HiTZ-IXA at PoliticES-IberLEF2023: Document and Sentence Level Text Representations for Demographic Characteristics and Political Ideology Detection

Joseba Fernandez de Landa, Rodrigo Agerri

HiTZ Basque Center for Language Technologies - Ixa NLP Group, University of the Basque Country UPV/EHU

Abstract

In this paper we describe our participation to the PoliticES 2023 shared task held at IberLEF 2023. The task focuses on extracting demographic and political information from tweets, and it is structured as an author profiling task. Our participation is focused on developing a multi-level textual representation that combines both the tweets text and user representations. This approach allows us to effectively capture and integrate social information, including demographic and ideological traits. Furthermore, our text-based features leverage document and sentence information, amalgamating specific and general aspects. The combination of both social and textual features results in a remarkable improvement in overall performance across the various text classification tasks proposed within the task. An additional benefit of our approach is its robustness and generalization capability, as it performs competitively using same features across all traits. Finally, we address potential memory constraints by efficiently managing extensive timelines or documents, segmenting them into individual sentences or tweets while keeping the document level information. Our technique offers promising results in effectively handling large-scale textual data in document classification tasks. Our system achieved the second highest score for the overall PoliticES task and the best score for predicting the *profession* category.

Keywords

Demographic Traits detection, Political Ideology detection, Author Profiling, Computational Social Science, Natural Language Processing

1. Introduction

This paper describes the HiTZ-IXA team participation in PoliticEs2023 shared task [1] organised in IberLEF 2023 [2], which consists of extracting demographic and political information from texts. Framed as an author profiling task, the objective is to extract Twitter user's characteristics based on 80 distinct text-based documents per author. By leveraging text-based Twitter data in Spanish language, the aim is to extract demographic traits including gender and profession, as well as, political ideology approached from both a binary and a multiclass perspective, from a given set of tweets.

There exists a significant interest in extracting demographics and ideology from Social Media,

IberLEF 2023, September 2023, Jaén, Spain
joseba.fernandezdelanda@ehu.eus (J. F. d. Landa); rodrigo.agerri@ehu.eus (R. Agerri)
0000-0001-6067-3571 (J. F. d. Landa); 0000-0002-7303-7598 (R. Agerri)

as it represents a means of gaining deeper insights into society. Twitter has become a source of spontaneously generated textual data for many human languages, and its use for doing demographic and ideological inferences increasingly common [3, 4]. As for most research topics in Natural Language Processing (NLP), recent works have experimented with Transformer-based [5] contextualized sentence embeddings for user level demographic prediction such as age or gender [6, 7]. Thus, Transformer-based approaches were the most common method in PoliticES 2022 shared task [8]. However, a limitation of those approaches is that they are usually centered on document or sentence level representations only.

In order to harness user-related data for text classification problems in social media, tasks such as stance detection have been approached from a perspective that utilizes both author and tweet characteristics to infer stance at tweet level [9, 10, 11] . In those works, authors are represented through Twitter interactions such as *friends*, *retweets*, *quotes*, or *replies*, since there is a lack of accessible data for utilizing author-level texts. These author representations are subsequently combined with textual representations derived from specific tweets, resulting in improved overall performance [11]. Therefore, taking this idea as a starting point, our approach for PoliticES focuses on the combination of user and tweet-level representations, but unlike previous aforementioned work, both derived exclusively from textual data.

Thus, in this paper we present a multi-level textual representation combining tweet and user representations in order to embed social information such as demographic and ideological traits. The contribution of the proposed text-based features lies in their ability to leverage document and sentence information, combining features derived from both sentence and author-based representations. Results demonstrate that such integration of user and tweet representation levels enhances the ability to capture meaningful information, thereby improving performance in various text classification tasks.

Our method is robust and exhibits good generalization capabilities obtaining good performance across all tasks using the same features. Furthermore, it shows the capacity to effectively manage extensive timelines or documents by dividing them into individual sentences or tweets, thereby mitigating potential memory constraints. The official results show that our system was ranked 2nd from 11 participants among the general task and in the 1st position for the *profession* category.

2. Related Work

In previous work related to this topic, we must highlight the previous task PoliticES 2022 [8], which focused on author profiling by employing text-based data in order to extract demographic and ideological traits. Most of the systems participanting in the task were based on Transformer [5] models. More specifically, there is a notable presence of monolingual models in Spanish, especially BETO [12] and MarIA [13]. The rest of the section provides an overview of the key features exhibited by the top four models presented on the 2022 task.

