

Language Independent Stance Detection: Social Interaction-based Embeddings and Large Language Models

Detección de Stance Independiente del Idioma: Representaciones Vectoriales basadas en Interacciones Sociales y Grandes Modelos de Lenguaje

Joseba Fernandez de Landa, Rodrigo Agerri

HiTZ Center - Ixa, University of the Basque Country UPV/EHU
joseba.fernandezdelanda@ehu.eus, rodrigo.agerri@ehu.eus

Abstract: The large majority of the research performed on stance detection has been focused on developing more or less sophisticated text classification systems, even when many benchmarks are based on social network data such as Twitter. This paper aims to take on the stance detection task by placing the emphasis not so much on the text itself but on the interaction data available on social networks. More specifically, we propose a new method to leverage social information such as *friends* and *retweets* by generating Relational Embeddings, namely, dense vector representations of interaction pairs. Our experiments on seven publicly available datasets and four different languages (Basque, Catalan, Italian, and Spanish) show that combining our relational embeddings with discriminative textual methods helps to substantially improve performance, obtaining state-of-the-art results for six out of seven evaluation settings, outperforming strong baselines based on Large Language Models, or other popular interaction-based approaches such as DeepWalk or node2vec.

Keywords: Stance Detection, Multilinguality, Social Networks, Interactions, Large Language Models, Natural Language Processing.

Resumen: La gran mayoría de los trabajos sobre la detección de *stance* (posicionamiento) se han centrado en clasificación de texto, incluso cuando los datos se recolectan de redes sociales como Twitter. Este artículo aborda la tarea de detección de *stance* haciendo énfasis, además de en los datos textuales de los mensajes, en los datos de interacción disponibles en las redes sociales. Proponemos un nuevo método para representar información social como *amigos* y *retuits* generando embeddings relacionales, es decir, representaciones vectoriales densas basadas en pares de interacción. Nuestros experimentos en siete conjuntos de datos públicamente disponibles y para cuatro idiomas (catalán, euskera, español e italiano) demuestran que la combinación de los embeddings relacionales con métodos textuales ayuda a mejorar el rendimiento, obteniendo resultados del estado del arte en seis de los siete escenarios de evaluación, superando otras aproximaciones basadas en grandes modelos de lenguaje u otros enfoques basados en interacciones como DeepWalk o node2vec.

Palabras clave: Detección de Stance, Multilingüismo, Redes Sociales, Interacciones, Modelos de Lenguaje, Procesamiento del Lenguaje Natural.

1 Introduction

Stance detection consists of identifying the viewpoint or attitude expressed by a piece of text with respect to a given target. With the enormous popularity of social networks,

users spontaneously share their opinions on social media, generating a valuable resource to investigate stance. This means that research on stance has a social impact, for example, to help addressing misinformation on

vaccines, or to better understand public opinion about topics such as climate change or migration. Furthermore, stance detection is considered an important intermediate task for fact-checking (Augenstein, 2021) or fake news detection (Pomerleau and Rao, 2017).

The SemEval 2016 task on stance detection in Twitter (Mohammad et al., 2016) presented a dataset with tweets expressing FAVOR, AGAINST and NEUTRAL stances with respect to five different targets, a trend followed by many other researchers (Derczynski et al., 2017; Taulé et al., 2018; Zotova, Agerri, and Rigau, 2021; Hardalov et al., 2022). However, despite many of them using Twitter-based source data, the large majority address the task by considering only the textual content of tweets (Augenstein et al., 2016; Schiller, Daxenberger, and Gurevych, 2021; Hardalov et al., 2021; Li, Zhao, and Caragea, 2021; Ghosh et al., 2019; Küçük and Can, 2020; Sobhani, Inkpen, and Zhu, 2017; Glandt et al., 2021a).

This shortcoming has been addressed by proposing new datasets (Cignarella et al., 2020; Agerri et al., 2021) that include different languages and social interaction data, such as *retweets* or *friends*. Although these new datasets have facilitated the development of new techniques for stance detection considering also interaction data, most of them employ manually engineered features tailored to each specific data type (Espinosa et al., 2020; Lai et al., 2021; Alkhalifa and Zubiaga, 2020), making it difficult to generalize across languages and targets. Recently, significant attention has been directed towards the use of Large Language Models (LLMs) as few-shot learners (Brown et al., 2020). However, the success of in-context learning techniques using LLMs has been mostly limited to English benchmarks such as SemEval 2016 (Taranukhin, Shwartz, and Milios, 2024; Gatto, Sharif, and Preum, 2023; Zhang et al., 2023a), probably because the pre-training of the large majority of publicly available LLMs has been focused mostly on English.

This paper focuses on stance detection of tweets by placing emphasis on the interaction data commonly available in social media. More specifically, we propose a new method to leverage social information such as friends and retweets by generating Relational Embeddings, namely, dense vector representa-

tions of interaction pairs. The development of our new method allows us to make the following contributions to language independent stance detection: (i) a new method to represent and exploit interaction data, such as *friends* and/or *retweets*, by generating relational embeddings based on one-to-one relations; (ii) comprehensive experiments on seven publicly available datasets and four different languages different to English show that our relational embeddings behave robustly across different targets and languages without any specific manual engineering; (iii) combining our method with text-based discriminative classifiers helps to systematically improve their results, outperforming also ensembles of pre-trained language models (Giorgioni et al., 2020) or strong in-context learning baselines using Large Language Models (Taranukhin, Shwartz, and Milios, 2024); (iv) we empirically demonstrate that our new Relational Embeddings clearly outperform popular graph-based approaches to encode interaction data, such as DeepWalk or node2vec; (v) exhaustive ablation and error analyses show that the method used to obtain the *retweet* data and the size of the users community is crucial for state-of-the-art performance using our technique; (vi) the new generated datasets with interaction data and code are publicly available¹.

Finally, while this paper is focused on stance detection, we believe that our Relational Embeddings can be successfully applied to a large number of Computational Social Science and NLP tasks based on social media, especially those related to political ideology, misinformation, and hate speech, but also for health-related applications such as the detection of early signs of epidemic outbreaks (Martín-Corral et al., 2022).

