Stance Prediction from Multimodal Social Media Data

Laís C. L. Cavalheiro **Matheus C. Pavan** **Ivandré Paraboni**
School of Arts, Sciences and Humanities
University of São Paulo
Av Arlindo Bettio 1000, São Paulo, Brazil
{laiscarraro,matheus.pavan,ivandre}@usp.br

Abstract

Stance prediction – the computational task of inferring attitudes towards a given target topic of interest – relies heavily on text data provided by social media or similar sources, but it may also benefit from non-text information such as demographics (e.g., users' gender, age, etc.), network structure (e.g., friends, followers, etc.), interactions (e.g., mentions, replies, etc.) and other non-text properties (e.g., time information, etc.). However, so-called hybrid (or in some cases multimodal) approaches to stance prediction have only been developed for a small set of target languages, and often making use of count-based text models (e.g., bag-of-words) and time-honoured classification methods (e.g., support vector machines). As a means to further research in the field, in this work we introduce a number of text- and non-text models for stance prediction in the Portuguese language, which make use of more recent methods based on BERT and an ensemble architecture, and ask whether a BERT stance classifier may be enhanced with different kinds of network-related information.

1 Introduction

Standard stance prediction concerns the inference of for/against attitudes towards a target topic of interest from text data. The task may be seen as analogous to sentiment (e.g., positive or negative) analysis, but stance and sentiment do not necessarily correlate (Aldayel and Magdy, 2021). For instance, consider the following statement:

> *'People who refuse the vaccine should be banned from entering the premises'*

In this example, given the intended target 'vaccination', the statement suggests a favourable stance. Still, the overall sentiment (particularly in the use of the word 'banned') may be regarded as being more on the negative side.

Computational models of stance prediction, which often take social media text as an input, have been applied to a wide range of topics, including moral or social issues (Pavan et al., 2023; Geiss et al., 2022), politics (Darwish et al., 2017; Lehmann and Derczynski, 2019; Cignarella et al., 2020), and others. The task has become particularly popular in the field since the SemEval-2016 shared task (Mohammad et al., 2016) and accompanying corpus, focusing on stance prediction from Twitter text in the English language.

In addition to using text data (Zhang et al., 2020; Allaway et al., 2021; Pavan et al., 2020), recent work in stance prediction has addressed the use of non-text data as well (Aldayel and Magdy, 2021). Studies of this kind are largely motivated by the notion of *homophily* (McPherson et al., 2001), that is, the concept of 'similarity breeds connection', and take into account a range of well-known non-linguistic stance predictors. These include, for instance, the use of demographics information (e.g., users' gender, age, etc.) (Lehmann and Derczynski, 2019; Geiss et al., 2022), network structure (e.g., social media friends, followers, etc.) (Lai et al., 2020b), interactions (e.g., mentions, replies, retweets, etc.) (Magdy et al., 2016; Darwish et al., 2017), and other network properties (e.g., number of replies, time information, etc.) (Espinosa et al., 2020), which are often combined with standard text models. Hybrid models of this kind, although not necessarily using images, audio or other media, will be hereby called *multimodal*.

Existing work in stance prediction based on hybrid data are often derived from two main NLP initiatives: (1) the Iberaval-2017 shared task (Taulé et al., 2017), devoted to the Catalan and Spanish languages, and (2) the EVALITA-2020 shared task (Cignarella et al., 2020) for the Italian language, in both cases providing social media corpora and accompanying non-text data. The Iberaval-2017

Proceedings of Recent Advances in Natural Language Processing, pages 242–248
Varna, Sep 4–6, 2023
https://doi.org/10.26615/978-954-452-092-2_027

corpus is however limited to text and gender information, whereas EVALITA-2020 is a truly hybrid corpus that provides both text and network-related information. Similar hybrid resources (most notably for English and a few other European languages) do exist (Lehmann and Derczynski, 2019; Lai et al., 2020b). However, we are not aware of any study in stance prediction based on hybrid data that has been devoted to our target language – Portuguese.