The first model, proposed by Carrasco and Rosillo [14], employs 512 token-blocks comprising tweets from the same author in the dataset, along with additional data, to fine-tune a combined model of BETO and MarIA. This combined model is used to predict labels at the token-block level. Subsequently, user characteristics are predicted using a majority vote strategy based on the aforementioned token-blocks.

In the second model, Villa-Cueva et al. [15] introduce PolitiBETO, a BETO model that is pre-trained on data derived from social media and news texts. Using this specific model, predictions are made at tweet level and then aggregated through a majority vote to infer author labels.

The third model [16] employs all author's tweets to extract author features. Word and character n-grams, as well as lexical and stylistic features, are used to feed the model, manually selecting them for each of the categories.

The fourth model [17] groups tweets belonging to the same author into clusters containing 8-12 tweets, grouping more information while also accommodating memory constraints. For each category different classification techniques are presented, and manual engineering is applied accordingly. Additionally, a voting system is employed to unify the labels of tweet clusters into user labels.

In summary, three out of the four leading models are based on the Transformer architecture, with the top two specifically relying on monolingual Spanish Transformer-based language models. In addition, tweets authored by the same user tend to be grouped by different techniques, either by concatenating them at the input stage or by merging the associated labels at the output stage.

The exploration of grouping techniques for tweets authored by the same user reflects the ongoing efforts to improve the handling of sequential and contextual information in social media data in Transformer-based classifiers. The decision to concatenate tweets at the input stage or merge associated labels at the output stage implies a deliberate consideration of how to capture and leverage the inherent relationships and dependencies within user-generated content. Nonetheless, in previous works the classifiers were not given both textual and user-based information as input, a shortcoming that our approach tries to address.

3. Datasets

The dataset employed on this work is an expansion of the PoliCorpus 2020 dataset [18] and the corpus utilized for the PoliticES 2022 shared task. It encompasses information extracted from Twitter accounts belonging to politicians, political journalists, and celebrities in Spain. Political accounts were selected among members of the Spanish government, the Congress and Senate of Spain, mayors of important Spanish cities, presidents of the autonomous communities, former politicians, and collaborators affiliated with political parties. Furthermore, journalists were selected from various Spanish news media such as *ABC*, *El País*, *El Diario*, *El Mundo* or *La Razón*, among others.

The objective of creating such dataset is to facilitate techniques for the extraction of demographic characteristics and political ideology from a provided user's collection of tweets. Demographic attributes encompass elements like gender and profession, while political ideology is approached both as a binary and a multiclass problem. Users are annotated by gender (*male* or *female*), profession (*politician, journalist* or *celebrity*), and by political alignment along two axes: a binary scale (*left* or *right*) and a multi-class scale (*left, moderate left, moderate right, right*). To ensure the users privacy, they created clusters of 80 tweets each, with each user-cluster containing tweets from different users that share all the traits under evaluation. In this way user-clusters are used instead of the users themselves with the objective of avoiding to incur in any legal and ethical issues.

In addition to users, the textual content is also anonymized. Thus, tweets that shared content from or mention news websites is filtered. Moreover, any mention to politicians on Twitter is substituted with the token *@user*, while mentions of other Twitter accounts are encoded as *@user*. Finally, References to political parties are also replaced with the token *@political_party*.

The dataset comprises approximately 2800 user-clusters, with each user-cluster containing diverse texts from different dates and topics. The user-clusters from the training and test sets are independent to prevent the possibility of identifying the authors. As shown by the train set quantitative description provided in Table 1, the train set contains 2250 user-clusters (180,000 tweets) whereas the test set includes 547 (43760 tweets).

Category	Class	Tweets	User-clusters	Distribution
Gender	male	119,440	1,493	66.36%
	female	60,560	757	33.64%
Profession	journalist	110,800	1,385	61.56%
	politician	60,160	752	33.42%
	celebrity	9,040	113	5.02%
Political Ideology: binary	left	100,400	1,255	55.78%
	right	79,600	995	44.22%
Political Ideology: multiclass	left	34,400	430	19.11%
	moderate left	66,000	825	36.67%
	moderate right	58,240	728	32.36%
	right	21,360	267	11.86%

Table 1
PoliticES 2023 train set statistics by category and class.