2 Related work

Recent studies have demonstrated that using LLMs on Stance Detection tasks can provide significant performance increases (Zhang et al., 2023b; Zhang et al., 2023a). Furthermore, combining the application of LLMs with Chain-of-Thought (CoT) prompting (Wei et al., 2022), and in-context learning in which the model generates intermediate reasoning steps to arrive at a final prediction,

¹https://github.com/joseba-fdl/relational_embeddings/

has also helped to substantially improve results (Kojima et al., 2022; Wang et al., 2023; Gatto, Sharif, and Preum, 2023).

Despite their high capabilities, the application of LLMs still faces several challenges, such as dealing with cases of implicit stance or avoiding hallucinations, even when employing advanced prompting strategies such as CoT reasoning (Gatto, Sharif, and Preum, 2023). To address these limitations, Stance Reasoner (Taranukhin, Shwartz, and Milios, 2024) improves the CoT method by including examples and reasoning as background knowledge to achieve generalizable predictions across different targets. However, these approaches are only focused on English.

Additionally, most stance detection research and datasets released do not include interaction data, despite being collected from social media sources such as Twitter. Küçük and Can (2020) lists stance-annotated datasets for 11 languages, whereas recent work on cross-domain and cross-lingual stance provide experimentation for 16 datasets and 15 languages (Hardalov et al., 2021; Hardalov et al., 2022). The focus, however, remains on the textual content of the tweets. This trend has recently changed with the release of, to the best of our knowledge, two datasets which, in addition to the stance labeled tweets, include interaction data such as *retweets* and *friends*: SardiStance (Cignarella et al., 2020) and VaxxStance (Agerri et al., 2021).

The winner (Espinosa et al., 2020) of the SardiStance shared task (Cignarella et al., 2020) used a weighted voting ensemble that combined two inputs: (a) psychological, sentiment and *friends* distances as features used to learn an XGBoost (Friedman, 2001) model, with (b) text classifiers based on the Transformer architecture (Devlin et al., 2019). Other systems combined textual data (emoticons, special characters, and word embeddings) with 2 dimensions extracted from the interactions distance matrix using Multidimensional Scaling (MDS) (Ferraccioli et al., 2020), or friendship-based graphs created with DeepWalk (Perozzi, Al-Rfou, and Skiena, 2014) and various types of textual embeddings (Alkhalifa and Zubiaga, 2020).

The VaxxStance shared task (Agerri et al., 2021) provided textual and interaction data (*friends* and *retweets*) to study stance detec-

tion on vaccines in Basque and Spanish. The one system that systematically outperformed the baselines (Lai et al., 2021) manually engineered a large number of features based on stylistic, tweet, and user data, lexicons, dependency parsing, and network information, which were specifically developed for these datasets and languages.

The most recent approaches tackling unsupervised stance detection using social media interactions as features use the force-directed algorithm (Fruchterman and Reinhold, 1991) or UMAP (McInnes et al., 2018). These algorithms transform interaction frequency vectors into features, reducing huge interaction matrices into low-dimensional features. Darwish et al. (2020) use both the force-directed algorithm and UMAP for unsupervised stance detection of Twitter users. UMAP is also used to get interaction-based features for automatically tagging Twitter users’ stance (Stefanov et al., 2020) and to explore political polarization in Turkey (Rashed et al., 2021).

Other works are based on node2vec (Grover and Leskovec, 2016) for user profiling and extracting user features for abuse detection (Mishra et al., 2018) and also for sentiment, stance and hate speech detection (Del Tredici et al., 2019). Commonly used algorithms for building interaction-based models like DeepWalk (Perozzi, Al-Rfou, and Skiena, 2014) and node2vec are based on generating Random Walks. However, those randomly generated walks create artificial interactions that may not occur in the gathered interaction pairs. Furthermore, selecting the structure of the random walks and deciding the number of context users to be predicted needs to be manually modeled and adapted.

In contrast to previous work based on in-context learning with LLMs, supervised text classification or interaction-based methods such as DeepWalk or node2vec, our Relation Embeddings method provides dense interaction-based representations of users, focusing on real interaction pairs. The training process is designed to predict a target user receiving a *retweet* or a *follow* from a source user, each instance an item-to-item prediction instead of context-to-item (CBOW) or item-to-context (Skip-gram) prediction. Additionally, we focus on all the interaction pairs, without generating artificial random interactions to train the model or manually

selecting the most salient users.

3 Method

We proposed a new method to generate vector-based representations of interactions in social networks, such as *friends* and *retweets*. These new representations, which we refer to as *Relational Embeddings* (RE), are then leveraged to propose two methods to perform stance detection: (i) building classifiers using just our relational embeddings (§3.2) and, (ii) combining RE with various classifiers based on textual data (§3.3).

3.1 Relational Embeddings

In this paper, the type of interactions used is *retweets* and *friends*, which are seen as relations between two users, one generating the action (source) and the other receiving it (target). Thus, the actions of *retweeting* or *following* other users are considered interaction pairs. Generally, these interactions should help to reveal users' preferences by capturing meaningful information from their performative actions.

The first step in our method consists of gathering the interactions from the users included in the labeled data, namely, the one-to-one *retweet* and *follow* actions between the users/authors of the tweets. It should be noted that a set of *retweet* and *follow* interactions can consist of independent one-to-one actions without direct relation between them. This is why in our model we consider each interaction pair as a single instance without any preprocessing or modification.

Using this interaction data, our model is then trained in an unsupervised manner to predict, in each instance, a target user from a given source user. Note that the instances used as input are real interaction pairs, namely, they do not correspond to sparse interaction frequency matrices or neighbors arising from interaction networks or without generating artificial ones as random walks do.