In addition to the language gap, we notice that existing work in stance prediction based on hybrid data often relies on text representations based on feature counts (e.g., bag-of-words), in many cases combined with support vector machine (SVM) or other similarly time-honoured classification methods. As discussed in (Espinosa et al., 2020), some of these choices may be explained by the challenges involved in combining large text representations with (usually) much smaller non-text models (e.g., representing interactions, demographics, etc.).

Based on these observations, in this work we investigate whether stance prediction using a more contemporary text representation – namely, built from BERT (Devlin et al., 2019) – may be enhanced with different kinds of network-related information. Using a large multimodal stance corpus in the Portuguese language as an input, we envisaged a number of experiments to assess stance prediction models in an ensemble architecture of text- and network-related classifiers. The contributions made in this work are as follows:

- Hybrid (text and non-text) stance prediction approach using BERT and ensemble of classifiers.
- Stance prediction models for the Portuguese language.
- Best-performing strategy combining text, structural, and interaction information.

The remainder of this article is structured as follows. Section 2 reviews existing work in stance prediction and related resources. Section 3 describes our current experiment setting by presenting the models under consideration and the corpus to be taken as train and test data. Section 4 presents the results of our experiments. Finally, Section 5 presents our final remarks and future opportunities of investigation.

2 Related work

Table 1 summarises a number of recent studies that are devoted to stance prediction using text and non-text data, or which introduce a dataset for the task. These are organised according to text source (p=political discourse, r=Reddit, t=Twitter), language (Ar=Arabic, Ca=Catalan, Da=Danish, En=English, Fr=French, It=Italian, Sp=Spanish), the number of learning instances, the choice of text features (w=word, p=part-of-speech tags, c=characters, we=word embeddings, bp=BERT class probabilities), the kind of non-text feature under consideration (dem=demographics, m=mentions, r=replies, rt=retweets, dom=domain- or task-specific information, fr=friends, fo=followers, h=time of posting, dist=distance to other users, int=interactions), and main computational method (SVM=support vector machine, LR=logistic regression, etc.). Further details are discussed below.

Generally speaking, existing work in the field is largely based on Twitter, a preference that may be motivated by the ease of access to text and non-text data through the Twitter API.

All of the selected studies are devoted to English or other European languages. Among these, there are several studies focuses on Catalan and Spanish, including the work in (Lai et al., 2017), which has been developed in the light of the Iberaval-2017 shared task (Taulé et al., 2017), and which obtained the overall best results among the participant systems. Similarly, several recent studies have focused on the Italian language, including (Espinosa et al., 2020), which was the overall best-performing system at the EVALITA-2020 task B (contextual stance detection) shared task (Cignarella et al., 2020).

Although presently not shown, we notice also that language and topic choices usually come hand-in-hand, that is, existing stance datasets tend to favour target topics that are of interest to a rather local audience. This trend has been observed since the SemEval-2016 English corpus (Mohammad et al., 2016), which includes US-specific topics such as Trump and Hillary Clinton among more general topics such as climate change, etc. Similarly, the Catalan/Spanish study in (Taulé et al., 2017) focuses on stances towards the Catalan independence movement; the Arabic study in (Darwish et al., 2017) addresses the issue of stance towards Saudi/Egyptian islands ownership; the Danish cor-