4. Methods

In this section we describe our two techniques to represent textual data at different levels. Firstly, we present a Transformer-based tweet-and-user representation (t&u) method, consisting of representing both levels in the same feature. Secondly, we also introduce a token-level user representation technique which we named word-to-user (w2u).Those representations are then used to feed a text classifier based on each category, but sharing the same unaltered features.

4.1. Tweet-and-user representations

First of all Transformer-based [5] language models are utilized to obtain contextual text-based representations at sentence or tweet level. These models focus on capturing context and meaning by analyzing the relationships among tokens in a text sequence. In contrast to static embedding methods like word2vec[19], which represent words with fixed vector values, Transformer-based models modify those vectors values depending on the surrounding words and their order. It should be noted that this approach is not limited to word-level representations, but rather it can also handle sequences of text, making it suitable for text classification tasks similar to ours. For our experiments, we have selected the following models:

• mBERT [20] is the multilingual version of BERT[20] pre-trained with the largest 104 languages in Wikipedia. Rather than simply predicting the next word in the sequence, the BERT model takes into consideration all of the words in the sequence, thereby developing a more dense and rich representation of the context. BERT pre-trains bidirectional representations from unlabeled text by considering both left and right context in all layers utilising next-sentence prediction and masked-language modeling.

• DistilmBERT [21], the multilingual version of DistilBERT, a smaller and faster Transformer model distilled from BERT. It retains over 95% of BERT's performance on the GLUE benchmark while having 40% fewer parameters and 60% faster inference speed.

• XLM-RoBERTa [22] is a multilingual version of RoBERTa [23], trained with the CC100 corpus, on a large multilingual dataset spanning 100 languages. It is an optimized BERT variant that benefits from training on a dataset ten times larger than BERT, employing dynamic masking, byte-pair encoding tokenization, and omitting the next-sentence prediction objective.

• XLM-T [24] is an extension of the XLM-RoBERTa base model, further trained with 198 million multilingual tweets. This model's focus on Twitter-based data makes it particularly relevant for evaluating performance in tasks specific to this social media platform.

• mDeBERTa [25], the multilingual version of DeBERTa, utilizes the same arquitecture of DeBERTa and, as XLM-RoBERTa, it was trained on the CC100 multilingual dataset, although just for 15 languages. DeBERTa improves BERT and RoBERTa models through disentangled attention and an enhanced mask decoder, outperforming RoBERTa on the majority of natural language understanding (NLU) tasks.

• BETO [12] is a BERT model trained on a Spanish corpus comprising 3 billion tokens. It is similar in size to BERT-base and was trained using the Whole Word Masking technique.

• PolitiBETO [15] is a BERT model specifically tailored for political tasks in social media corpora. It is created through a two-stage domain adaptation process applied to the BETO model, incorporating the language structure found on Twitter and in newspapers.

• MarIA [13] or roberta-large-bne is based on the RoBERTa-large model. It has been pre-trained on a Spanish corpus totaling 570GB of clean and deduplicated text sourced from the National Library of Spain.

We use the listed Transformer models to extract the features and evaluate their performance for the specific task. To represent each tweet, we utilize the last hidden state corresponding to the start-of-sequence token as an aggregate document representation, following the approach described in [20]. This hidden state serves as a tweet vector, which acts as a textual feature representation for the tweet. In order to maintain the same representations for all the categories the Transformers models are used without fine-tuning, meaning that we are using default frozen weights.

In order to extract author representations we average local elements to generate a global representation, as done in previous approaches [26, 27, 28]. In other words, the user representation is obtained by extracting the mean vector of all the tweet vectors authored by each user. Once we get the user representations, each of the tweet representations is concatenated with its author representation vector, generating a tweet-and-user representations for each of the tweets in the dataset.

Finally, the combined tweet-and-user representations for each tweet are used to train a Logistic Regression classifier without any additional tuning. Thus, the same features, without any modification, are used to train a classifier for each of the categories in the PoliticES dataset. After predicting the labels at tweet level, a majority voting strategy is employed to infer the user label by considering the various tweet labels associated with the same author.