In order to obtain our relational representations, we use a single hidden-layer neural network (Figure 1). The network is used to train a dense interaction representation model using the *friends* and/or *retweet* based data. Each user is encoded as a one-hot vector of size U , where U is the number of users among interaction pairs (I) in a specific dataset. Given a one-hot vector U , the aim




Figure 1: One hidden layer artificial neural network.

of the single hidden-layer feedforward neural network consists of predicting the target user. The dimensions of the hidden layer (D) determine the size of relational vectors representing the target user, which correspond to the number of learned features. During training, the weights W and W' are modified to minimize the loss function due to back-propagation. According to Equation 1, the summation goes over all the interaction pairs (I) in the training corpus, computing the log probability of correctly predicting the target user (u_{target}) from the source user (u_{source}) for each interaction (i). The training process is done by sub-sampling the most frequent instances and with negative sampling (Mikolov et al., 2013). Finally, the W matrix is used to retrieve the interaction vectors representing each user, generating the relational embedding, from which the relation vector for each user is obtained. In this model, users with similar interactions should have similar representations, turning many interaction pairs into dense relational representations of D dimensions.

$$\frac{1}{I} \sum_{i=1}^I \log p(u_{target}|u_{source}) \quad (1)$$

3.2 Interaction-based Classifier with Relational Embeddings

Our first system consists of a linear classifier taking as input only the relational embeddings described in the previous section. Building such a system will allow us to understand the performance of the generated Relational Embedding models on their own.

Each of the tweets from a dataset will be represented by its author's (user) relational vector, which represents the interactions of its author. By doing so, we effectively project the relations of the author to the tweet level, generating a link between the relational data and the stance labels. In this step, some users

may be repeated across the data, but their assigned stance label will be that of the corresponding tweet. It should be noted that, although possible, it is quite uncommon to have a user with tweets labeled differently across the data. Thus, each tweet is converted into relational vectors, represented by the specific user’s vector weights in the relational embedding model. Those users not present in the model are represented as vectors of zeros. This is usually due to the inability to retrieve user interaction data, either because the user has disappeared from Twitter or because their profiles are kept private. As shown in Figure 2, the final relational vectors for each tweet are used to train an SVM (RBF kernel) classifier without any additional input.




Figure 2: Relational Embeddings + SVM model architecture.

3.3 Combining Textual and Interaction Data

In order to combine textual and interaction data, we use both the texts conveyed by a given user and its associated social media interactions as data input. More specifically, we propose an ensemble method combining textual and interaction-based features before learning. This ensemble method is applied to combine interaction and textual representations with SVM and Transformers.

In Figure 3, we obtain a FastText (*FTEmb*) dense or TFIDF sparse word vectors to represent each tweet. The interaction embedding of each author will then be concatenated to the textual vector, adding a vector of zeros to the textual vector if no relational information is available. Finally, the concatenation of the textual and relational vectors is used to learn an SVM (RBF kernel) model.

As shown in Figure 4, Transformer models and interaction embeddings are combined by concatenating user vectors from the interaction embeddings with the Transformer’s CLS representations of the tweets. When there is no user information related to interactions, a vector of zeros is concatenated to the CLS




Figure 3: SVM-based combined models architecture.

representation. Finally, we add a linear classification layer on top of the CLS token vector concatenated with the relational vector and fine-tune the model in an end-to-end manner.




Figure 4: Transformer-based combined model architecture.

4 Stance Detection Datasets

In order to be able to experiment with the interaction data, the datasets should include, in addition to the labeled textual data, interactions of the users that published the tweets, such as each user’s *friends* and *retweets*. To the best of our knowledge, there are currently only two publicly available datasets with such contextual information, namely, SardiStance (Cignarella et al., 2020) and VaxxStance (Agerri et al., 2021).

We would like to note that we tried to obtain user information for SemEval 2016 (Mohammad et al., 2016) and other English datasets (Lai et al., 2020; Conforti et al., 2020; Glandt et al., 2021b), without much success. In the case of SemEval 2016, perhaps the most popular stance detection dataset, we only managed to retrieve less than 30% of users, not nearly enough for meaningful experimentation with interaction data. For other datasets, it was simply not possible to extract the tweets based on the IDs published. However, we did manage to extract the interactions (over 80%) for the Catalonia Independence Corpus (Zotova, Agerri, and Rigau, 2021) (CIC).

The number of labeled tweets and their distribution between train and test sets can be seen in Table 1. This choice of data offers a varied relation user-tweets (very low in SardiStance, quite high in CIC), which would also allow us to test the robustness of our method.

	Tweets			Interactions		
	Train	Test	Total	Users	RTs	Friend
SE2016	2,914	1,249	4,163	-	-	-
C-ca	8,038	2,010	10,048	691		
C-ca*	8,056	1,992	10,048	691	10M [†]	24M [†]
C-es	8,036	2,011	10,047	334		
C-es*	8,016	2,031	10,047	334		
S	2,132	1,110	3,242	2,827	575K	3M
V-eu	1,072	312	1,384	210	190K	170K
V-es	2,003	694	2,697	1,675	9K 552K [†]	2.1M

Table 1: Datasets: SE2016 (SemEval 2016), C (CIC), S (SardiStance), V (VaxxStance); * means no overlap of users across train and test; RT (retweets). [†] mark represents supplementary interaction-based data added by us.

4.1 SemEval 2016

The SemEval2016 dataset (Mohammad et al., 2016), consists of English tweets labeled for stance (AGAINST, FAVOR, and NONE) from an initial pool of 2M collected tweets. In the supervised track, 4163 tweets are provided with manual annotation for five targets: “Atheism”, “Climate Change is a Real Concern”, “Feminist Movement”, “Hillary Clinton”, and “Legalization of Abortion”. The annotated tweets were ordered by their timestamps. Interestingly, in addition to stance, annotations are provided to express whether the target is explicitly mentioned in the tweet. Finally, no user-based information is included in the datasets, just the tweets.

4.2 SardiStance

This dataset contains tweets in Italian about the Sardines movement (Cignarella et al., 2020). In addition to the textual data, this dataset also provides social and user information, such as the authors’ friends and the retweets. We were unable to extract any supplementary data because both tweet and user identifiers are encrypted.