Study	Source	Language	Inst.	Text	Non-text	Method
(Magdy et al., 2016)	t	En	336.3K	w	dem,m,r,rt	SVM
(Taulé et al., 2017)	t	Ca,Sp	10.8K		dem,m	corpus release
(Lai et al., 2017)	t	Ca,Sp	10.8K	w,p,c	m	SVM,LR
(Darwish et al., 2017)	t	Ar	33K	w	rt,m	similarity
(Lehmann and Derczynski, 2019)	p	Da	898	we	dem,dom	LSTM,MLP
(Lai et al., 2020a)	t	En,Fr,It,Sp,Ca	14.4K	w,p,c	m,dom	LSTM,CNN
(Cignarella et al., 2020)	t	It	3.2K		fr,fo,rt,m,h	corpus release
(Espinosa et al., 2020)	t	It	3.2K	bp	fr,fo,h,dist	voting
(Lai et al., 2020b)	t	En	1.8K	w,c	fr,fo,h,dom	SVM
(Geiss et al., 2022)	r	En	2,717K	w,we	int,dem	SVM

Table 1: Stance prediction methods using non-text features.

pus in (Lehmann and Derczynski, 2019) focuses on Danish immigration policies; the Italian shared task in (Cignarella et al., 2020) focuses on stance towards the Sardine's movement, and so forth. This relation between language and topic is likely to be necessary as a means to model meaningful tasks, and to obtain the necessary amount of data. However, this may also suggest that in studies focused on a particular language – as in the present work, focused on Portuguese – the benefits afforded by using corpora developed for other languages may be limited.

Dataset sizes are often within a few thousand instances, which is close to the text-only SemEval-2016 stance corpus in (Mohammad et al., 2016) with 4.2K labelled instances. We notice, however, that the two largest stance corpora in the present review – used in (Magdy et al., 2016) and (Geiss et al., 2022) – are not manually annotated at the text level, resorting to label propagation or similar methods instead.

Text data is usually modelled in a bag-of-words or similar approach (e.g., using words, part-of-speech tags, or character n-grams), which may be explained by the need to combine well-balanced sets of text and non-text features as discussed in (Espinosa et al., 2020). Moreover, since some of these studies are more focused on the dataset (and not necessarily on any particular classifier method), the use of simpler text representations tends to be preferred. As a result, the use of word embeddings is less common, and the use of more recent transformed-based language models such as BERT (Devlin et al., 2019) appears only in one study, the aforementioned work in (Espinosa et al., 2020). Even in this case, however, the language model is used only as a means to obtain class probabilities to be taken as learning features, rather than modelling a full text representation directly.

Regarding the use of non-text features, Twitter-related network features are common and, to a lesser extent, this is true also of demographics (mostly gender), and domain-dependent information (e.g., information related to political affiliation, opposition, etc.). In the case of Reddit-based studies, we notice that non-text information is largely limited to interactions between users, whereas Twitter-based studies may also make use of friends and followers information, among others.

Computational methods for stance prediction based on hybrid data are often simple, well-known classifiers such as SVM or logistic regression. This may be explained by the relatively low dimension of these models if compared to what would normally be required for, e.g., text modelling. Moreover, in the case of shared tasks participant systems, it may be the case that execution times are also a concern, which might have favoured the use of these methods over more computationally-expensive (e.g., neural) alternatives.

3 Stance prediction based on hybrid data

The main objective of the present work is to investigate which combinations of network-related information, if any, may be added to an otherwise standard text-based model to improve stance classification results in the Portuguese language. As in the work described in (Espinosa et al., 2020), which is devoted to the Italian language, we will also focus on Twitter data, and on the use of friends, followers and mentions (e.g., users with whom they interact on Twitter) information. However, instead of resorting to BERT class probabilities as learning features, we shall make use of a full text representation built from BERT, and will investigate a range of simple ensemble approaches to combine (large) text and (comparatively small) non-text feature sets as a single model.

3.1 Models

We envisaged an ensemble approach to stance prediction that combines text features (or 't' for short) with one or more sources of network-related information, namely, friends ('fr'), followers ('fo'), mentions ('me'), or any combination of these, making seven binary (for/against) classifier alternatives as follows:

- t+fr: *text+friends*
- t+fo: *text+followers*
- t+me: *text+mentions*
- t+fr+fo: *text+friends+followers*
- t+fr+me: *text+friends+mentions*
- t+fo+me: *text+followers+mentions*
- t+fr+fo+me: *text+friends+followers+mentions*

In all our classifier alternatives, the basic text component is a standard BERT classifier described in (da Costa et al., 2023). This consists of token and Bi-LSTM layers followed by multi-head self-attention and a dense layer using sigmoid as an activation function and dropout. The token embedding layer is built from a BERT model pre-trained on Brazilian Portuguese Twitter data called BERTabaporu (da Costa et al., 2023)[1].