4.2. Word-to-user representations

The *word-to-user* model is trained in a unsupervised manner to predict a target user from a given word-token. The input of the model is the raw text data without any preprocessing or modification. In order to obtain our user representations, we use a single hidden-layer neural network. The network is used to train a dense interaction representation model using the tokens from the users' text. The aim of the single hidden-layer feedforward neural network consists of predicting the target user from a given word token that appears in the corpus. The dimensions of the hidden layer determine the size of the final user representation vectors, corresponding to the number of learned features. During training, all the word tokens present in the training and test corpus are used, computing the log probability of correctly predicting the target user from the given word token. The training process is done by sub-sampling the most frequent instances and with negative sampling [19]. A number of experiments were undertaken in order to obtain the optimal dimensionality of the *word-to-user* representations. Once all the users are embedded, the resultant vector is used to represent a given user.

The final step consists of using the extracted word-to-user representation vectors from each user to train a Logistic Regression classifier without any additional tuning. This means that the same features are used across each of the categories in the dataset.

5. Experiments on the Development Data

We experimented with the development data in order to obtain the optimal configuration for the two methods described above. With respect to *tweet-and-user*, the objective is to establish which model is best to extract the contextual representations which are used as features. For the *word-to-user* approach we want to know which dimensionality to represent word-to-user features provides the best results.

5.1. Tweet-and-user

In order to select the best configuration for tweet-and-user representations, we use different feature extraction methods using the language models described in the previous section. The obtained tweet-and-user combined vectors are then used as input to train Logistic Regression classifiers, while employing majority voting. Development results (Table 2) show that the best configuration given by the MarIA model. Therefore, the features extracted from the MarIA model will be used to train the final model.

To conduct an ablation study, we utilize the tweet vectors to train a Logistic Regression classifier independently of the user vectors, while employing majority voting. This tweet-only approach allows us to assess the influence of incorporating tweet and user level information

Category	Language Model mB	dmB	mR	mRl	mRt	mD	Be	pBe	Ma
gender	0.814	**0.817**	0.782	0.800	0.795	0.621	0.768	0.782	0.806
profession	0.725	0.818	0.818	0.599	0.827	0.724	**0.851**	0.701	0.743
ideology binary	0.666	0.719	0.776	0.786	0.761	0.661	0.745	0.749	**0.842**
ideology multiclass	0.582	0.618	0.610	0.593	0.647	0.504	0.588	0.628	**0.678**
average	0.697	0.743	0.747	0.695	0.758	0.628	0.738	0.715	**0.767**

Table 2
Tweet and user level classification. F1 macro score results on development set. Algorithms used to generate the features: mB (mBERT), dmB (DistilmBERT), mR (XLM-RoBERTa-base), mRl (XLM-RoBERTa-large), mRt (XLM-T), mD (mDeBERTa), Be (BETO), pBe (PolitiBETO), Ma (MarIA). Values in bold represent best results for each category.

in extracting valuable social insights. Therefore, the train set is used to train the classifier, while the development set is used to evaluate different algorithms for feature extraction. When comparing the results of tweet-and-user combined approach (Table 2) to the results of the tweet-only approach (Table 3), it can be observed that the latter exhibits a significant loss in performance.

Category	Language Model mB	dmB	mR	mRl	mRt	mD	Be	pBe	Ma
gender	0.473	0.500	0.473	0.439	0.473	0.384	0.518	**0.525**	**0.525**
profession	0.526	0.526	0.526	0.533	0.535	0.503	**0.572**	0.535	0.546
ideology binary	0.675	0.675	0.625	0.660	0.720	0.486	0.670	0.679	**0.774**
ideology multiclass	0.391	0.313	0.294	0.292	0.401	0.263	0.446	0.358	**0.515**
average	0.516	0.504	0.480	0.481	0.532	0.409	0.552	0.524	**0.590**

Table 3
Tweet level classification as ablation test. F1 macro score results on development set. Algorithms used to generate the features: mB (mBERT), dmB (DistilmBERT), mR (XLM-RoBERTa), mRl (XLM-RoBERTa-large), mRt (XLM-T), mD (mDeBERTa), Be (BETO), pBe (PolitiBETO), Ma (MarIA). Values in bold represent best results for each category.

5.2. Word-to-user

In order to select the best configuration for word-to-user representations, we trained different Logistic Regression classifiers with different dimensions. Thus, results on the development dataset (Table 4) show that the best configuration is given by obtaining the word-to-user vector represention in 200 dimensions. It is remarkable the high performance achieved with this configuration, clearly outperforming the Transformer-based tweet-and-user approach on development set.