4.3 VaxxStance

The dataset was independently collected for two languages: Basque and Spanish. No user overlap across train and test sets occurs in the data. Interactions, such as *friends* from the users and *retweets* made to the labeled tweets, are also included. The Basque version (VaxxStance-eu) also includes retweets retrieved from the users’ timelines, as there are few tweets retrieved from the labeled tweets. In order to get more interaction-based data, we extracted 552K supplemen-

tary retweets from the users’ timelines of the Spanish subset (VaxxStance-es), emulating the extra retweet collection as the authors did for the Basque version.

4.4 Catalonia Independence Corpus

The Catalonia Independence Corpus includes coetaneous tweets in both Spanish and Catalan (Zotova et al., 2020), is multilingual, quite large (10K tweets), and reasonably balanced. In the original CIC data, 92.50% of the users in the Catalan set occur also in the test set, whereas for Spanish the proportion is even higher, namely, 99.72%. In order to avoid any possible overfit to the author’s style, a second version of the dataset (Zotova, Agerri, and Rigau, 2021) distributes the tweets in such a way that their authors do not appear across the training, development, and test sets (CIC*).

5 Baselines

We compare state-of-the-art in-context learning with LLMs methods with supervised text-based statistical and Transformer encoder-only classifiers. Furthermore, as explained in Section 3, supervised textual methods are combined with interaction-based features extracted from the Twitter social network. Thus, in addition to our Relational Embeddings, we also experiment with interaction-based methods such as node2vec and DeepWalk for direct comparison with our proposal.

In-context learning with LLMs While traditional supervised text-based methods fine-tune a language model on some training data, in-context learning is based on proving the language model with optimized prompts (Brown et al., 2020). The benefits are two-fold: no training is required, just a handful of examples (few-shot), and by guiding the language model to learn in-context it facilitates the generalization capability of the model to unseen targets. Taking into account their multilinguality claims, we prove the LLMs Mistral 7B v0.2 instruct (Jiang et al., 2023) and Llama3 8B instruct (Dubey et al., 2024) with two different types of prompts plus examples and their translations into the specific language of the dataset (full examples provided in Annex B).

- **Few-Shot** A prompt including 6 in-context examples, two for each label, on targets dif-

ferent to the ones included in the evaluation set. The examples are carefully crafted to include features that would maximize the generalization of the models (Taranukhin, Shwartz, and Milios, 2024).

- **Stance Reasoner** The few-shot prompt is augmented with chain-of-thought (Wei et al., 2022) to generate intermediate reasoning steps that lead to a label prediction (Taranukhin, Shwartz, and Milios, 2024).

Supervised text-based methods In order to combine interaction-based with textual-only supervised approaches, we choose three commonly used text classification methods for stance detection (AlDayel and Magdy, 2021; Küçük and Can, 2020; Hardalov et al., 2021; Zotova, Agerri, and Rigau, 2021) as baselines to compare with our relational embeddings models:

- **Word Embeddings:** We use FastText CommonCrawl models trained using the C-BOW architecture and 300 dimensions on a vocabulary of 2M words (Grave et al., 2018). For classification, each tweet is represented as the average of its word vectors (Kenter, Borisov, and de Rijke, 2016) which is then used to train an SVM (RBF kernel) classifier.

- **TFIDF:** TFIDF (Term Frequency Inverse Document Frequency) vectorization is applied in order to reduce word vector dimensionality by lowering the impact of words that occur too frequently in the selected corpus. TFIDF vectorizing is applied over the text of the tweets, selecting most salient features and reducing sparsity. The obtained TFIDF vectors are used to learn a SVM (RBF kernel) model.

- **XLM-RoBERTa:** We use XLM-RoBERTa (Conneau et al., 2020) for text classification, a masked encoder-only language model pre-trained for 100 languages on 2.5 TB of CommonCrawl text. This model has been widely tested for multilingual stance detection with state-of-the-art performance (AlDayel and Magdy, 2021; Ghosh et al., 2019; Küçük and Can, 2020; Espinosa et al., 2020; Zotova, Agerri, and Rigau, 2021).

Interaction-based Methods In addition to the Relational Embeddings proposed in Section 3.1, we also compare with two commonly used graph-based approaches to extract user representations of the tweets' au-

thors which are then used to train an SVM (RBF kernel) classifier, as described in Section 3.2.

- **DeepWalk** (Perozzi, Al-Rfou, and Skiena, 2014): Given a node(s) in the network, this algorithm learns feature representations to predict their context or neighbors. In this item-to-context (Skip-gram) predicting task, the neighbors to predict may be artificially generated by simulating random walks among the connected nodes.

- **Node2vec** (Grover and Leskovec, 2016): Similar to DeepWalk, it adds two parameters to control the structure of the network during the generated random walks. The control parameters focus on the probability of revisiting points and on the probability of visiting further points.

6 Experiments

The experimental setup consists of (i) evaluating the textual baselines on SemEval 2016 and, (ii) experimenting with interaction data on seven datasets including three different topics (Independence of Catalonia, anti-vaxxers, Sardines movement) and four languages (Basque, Catalan, Italian and Spanish). This choice of datasets will allow us to first evaluate, using SemEval 2016, popular supervised textual classifier methods based on statistical and Transformers encoder-only models with respect to state-of-the-art in-context learning techniques with Large Language Models (LLMs). Once we have established the performance of each textual classifier for English in SemEval 2016, we will then combine the textual supervised methods with interaction data (as explained in Section 3.3) and compare them with in-context learning with LLMs in CIC, VaxxStance and SardiStance.

6.1 Experimental Settings

Regarding which interactions to choose to complement textual classifiers, for each dataset (CIC, SardiStance, and VaxxStance) three different types of interaction-based embeddings were trained, based on the data source: (i) *retweets*, (ii) *friends* and, (iii) *mixed* embeddings.

Retweets are used to share specific content from other users' publications. The reiteration of these actions may demonstrate attachment to a user or its content, actively showing the specific preferences of the source user.

