um estudo da aceitação das vacinas contra a covid-19 na rede social twitter utilizando ferramentas de análise sentimental. In *II Workshop de Tecnologia da Fatec Ribeirão Preto*, volume 1-2.

Walker, M. A., Anand, P., Abbott, R., Tree, J. E. F., Martell, C., and King, J. (2012). That is your evidence?: Classifying stance in online political debate. *Decision Support Systems*, 53(4):719–729.

Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, , Gouws, S., Kato, Y., Kudo, T., Kazawa, H., and ... Dean, J. (2016). Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv*.

Zhang, S., Qiu, L., Chen, F., Zhang, W., Yu, Y., and Elhadad, N. (2017). We make choices we think are going to save us. In *Proceedings of the 26th International Conference on World Wide Web Companion - WWW '17 Companion*. ACM Press.

Zheng, A. and Casari, A. (2018). *Feature Engineering for Machine Learning: Principles and Techniques for Data Scientists*. O'Reilly Media.