To have a better understanding of the word-to-user (200 dimensions) user-level representations, we plotted the users' representations and the corresponding class. Figure 1 and 2 show 2 dimensional visualizations obtained by applying a t-SNE dimensionality reduction to word-to-user representations arisen from the development data; each color represents a different class. Demographic characteristics are plotted on Figure 1 while ideological binary and multiclass

Category	Dimensions 10	20	50	100	200	300
gender	0.690	0.789	0.830	0.798	**0.855**	0.830
profession	0.601	0.841	0.862	0.732	**0.925**	0.795
ideology binary	0.818	0.850	0.820	**0.919**	0.910	0.851
ideology multiclass	0.577	0.652	0.678	**0.696**	0.681	0.657
average	0.672	0.783	0.798	0.786	**0.843**	0.783

Table 4
F1 macro score results on development set for w2u trained over different amount of dimensions. Values in bold represent best results for each category.

representations can be seen on Figure 2.

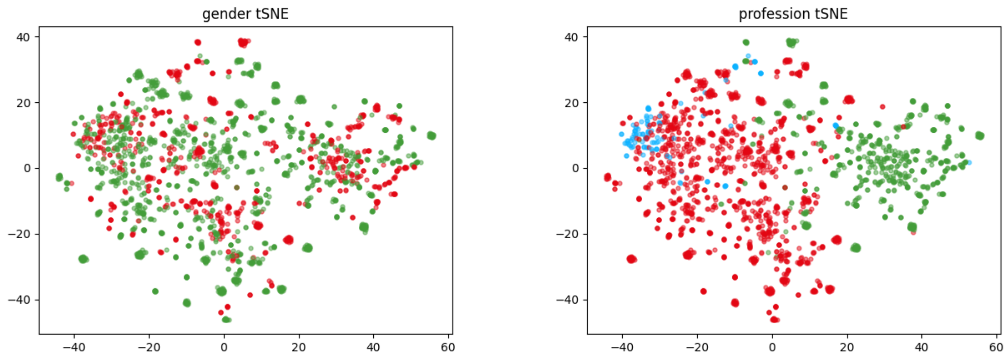

Figure 1: Visualization of t-SNE 2 dimension reduction of w2u (200 dim) user representations for gender (left) and profession (right) Demographic traits on development set.

Regarding Demographic traits, it can be seen that the classes present on profession are clearly defined (Figure 1 right), while the representations of gender (Figure 1 left) seem to be more sparse. With respect to political ideology, the binary framework (Figure 2 left) shows clearer communities than the more sparse multiclass framework (Figure 2 right). Thus, the evaluation results and the visual representations would seem to correlate, as the categories with clearer communities (profession and ideology binary) are also the categories that obtained better classification results on the development data.

6. Results on the Official Test Data

As a result of the experiments performed in the previous section, the *tweet-and-user* method will be using the MarIA model to obtain the combined vectors which are the input to Logistic Regression classifiers for each of the traits (while employing majority voting). In this setting, the classifiers are trained on the training data and evaluated on the official test set. The same procedure is applied for the *word-to-user* method, which will be based on training, for each

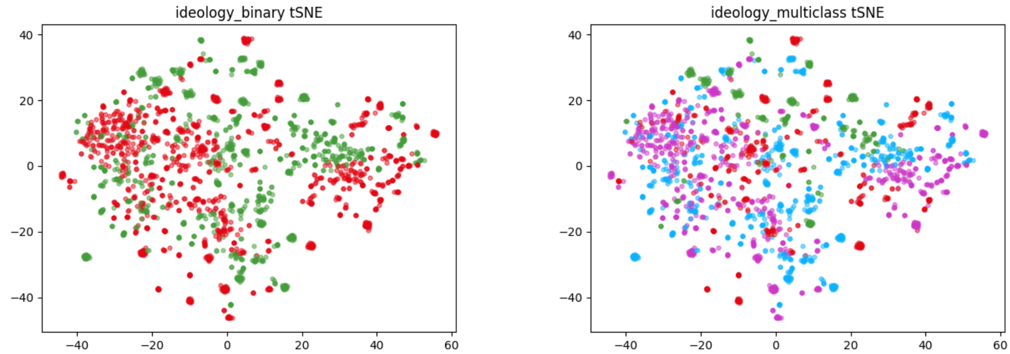

Figure 2: Visualization of t-SNE 2 dimension reduction of w2u (200 dim) user representations for binary (left) and multiclass (right) Political Ideology on development set.

of the traits, Logistic Regression classifiers which take as input features the 200 dimensional *word-to-user vectors*.