Model	LA	AT	CC	FM	HC	Overall
<i>Supervised</i>						
TFIDF	62.2	53.8	41.6	55.5	58.1	62.9
FTEmb	60.4	59.8	40.0	53.8	47.7	65.8
XLM-R	<u>63.7</u>	72.4	<u>45.5</u>	<u>58.3</u>	<u>73.9</u>	<u>72.5</u>
<i>Few-shot with LLMs</i>						
Few-Shot						
Mistral 7B	66.7	48.9	<u>65.4</u>	64.7	76.1	70.7
LLaMA3 8B	64.3	56.1	57.9	68.6	78.3	70.8
Stance Reasoner						
Mistral 7B	72.3	<u>72.0</u>	60.0	74.0	76.1	75.4
LLaMA3 8B	63.2	69.3	64.9	68.0	76.8	73.7
Previous SOTA	74.3	44.6	65.7	65.0	71.3	75.1

Table 2: Overall F_1 macro results on the SemEval 2016 task 6a dataset (Mohammad et al., 2016) and for each target in the test split, namely, LA - *Legalization of Abortion*, AT - *Atheism*, CC - *Climate Change is a Real Concern*, FM - *Feminist Movement*, HC - *Hillary Clinton*. The best results are highlighted in bold. The best results for each group in underlined. Few-shot with LLMs: our implementation with prompts from Taranukhin et al. (2024).

Furthermore, although retweet actions are more likely to encode latent information related to community or polarization (Conover et al., 2011; Zubiaga et al., 2019), we also wanted to include *friends* related data, which is a result of a *following* action. This passive action allows the source user to be aware of what is being said without sharing or promoting any content. Finally, we also combine both *retweets* and *friends* in a *mixed* representation to test whether merging passive and active interaction types in the same interaction space helps to embed social information.

The best interaction type to build the embeddings for each dataset was chosen by evaluating them with the classifiers built with interaction-based representations (RE, DW and N2V) via 5-fold cross-validation on the training data. The results showed that RE *retweet* was the best interaction data for CIC and VaxxStance-eu, whereas the *mixed* embeddings were the best for SardiStance and VaxxStance-es. With respect to DW and N2V, *retweets* were best for CIC and VaxxStance, while *mixed* performed better for SardiStance.

The procedure to choose the rest of the hyperparameters (dimensionality of the interaction embeddings etc.) for every method is described in Appendix A. Finally, as it is customary for this task, despite training and predictions being done for the 3 classes, evaluation is performed by calculating the averaged F1-score over the AGAINST and FA-

VOR classes (Mohammad et al., 2016).

6.2 Evaluation Results

The evaluation of stance detection systems on the SemEval 2016 dataset (Table 2) highlights the effectiveness of few-shot approaches leveraging large language models (LLMs). In particular, the Stance Reasoner configuration with Mistral and LLaMA3 consistently surpasses the fine-tuned XLM-RoBERTa’s performance, outperforming also the previous state-of-the-art (SOTA) results obtained by Stance Reasoner (Taranukhin, Shwartz, and Milios, 2024). These results demonstrate the strength of in-context learning for English using modern LLMs.

Moving on to the experiments with datasets with available interaction data, in Table 3 we report the results of evaluating stance detection systems on CIC (Catalan and Spanish) SardiStance (Italian) and VaxxStance (Basque and Spanish). If we look first at the text-only approaches we observe that, unlike for SemEval 2016, supervised approaches clearly outperform few-shot in-context learning with LLMs methods. We hypothesize that this is due to the English-centric development of LLMs. The ‘translated’ rows refer to the results obtained when the tweets included in the few-shot configuration are translated to the target languages. These results show that while LLMs benefit from having the examples included in the prompt in the target languages, the gains are much lower with the Stance Reasoner prompt, in which the CoT reasoning compo-

Model	C-ca	C-ca*	C-es	C-es*	S	V-es	V-eu	avg.
<i>Few-shot with LLMs</i>								
Few-shot								
Mistral 7B	20.1	22.0	28.8	28.5	48.1	63.8	34.1	35.1
Mistral 7B (translated)	28.8	32.9	32.4	33.5	56.3	64.0	32.7	40.1
LLaMA3 8B	49.9	56.9	48.8	51.2	59.0	50.6	<u>47.5</u>	52.0
LLaMA3 8B (translated)	<u>55.4</u>	<u>59.7</u>	<u>49.5</u>	<u>52.0</u>	<u>63.1</u>	54.1	43.3	<u>53.9</u>
Stance Reasoner								
Mistral 7B	47.4	46.6	44.2	45.8	50.9	<u>65.0</u>	41.6	48.8
Mistral 7B (translated)	49.0	51.1	43.0	44.5	55.4	64.0	41.2	49.7
LLaMA3 8B	<u>50.6</u>	<u>55.7</u>	<u>51.6</u>	<u>52.7</u>	60.4	59.1	<u>46.8</u>	<u>53.9</u>
LLaMA3 8B (translated)	49.6	53.2	48.9	50.7	<u>60.5</u>	59.9	44.0	52.4
<i>Supervised</i>								
Text-based								
TFIDF	75.3	71.6	73.7	73.1	<u>63.4</u>	76.5	<u>54.4</u>	<u>69.7</u>
FTEmb	61.6	62.6	57.4	59.8	56.1	71.3	47.7	59.5
XLM-R	<u>77.6</u>	<u>74.6</u>	<u>74.2</u>	<u>73.9</u>	57.2	<u>82.5</u>	41.2	68.7
Interaction-based								
RE (ours)	<u>82.2</u>	<u>70.2</u>	<u>85.2</u>	<u>84.4</u>	<u>71.7</u>	<u>85.5</u>	<u>48.4</u>	<u>75.4</u>
N2V	69.7	61.8	65.9	49.3	65.7	76.2	29.0	59.7
DW	69.1	68.0	66.5	64.0	66.4	79.5	25.7	62.7
Combined								
RE + FTEmb	<u>82.2</u>	69.2	88.6	<u>87.7</u>	74.0	89.1	73.2	80.6
RE + TFIDF	<u>82.2</u>	<u>80.2</u>	<u>92.5</u>	86.6	<u>74.6</u>	<u>90.2</u>	<u>75.3</u>	<u>83.1</u>
RE + XLM-R	78.8	75.9	76.8	77.2	60.2	81.5	51.8	71.7
N2V + FTEmb	71.7	63.7	71.3	57.7	70.3	80.4	48.4	66.2
N2V + TFIDF	77.6	74.1	80.1	72.6	70.9	85.8	54.1	73.6
N2V + XLM-R	77.2	73.7	72.9	74.6	56.3	80.6	39.2	67.8
DW + FTEmb	72.9	67.3	71.6	72.9	67.7	82.7	48.5	69.1
DW + TFIDF	79.2	75.9	81.1	79.7	70.1	86.0	54.7	75.2
DW + XLM-R	78.0	75.1	73.5	72.8	55.6	78.7	47.3	68.7
Previous SOTA	74.7	74.9	74.7	71.8	74.4	89.1	77.7	76.7