Friends and followers features correspond to the lists of all friends and followers of every individual as provided by the Twitter API. Mentions features correspond to the list of all usernames mentioned in the corpus timelines (and which may or may not coincide with a friend, a follower, or both), which are marked by the '@' character in Twitter text data. Features are modelled in a so-called bag-of-users approach using tf-idf counts (i.e., building bag-of-friends, bag-of-followers, and a bag-of-mentions vectors, respectively). Feature vectors are taken as an input to a logistic regression classifier with parameters C, tol and $penalty$ optimised through grid search for each task.

Regarding the ensemble approach, predictions made by the individual model components are combined as a single output by majority voting. In the case of a tie, a random (for/against) prediction is made. As an example, Figure 1 illustrates the architecture of the full model *t+fr+fo+me*.

[1]https://huggingface.co/pablocosta/bertabaporu-large-uncased

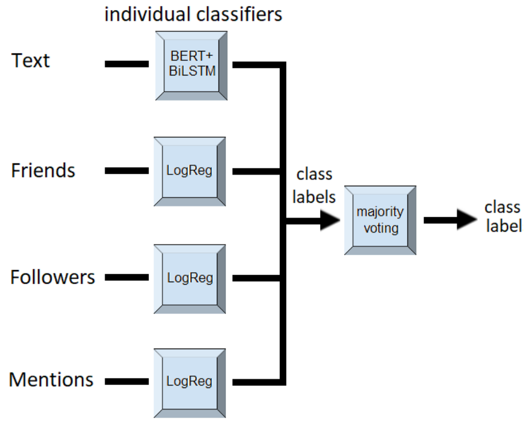

Figure 1: Example model architecture.

The use of BERT prediction in the (t)ext component of the ensemble is comparable to the use of BERT class probabilities in (Espinosa et al., 2020). In the present work, however, we use the actual label (for/against) predictions rather than class probabilities.

3.2 Data

The present work uses the Twitter corpus UstanceBR r2, whose preliminary version (r1) appeared in (Pavan and Paraboni, 2022). The corpus conveys stances towards six topics (two Brazilian presidents, two Covid-related treatments, and two local institutions) often regarded as having either a liberal or a conservative political leaning.

The corpus contains about 46.8K manually labelled tweets in the Portuguese language, and network-related information representing their friends, followers, and mentions. The text portion of the corpus has been previously applied to text-only stance prediction in (Pavan and Paraboni, 2022), but the use of its non-text portion in a hybrid setting – as in the present work – is novel.

Table 2 summarises text- and network-related corpus descriptive statistics across target topics by presenting their number of instances (tweets), tokens, friends (Fr), followers (Foll), and mentions (Ment).

4 Evaluation

All models were created and evaluated using the original train-test split provided by the corpus. Evaluation was carried out by computing average F1 scores. Statistical significance was assessed by using a McNemar test (McNemar, 1947) to compare model pairs. Table 3 summarises stance pre-

Target	Inst.	Tokens	Fr.	Foll.	Ment.
Lula	8,320	422K	463K	677K	98K
Bolsonaro	9,414	259K	346K	536K	60K
Hydrox.	7,995	278K	577K	732K	406K
Sinovac	7,973	253K	821K	1164K	488K
Church	7,137	322K	962K	1547K	183K
Globo TV	6,013	215K	743K	1168K	122K

Table 2: Data descriptive statistics.