	Gen	Prof	Pib	Pim	**AVG**
baseline	66.34	60.24	79.77	54.72	65.27
w2u	75.88	78.67	87.27	63.78	76.40
t&u	73.79	85.48	85.57	59.36	76.05
t&u + w2u	**79.56**	**86.08**	**87.75**	**63.98**	**79.34**

Table 5
F1 macro average test scores for *gender* (Gen), *profession* (Prof), *ideology binary* (Pib), *ideology multiclass* (Pim) categories and overall average (AVG). Algorithms used to generate the features: word-to-user (w2u) 200d, tweet-and-user (t&u) based on MarIA and a combination between both (t&u + w2u).

In reference to the test results presented in Table 5, both the tweet-and-user (t&u) and word-to-user (w2u) approaches exhibit similar average scores, although variations can be observed among the traits. Consequently, we made the decision to merge both feature extraction techniques to assess their combined performance. In fact, it turned out that fusion of tweet-and-user and word-to-user (t&u + w2u) yielded the highest test scores across all categories and the overall average. Moreover, this combination outperformed w2u and t&u individually in all the assessed categories. These findings imply that the combination of features generated by tweet-and-user and word-to-user can result in improved performance when predicting gender, profession, and ideological aspects. Furthermore, the scores obtained provide valuable insights into the effectiveness of the algorithms and their ability to generalize across diverse categories.

7. Conclusion

This paper demonstrates the benefits of combining author and sentence level textual representations for political ideology detection and characterization of users with respect to demographic traits. More specifically, for our participation to the PoliticES 2023 shared task we have experimented with different level Transformer-based textual features as well as with user features directly arisen from word tokens. This combination of features has allowed us to obtain the second-best overall results in the task using a general approach, namely, without performing any specific feature-engineering for any of the traits.

References

[1] J. A. García-Díaz, S. M. Jiménez-Zafra, M. T. Martín-Valdivia, F. García-Sánchez, L. A. Ureña-López, R. Valencia-García, Overview of PoliticEs at IberLEF 2023: Political ideology detection in Spanish texts, Procesamiento del Lenguaje Natural 71 (2023).

[2] S. M. Jiménez-Zafra, F. Rangel, M. Montes-y Gómez, Overview of IberLEF 2023: Natural Language Processing Challenges for Spanish and other Iberian Languages, in: Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2023), co-located with the 39th Conference of the Spanish Society for Natural Language Processing (SEPLN 2023), CEUR-WS.org, 2023.

[3] D. Nguyen, A. S. Doğruöz, C. P. Rosé, F. de Jong, Computational Sociolinguistics: A Survey, Computational Linguistics 42 (2016) 537–593. doi:10.1162/COLI_a_00258.

[4] N. Cesare, C. Grant, E. O. Nsoesie, Detection of user demographics on social media: A review of methods and recommendations for best practices, arXiv preprint arXiv:1702.01807 (2017).

[5] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, in: Advances in Neural Information Processing Systems, 2017, pp. 5998–6008.

[6] J. Fernandez de Landa, R. Agerri, Social analysis of young basque-speaking communities in twitter, Journal of Multilingual and Multicultural Development 0 (2021) 1–15.

[7] M. Abdul-Mageed, C. Zhang, A. Rajendran, A. Elmadany, M. Przystupa, L. Ungar, Sentence-level bert and multi-task learning of age and gender in social media, arXiv preprint arXiv:1911.00637 (2019).

[8] J. A. García-Díaz, S. M. Jiménez-Zafra, M.-T. M. Valdivia, F. García-Sánchez, L. A. Ureña-López, R. Valencia-García, Overview of PoliticEs 2022: Spanish author profiling for political ideology, Procesamiento del Lenguaje Natural 69 (2022) 265–272.

[9] R. Agerri, R. Centeno, M. Espinosa, J. F. de Landa, Álvaro Rodrigo, Vaxxstance@iberlef 2021: Overview of the task on going beyond text in cross-lingual stance detection, Procesamiento del Lenguaje Natural 67 (2021) 173–181.