Table 3: F_1 macro scores (Mohammad et al., 2016). Best results highlighted in bold; best results per group underlined. Few-shot with LLMs: our implementation using prompts from Taranukhin et al. (2024) with Mistral v0.2 7B and Llama 3 8B, both instruct. Supervised: using interaction-based systems (RE, N2V, and DW), text-based systems (FTEmb, TF-IDF and XLM-RoBERTa) and their combinations. Previous SOTA: CIC (Zotova et al., 2020; Zotova, Agerri, and Rigau, 2021), SardiStance (Espinosa et al., 2020) and VaxxStance (Lai et al., 2021).

nent plays a crucial role.

The results using interaction-based approaches underline the usefulness of encoding social interactions for stance detection. This is particularly true of the Relational Embeddings, which outperform every text-based method using only relational vectors and SVM as a learning algorithm. Relational Embeddings also outperform interaction-based alternatives like node2vec and DeepWalk, emphasizing the value of modeling interactions as relational pairs rather than arbitrary random walks.

Most importantly, combining text-based supervised methods with our Relational Embeddings systematically helps to improve results for every language and dataset, obtaining in most cases state-of-the-art results. Thus, for SardiStance and VaxxStance, Rela-

tional Embeddings achieve performance comparable to SOTA without manual tuning or extensive additional resources, in contrast to previous systems that relied heavily on manually engineered features and external data (Lai et al., 2021). Finally, the combination of Relational Embeddings with TFIDF (RE+TFIDF) consistently shows the best performance, even outperforming systems leveraging pre-trained language models and state-of-the-art LLMs. These findings demonstrate the potential of combining supervised textual and relational features for robust and language-independent stance detection.

7 Discussion

In order to have a clearer understanding of the results obtained by the different

interaction-based methods, we visualize their generated representations via PCA reduction. The reported results have shown that Relational Embeddings, which use actual interaction pairs rather than random walk methods like N2V and DW, lead to superior results. Visualizations of user representations, reduced to 2D with PCA, reveal a clear link between the readability of relational embeddings and system performance (Figure 5).

CIC data shows distinct FAVOR and AGAINST communities, while SardiStance data displays more overlap. VaxxStance-eu embeddings struggle to separate stances, likely due to the small, highly interactive Basque Twitter community. Topics with strong political homophily, such as Catalan independence, generate clearer embeddings than less polarized or smaller datasets like VaxxStance.

Comparisons in Appendix A show that DW and N2V embeddings produce less coherent visualizations, reflecting their lower performance.

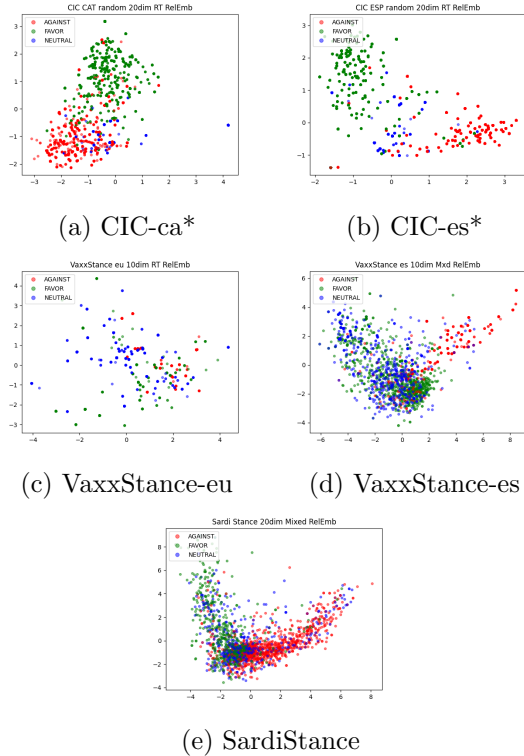


Figure 5: Relational embedding representation of training set users (PCA dimension reduction to 2).

8 Conclusion and Future Work

Stance Detection is a challenging task as the selection of the target significantly influences

the assignment of labels. For instance, the same piece of text may express support for one target while simultaneously opposing another, depending on the context and framing. This dependency makes it critical to perform the task in a generalizable way, as even a small change may affect the outcome. LLMs with in-context learning offer a promising approach to achieving generalizable Stance Detection. However, they face notable limitations, including a substantial error rate (approximately 25%) and significant difficulties in handling non-English languages.

In this paper, we propose Relational Embeddings, a new method to represent interaction data such as *retweets* and *friends*. Relational Embeddings help to reduce the sparsity of interaction data by behaving like dense graphs, being able to embed information related to stance from different data sources without any manual engineering. While this technique is language-independent and fast to train and apply, our results demonstrate that Relational Embeddings behave robustly across different datasets, targets, and languages, substantially and consistently improving results by combining them with text-based supervised methods. The results show that using Relational Embeddings also outperforms most text classification baselines. Furthermore, a direct comparison with previous interaction-based approaches such as DeepWalk and node2vec shows the superiority of our approach.

The results and analysis performed show that we should pay more attention to social network data, aiming to address the shortcomings discussed by further researching different strategies to leverage such interaction data. Future work may include analyzing the Relational Embeddings performance on zero-shot and cross-lingual settings, moving on towards a method that, by using user-based Relational Embeddings, helps to drastically reduce the need of annotated data at tweet level.