Model	Lula	Bols	Hydr	Sino	Church	Globo
t+fr	0.88	0.91	0.86	0.86	0.82	0.81
t+fo	0.85	0.90	0.85	0.82	0.70	0.79
t+me	0.87	0.91	0.89	0.87	0.83	**0.82**
t+fr+fo	0.91	0.94	0.88	0.84	**0.86**	0.75
t+fr+me	**0.92**	**0.96**	**0.94**	**0.92**	0.83	0.79
t+fo+me	**0.92**	0.95	0.93	0.90	0.85	0.79
t+fr+fo+me	0.91	0.95	0.90	0.87	0.84	0.77

Table 3: Stance prediction F1 results using (t)ext, (fr)iend, (fo)llower, and (me)ntion features. The highest F1 score for each task is highlighted.

diction results obtained by the models described in the previous sections across the six target topics.

Results show that simply using all available information, as provided by the full model *t+fr+fo+me*, is not the best strategy at all. On the contrary, it is the use of text, friends and mentions information alone, as provided by *t+fr+me*, that actually delivers best results across most topics, although clearly the advantage over the second best alternative may in some cases be minimal.

In order to verify the possible overall advantage of *t+fr+me*, this and the other top-performing model for each topic were compared by using a McNemar test. In doing so, the advantage afforded by *t+fr+me* over the selected alternative strategy across topics was found to be statistically significant as follows: Lula ($\chi = 9.8$, p <0.01), Bolsonaro ($\chi = 3.0$, p <0.05), Hydroxychloroquine ($\chi = 8.0$, p <0.05), Sinovac ($\chi = 4.0$, p <0.001), Church (not significant), and Globo TV ($\chi = 21.0$, p <0.001). This outcome further suggests a general preference of *t+fr+me* over the alternatives.

5 Final remarks

This paper has addressed the issue of stance prediction by reporting a number of experiments that combine text data with network-related information – represented by Twitter friends, followers and mentions – in a voting ensemble architecture. Results show that the use of friends and mentions, but not followers, obtained overall best results for the present setting.

The present work leaves a number of opportunities for further improvement. First, we notice that the present ensemble architecture is limited to the use of a simple majority voting method, and that more sophisticated strategies may increase overall model accuracy. This may be the case, for instance, of stacking (Wolpert, 1992), and others.

Second, the current bag-of-users approach – which has been taken as the basis for the present network-related models – may be replaced for a dense network representation provided by node embeddings. Models of this kind, which may be computed from, e.g., node2vec (Grover and Leskovec, 2016), would make the current friends, followers and mentions models more informative (or less sparse), and this may have a positive impact on the current results.

Regarding the text portion of the model, the current task may benefit from multiple, well-known NLP methods and applications. Among these, we may consider the use of hate speech detection methods (Basile et al., 2019; Mishra et al., 2019; da Silva et al., 2020), authorship attribution (Custódio and Paraboni, 2021; Barlas and Stamatatos, 2021), or author profiling (López-Santillán et al., 2020; Rangel et al., 2020). The latter, comprising the computational task of determining individuals' demographics from text, may help determine their stance towards a particular topic by taking into account, for instance, information regarding their political orientation (Flores et al., 2022), personality traits (Verhoeven et al., 2016; dos Santos et al., 2019), moral values (dos Santos and Paraboni, 2019; Pavan et al., 2023), and others.

Finally, it is worth noting that, for simplicity, our task definition has been presently limited to binary (for/against) stance classification. In more realistic settings, however, it would be arguably useful to consider the intermediate (or neutral) class as well. This possible extension is also left as a suggestion of future work.

Acknowledgements

The present research has been supported by the São Paulo Research Foundation (FAPESP grant #2021/08213-0).

References

Abeer Aldayel and Walid Magdy. 2021. Stance detection on social media: State of the art and trends. *Information Processing & Management*, 58(4):102597.