[10] A. T. Cignarella, M. Lai, C. Bosco, V. Patti, P. Rosso, SardiStance@EVALITA2020: Overview of the Task on Stance Detection in Italian Tweets, in: V. Basile, D. Croce, M. Di Maro, L. C. Passaro (Eds.), Proceedings of the 7th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020), CEUR-WS.org, 2020.

[11] J. Fernandez de Landa, R. Agerri, Relational embeddings for language independent stance detection, arXiv e-prints (2022) arXiv–2210.

[12] J. Cañete, G. Chaperon, R. Fuentes, J.-H. Ho, H. Kang, J. Pérez, Spanish pre-trained bert model and evaluation data, in: PML4DC at ICLR 2020, 2020.

[13] A. G. Fandiño, J. A. Estapé, M. Pàmies, J. L. Palao, J. S. Ocampo, C. P. Carrino, C. A. Oller, C. R. Penagos, A. G. Agirre, M. Villegas, Maria: Spanish language models, Procesamiento del Lenguaje Natural 68 (2022). doi:10.26342/2022-68-3.
[14] S. S. Carrasco, R. C. Rosillo, Loscalis at PoliticEs 2022: Political author profiling using BETO and maria, in: Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2022) co-located with the Conference of the Spanish Society for Natural Language Processing (SEPLN 2022), A Coruña, Spain, September 20, 2022, volume 3202 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022.
[15] E. Villa-Cueva, I. González-Franco, F. Sanchez-Vega, A. P. López-Monroy, NLP-CIMAT at PoliticEs 2022: PolitiBETO, a Domain-Adapted Transformer for Multi-class Political Author Profiling, in: Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2022), CEUR Workshop Proceedings, CEUR-WS, 2022.
[16] A. Mosquera, Alejandro mosquera at PoliticEs 2022: Towards robust spanish author profiling and lessons learned from adversarial attacks, in: Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2022) co-located with the Conference of the Spanish Society for Natural Language Processing (SEPLN 2022), A Coruña, Spain, September 20, 2022, volume 3202 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022.
[17] E. Santibáñez-Cortés, A. Carrillo-Cabrera, Y. A. Castillo-Castillo, D. Moctezuma, V. Muñiz-Sánchez, Cimat_2021 at PoliticEs 2022: Ensemble based classification algorithms for author profiling in spanish language, in: Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2022) co-located with the Conference of the Spanish Society for Natural Language Processing (SEPLN 2022), A Coruña, Spain, September 20, 2022, volume 3202 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022.
[18] J. A. García-Díaz, R. Colomo-Palacios, R. Valencia-García, Psychographic traits identification based on political ideology: An author analysis study on Spanish politicians' tweets posted in 2020, Future Generation Computer Systems 130 (2022) 59–74.
[19] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, arXiv preprint arXiv:1301.3781 (2013).
[20] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, 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, Volume 1 (Long and Short Papers), Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 4171–4186. URL: https://aclanthology.org/N19-1423. doi:10.18653/v1/N19-1423.
[21] V. Sanh, L. Debut, J. Chaumond, T. Wolf, Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter, in: NeurIPS EMC2 Workshop, 2019.
[22] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, V. Stoyanov, Unsupervised cross-lingual representation learning at scale, in: ACL, 2020.
[23] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, V. Stoyanov, RoBERTa: A Robustly Optimized BERT Pretraining Approach, arXiv preprint arXiv:1907.11692 (2019).
[24] F. Barbieri, L. Espinosa-Anke, J. Camacho-Collados, A multilingual language model toolkit for twitter, arXiv preprint arXiv:2104.12250 (2021).

[25] P. He, X. Liu, J. Gao, W. Chen, Deberta: Decoding-enhanced bert with disentangled attention, in: International Conference on Learning Representations, 2021.
[26] I. R. Hallac, S. Makinist, B. Ay, G. Aydin, user2vec: Social media user representation based on distributed document embeddings, in: 2019 International Artificial Intelligence and Data Processing Symposium (IDAP), 2019, pp. 1–5. doi:10.1109/IDAP.2019.8875952.
[27] T. Kenter, A. Borisov, M. de Rijke, Siamese CBOW: Optimizing word embeddings for sentence representations, in: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Berlin, Germany, 2016, pp. 941–951. URL: https://aclanthology.org/P16-1089. doi:10.18653/v1/P16-1089.
[28] Q. Le, T. Mikolov, Distributed representations of sentences and documents, in: International conference on machine learning, PMLR, 2014, pp. 1188–1196.