Acknowledgements

We are thankful to the following MCIN/AEI/10.13039/501100011033 projects: (i) DeepKnowledge (PID2021-127777OB-C21) and by FEDER, EU; (ii) Disargue (TED2021-130810B-C21) and European Union NextGenerationEU/PRTR; (iii) DeepMinor (CNS2023-144375) and

European Union NextGenerationEU/PRTR.

References

- Agerri, R., R. Centeno, M. Espinosa, J. F. de Landa, and Álvaro Rodrigo. 2021. Vaxxstance@iberlef 2021: Overview of the task on going beyond text in cross-lingual stance detection. *Procesamiento del Lenguaje Natural*, 67:173–181.
- AlDayel, A. and W. Magdy. 2021. Stance detection on social media: State of the art and trends. *Information Processing & Management*, 58(4):102597.
- Alkhalifa, R. and A. Zubiaga. 2020. QMUL-SDS@SardiStance: Leveraging Network Inter-actions to Boost Performance on Stance Detection using Knowledge Graphs. In *Proceedings of the 7th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020)*. CEUR Workshop Proceedings.
- Augenstein, I. 2021. Towards explainable fact checking. *ArXiv*, abs/2108.10274.
- Augenstein, I., T. Rocktäschel, A. Vlachos, and K. Bontcheva. 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, November. Association for Computational Linguistics.
- Brown, T. B., B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. teusz Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901.
- Cignarella, A. T., M. Lai, C. Bosco, V. Patti, and P. Rosso. 2020. SardiStance@EVALITA2020: Overview of the Task on Stance Detection in Italian Tweets. In V. Basile, D. Croce, M. Di Maro, and L. C. Passaro, editors, *Proceedings of the 7th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020)*. CEUR-WS.org.
- Conforti, C., J. Berndt, M. T. Pilehvar, C. Giannitsarou, F. Toxvaerd, and N. Collier. 2020. Will-they-won’t-they: A very large dataset for stance detection on twitter. In *ACL*.
- Conneau, A., K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov. 2020. Unsupervised Cross-lingual Representation Learning at Scale. In D. Jurafsky, J. Chai, N. Schluter, and J. Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451.
- Conover, M. D., J. Ratkiewicz, M. Francisco, B. Gonçalves, F. Menczer, and A. Flammini. 2011. Political Polarization on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*.
- Darwish, K., P. Stefanov, M. Aupetit, and P. Nakov. 2020. Unsupervised user stance detection on twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, pages 141–152.
- Del Tredici, M., D. Marcheggiani, S. Schulte im Walde, and R. Fernández. 2019. You shall know a user by the company it keeps: Dynamic representations for social media users in NLP. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4707–4717, Hong Kong, China, November. Association for Computational Linguistics.
- Derczynski, L., K. Bontcheva, M. Liakata, R. Procter, G. W. S. Hoi, and A. Zubiaga. 2017. SemEval-2017 task 8: RumourEval: Determining rumour veracity and support for rumours. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 69–76.
- Devlin, J., M. Chang, K. Lee, and K. 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.
- Dubey, A., A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan, et al. 2024. The Llama 3 Herd of Models. *ArXiv*, abs/2407.21783.
- Espinosa, M. S., R. Agerri, A. Rodrigo, and R. Centeno. 2020. DeepReading@SardiStance: Combining Textual, Social and Emotional Features. In *Proceedings of the 7th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020)*. CEUR Workshop Proceedings.
- Ferraccioli, F., A. Sciandra, M. D. Pont, P. Girardi, D. Solari, and L. Finos. 2020. TextWiller@SardiStance, HaSpeede2: Text or Context? A smart use of social network data in predicting polarization. In *Proceedings of the 7th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020)*. CEUR Workshop Proceedings.
- Friedman, J. H. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232.
- Fruchterman, T. M. J. and E. M. Reingold. 1991. Graph drawing by force-directed placement. *Software: Practice and Experience*, 21.
- Gatto, J., O. Sharif, and S. Preum. 2023. Chain-of-thought embeddings for stance detection on social media. In H. Bouamor, J. Pino, and K. Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4154–4161, Singapore, December. Association for Computational Linguistics.
- Ghosh, S., P. Singhanian, S. Singh, K. Rudra, and S. Ghosh. 2019. Stance detection in web and social media: a comparative study. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 75–87. Springer.
- Giorgioni, S., M. Politi, S. Salman, R. Basili, and D. Croce. 2020. Unitor @ sardistance2020: Combining transformer-based architectures and transfer learning for robust stance detection. In *EVALITA*.
- Glandt, K., S. Khanal, Y. Li, D. Caragea, and C. Caragea. 2021a. Stance detection in covid-19 tweets. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, volume 1.
- Glandt, K., S. Khanal, Y. Li, D. Caragea, and C. Caragea. 2021b. Stance detection in covid-19 tweets. In *ACL/IJCNLP*.
- Grave, E., P. Bojanowski, P. Gupta, A. Joulin, and T. Mikolov. 2018. Learning word vectors for 157 languages. In *Proceedings of the International Conference on Language Resources and Evaluation (LREC 2018)*.
- Grover, A. and J. Leskovec. 2016. node2vec: Scalable feature learning for networks. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Hardalov, M., A. Arora, P. Nakov, and I. Augenstein. 2021. Cross-domain label-adaptive stance detection. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9011–9028, Online and Punta Cana, Dominican Republic, November. Association for Computational Linguistics.
- Hardalov, M., A. Arora, P. Nakov, and I. Augenstein. 2022. Few-shot cross-lingual stance detection with sentiment-based pre-training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10729–10737.
- Jiang, A. Q., A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. d. l. Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Kenter, T., A. Borisov, and M. de Rijke. 2016. 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)*, pages 941–951, Berlin, Germany, August. Association for Computational Linguistics.
- Kojima, T., S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa. 2022. Large language models are zero-shot reasoners. In *ICML 2022 Workshop on Knowledge Retrieval and Language Models*.
- Küçük, D. and F. Can. 2020. Stance Detection: a Survey. *ACM Computing Surveys (CSUR)*, 53(1):1–37.
- Lai, M., A. T. Cignarella, L. Finos, and A. Sciandra. 2021. Wordup! at vaxxstance 2021: Combining contextual information with textual and dependency-based syntactic features for stance detection. In *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2021)*, CEUR Workshop Proceedings.
- Lai, M., V. Patti, G. Ruffo, and P. Rosso. 2020. #brexit: Leave or remain? the role of user’s community and diachronic evolution on stance detection. *Journal of Intelligent & Fuzzy Systems*, 31(2):2341–2352.
- Li, Y., C. Zhao, and C. Caragea. 2021. Improving stance detection with multi-dataset learning and knowledge distillation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6332–6345, Online and Punta Cana, Dominican Republic, November. Association for Computational Linguistics.
- Martín-Corral, D., M. García-Herranz, M. Cebrian, and E. Moro. 2022. Social media sensors to detect early warnings of influenza at scale. In *medRxiv*.
- McInnes, L., J. Healy, N. Saul, and L. Großberger. 2018. Umap: Uniform manifold approximation and projection. *J. Open Source Softw.*, 3:861.
- Mikolov, T., K. Chen, G. Corrado, and J. Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Mishra, P., M. Del Tredici, H. Yannakoudakis, and E. Shutova. 2018. Author profiling for abuse detection. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1088–1098, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Mohammad, S., S. Kiritchenko, P. Sobhani, X. Zhu, and C. Cherry. 2016. SemEval-2016 task 6: Detecting stance in tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41, San Diego, California, June. Association for Computational Linguistics.
- Perozzi, B., R. Al-Rfou, and S. Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, page 701–710.
- Pomerleau, D. and D. Rao. 2017. The fake news challenge: Exploring how artificial intelligence technologies could be leveraged to combat fake news. *Fake news challenge*.
- Rashed, A., M. Kutlu, K. Darwish, T. Elsayed, and C. Bayrak. 2021. Embeddings-based clustering for target specific stances: The case of a polarized turkey. In *ICWSM*.
- Schiller, B., J. Daxenberger, and I. Gurevych. 2021. Stance detection benchmark: How robust is your stance detection? *KI-Künstliche Intelligenz*, 53:329–341.
- Sobhani, P., D. Inkpen, and X. Zhu. 2017. A dataset for multi-target stance detection. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 551–557.
- Stefanov, P., K. Darwish, A. Atanasov, and P. Nakov. 2020. Predicting the topical stance and political leaning of media using tweets. In *ACL*.
- Taranukhin, M., V. Shwartz, and E. Milios. 2024. Stance reasoner: Zero-shot stance detection on social media with explicit reasoning. In N. Calzolari, M.-Y. Kan, V. Hoste, A. Lenci, S. Sakti, and N. Xue, editors, *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15257–15272, Torino, Italia, May. ELRA and ICCL.