Emily Allaway, Malavika Srikanth, and Kathleen McKeown. 2021. Adversarial learning for zero-shot stance detection on social media. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4756–4767, Online. Association for Computational Linguistics.

Georgios Barlas and Efstathios Stamatatos. 2021. A transfer learning approach to cross-domain authorship attribution. *Evolving Systems*, 12:625–643.

Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Rangel, Paolo Rosso, and Manuela Sanguinetti. 2019. SemEval-2019 Task 5: Multilingual Detection of Hate Speech Against Immigrants and Women in Twitter. In *Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval-2019)*. Association for Computational Linguistics.

Alessandra Cignarella, Mirko Lai, Cristina Bosco, Viviana Patti, and Paolo Rosso. 2020. SardiStance @ EVALITA2020: Overview of the task on stance detection in italian tweets. In *CEURS Proceedings vol. 2765*, pages 177–186, online. CEUR-WS.org.

Pablo Botton da Costa, Matheus Camasmie Pavan, Wesley Ramos dos Santos, Samuel Caetano da Silva, and Ivandré Paraboni. 2023. BERTabaporu: assessing a genre-specific language model for Portuguese NLP. In *Recents Advances in Natural Language Processing (RANLP-2023)*, Varna, Bulgaria.

José Eleandro Custódio and Ivandré Paraboni. 2021. Stacked authorship attribution of digital texts. *Expert Systems with Applications*, 176:114866.

Kareem Darwish, Walid Magdy, and Tahar Zanouda. 2017. Improved stance prediction in a user similarity feature space. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, pages 145–148, New York, USA. Assoc. for Computing Machinery.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

Maria S. Espinosa, Rodrigo Agerri, Alvaro Rodrigo, and Roberto Centeno. 2020. SardiStance: Combining Textual, Social and Emotional Features. In *7th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020)*. CEUR-WS.org.

Arthur Marçal Flores, Matheus Camasmie Pavan, and Ivandré Paraboni. 2022. User profiling and satisfaction inference in public information access services. *Journal of Intelligent Information Systems*, 58(1):67–89.

Henri-Jacques Geiss, Flora Sakketou, and Lucie Flek. 2022. OK boomer: Probing the socio-demographic divide in echo chambers. In *10th International Workshop on Natural Language Processing for Social Media*, pages 83–105, Seattle, Washington USA. Assoc. for Computational Linguistics.

Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable Feature Learning for Networks. In *KDD16: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 855–864, San Francisco, USA. Association for Computing Machinery.

Mirko Lai, Alessandra Cignarella, Delia Hernandez Farias, Cristina Bosco, Viviana Patti, and Paolo Rosso. 2020a. Multilingual stance detection in social media political debates. *Computer Speech and Language*, 63.

Mirko Lai, Alessandra Cignarella, and Delia Hernandez-Farias. 2017. itacos at ibereval2017: Detecting stance in Catalan and Spanish tweets. In *IberEval-2017 proceedings*, Murcia, Spain. CEUR-WS.org.

Mirko Lai, Viviana Patti, Giancarlo Ruffo, and Paolo Rosso. 2020b. #Brexit: Leave or remain? the role of user's community and diachronic evolution on stance detection. *Journal of Intelligent & Fuzzy Systems*, 39(2):2341–2352.

Rasmus Lehmann and Leon Derczynski. 2019. Political stance in Danish. In *22nd Nordic Conference on Computational Linguistics*, pages 197–207, Turku, Finland. Linköping University Electronic Press.

Roberto López-Santillán, Manuel Montes-Y-Gómez, Luis Carlos González-Gurrola, Graciela Ramírez-Alonso, and Olanda Prieto-Ordaz. 2020. Richer document embeddings for author profiling tasks based on a heuristic search. *Information Processing & Management*, 57(4).

Walid Magdy, Kareem Darwish, Norah Abokhodair, Afshin Rahimi, and Timothy Baldwin. 2016. #ISISisNotIslam or #deportallmuslims? predicting unspoken views. In *8th ACM Conference on Web Science*, pages 95–106, New York, NY, USA. Assoc. for Computing Machinery.

Quinn McNemar. 1947. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2):153–157.

Miller McPherson, Lynn Smith-Lovin, and James M. Cook. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444.

Pushkar Mishra, Marco Del Tredici, Helen Yannakoudakis, and Ekaterina Shutova. 2019. Abusive language detection with graph convolutional networks. In *Proceedings of NAACL-HLT 2019*, pages 2145–2150, Minneapolis, USA. Association for Computational Linguistics.

Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. SemEval-2016 task 6: Detecting stance in tweets. In *10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41, San Diego, California. Assoc. for Computational Linguistics.

Matheus Camasmie Pavan and Ivandré Paraboni. 2022. Cross-target stance classification as domain adaptation. In *Advances in Computational Intelligence - MICAI 2022 - Lecture Notes in Artificial Intelligence vol 13612*, pages 15–25, Cham. Springer Nature Switzerland.

Matheus Camasmie Pavan, Vitor Garcia dos Santos, Alex Gwo Jen Lan, Jo ao Trevisan Martins, Wesley Ramos dos Santos, Caio Deutsch, Pablo Botton da Costa, Fernando Chiu Hsieh, and Ivandré Paraboni. 2023. Morality classification in natural language text. *IEEE transactions on Affective Computing*, 14(1):857–863.

Matheus Camasmie Pavan, Wesley Ramos dos Santos, and Ivandré Paraboni. 2020. Twitter Moral Stance Classification using Long Short-Term Memory Networks. In *9th Brazilian Conference on Intelligent Systems (BRACIS). LNAI 12319*, pages 636–647. Springer.

F. Rangel, A. Giachanou, B. Ghanem, and P. Rosso. 2020. Overview of the 8th Author Profiling Task at PAN 2020: Profiling Fake News Spreaders on Twitter. In *CLEF 2020 Labs and Workshops, Notebook Papers - CEUR Workshop Proceedings vol. 2696*. CEUR-WS.org.

Wesley Ramos dos Santos and Ivandré Paraboni. 2019. Moral Stance Recognition and Polarity Classification from Twitter and Elicited Text. In *Recents Advances in Natural Language Processing (RANLP-2019)*, pages 1069–1075, Varna, Bulgaria. INCOMA Ltd.

Wesley Ramos dos Santos, Ricelli Moreira Silva Ramos, and Ivandré Paraboni. 2019. Computational personality recognition from facebook text: psycholinguistic features, words and facets. *New Review of Hypermedia and Multimedia*, 25(4):268–287.

Samuel Caetano da Silva, Thiago Castro Ferreira, Ricelli Moreira Silva Ramos, and Ivandré Paraboni. 2020. Data driven and psycholinguistics motivated approaches to hate speech detection. *Computación y Systemas*, 24(3):1179–1188.

Mariona Taulé, Maria Antònia Martí, Francisco Manuel Rangel Pardo, Paolo Rosso, Cristina Bosco, and Viviana Patti. 2017. Overview of the task on stance and gender detection in tweets on catalan independence at IberEval 2017. In *IberEval-2017 proceedings*, pages 157–177, Murcia, Spain. CEUR-WS.org.

B. Verhoeven, W. Daelemans, and B. Plank. 2016. TwiSty: a multilingual Twitter Stylometry corpus for gender and personality profiling. In *10th International Conference on Language Resources and Evaluation (LREC-2016)*, pages 1632–1637, Portoroz, Slovenia. ELRA.

David H. Wolpert. 1992. Stacked generalization. *Neural networks*, 5(2):241–259.

Bowen Zhang, Min Yang, Xutao Li, Yunming Ye, Xiaofei Xu, and Kuai Dai. 2020. Enhancing cross-target stance detection with transferable semantic-emotion knowledge. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3188–3197, Online. Association for Computational Linguistics.