- Taulé, M., F. Rangel, M. A. Martí, and P. Rosso. 2018. Overview of the task on multimodal stance detection in tweets on catalan loct referendum. In *IberEval 2018. CEUR Workshop Proceedings. CEUR-WS.org*, pages 149–166, Sevilla, Spain.
- Wang, X., J. Wei, D. Schuurmans, Q. V. Le, E. H. Chi, S. Narang, A. Chowdhery, and D. Zhou. 2023. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*.
- Wei, J., X. Wang, D. Schuurmans, M. Bosma, brian ichter, F. Xia, E. H. Chi, Q. V. Le, and D. Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*.
- Zhang, B., X. Fu, D. Ding, H. Huang, Y. Li, and L. Jing. 2023a. Investigating chain-of-thought with chatgpt for stance detection on social media. *arXiv preprint arXiv:2304.03087*.
- Zhang, Z., A. Zhang, M. Li, and A. Smola. 2023b. Automatic chain of thought prompting in large language models. In *The Eleventh International Conference on Learning Representations*.
- Zotova, E., R. Agerri, M. Nuñez, and G. Rigau. 2020. Multilingual stance detection in tweets: The Catalonia independence corpus. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1368–1375, Marseille, France, May. European Language Resources Association.
- Zotova, E., R. Agerri, and G. Rigau. 2021. Semi-automatic generation of multilingual datasets for stance detection in Twitter. *Expert Systems with Applications*, 170:114547.
- Zubiaga, A., B. Wang, M. Liakata, and R. Procter. 2019. Political homophily in independence movements: Analyzing and classifying social media users by national identity. *IEEE Intelligent Systems*, 34:34–42.

A Appendix

The dimensionality for the interaction-based embeddings was chosen by training a SVM classifier via grid search with 5-fold cross-validation. Based on results reported in previous work on reducing huge interaction matrices into low dimensional features for stance detection (Darwish et al., 2020; Stefanov et al., 2020), dimensions for interaction-based embeddings were selected between 10 or 20 dimensions. The best performing RE for CIC and SardiStance were of dimension 20, while for VaxxStance they were of dimension 10. With respect to DW and N2V, the best dimensionality for CIC-es, CIC-es* and SardiStance was 10, whereas the best one for CIC-ca, CIC-ca* and VaxxStance correspond to 20. Moreover, for DW and N2V we set the usual default values for the hyperparameters for these algorithms: `walks_per_node` = 10, `walk_length` = 80, `window` or `context_size` = 10, and the optimization is run for a single epoch (Perozzi, Al-Rfou, and Skiena, 2014; Grover and Leskovec, 2016). Specifically for `node2vec`, we have set `p`=1 and `q`=0.5 in order to enhance network community-related information (Grover and Leskovec, 2016).

Grid search with 5-fold cross-validation was also used to optimize C and Gamma hyperparameters for every SVM system (RE, DW, N2V, FTEmb, TFIDF, and all the combinations).

For XLM-RoBERTa, hyperparameter tuning was done by splitting the training set into training and development sets (80/20). Results on the development set allowed us to obtain the following hyperparameters: 128 maximum sequence length, 16 batch size, 2e-5 learning rate, and 5 epochs.

Figure 6 provides a direct comparison between the RE, DW and N2V user representations for each of the training datasets.

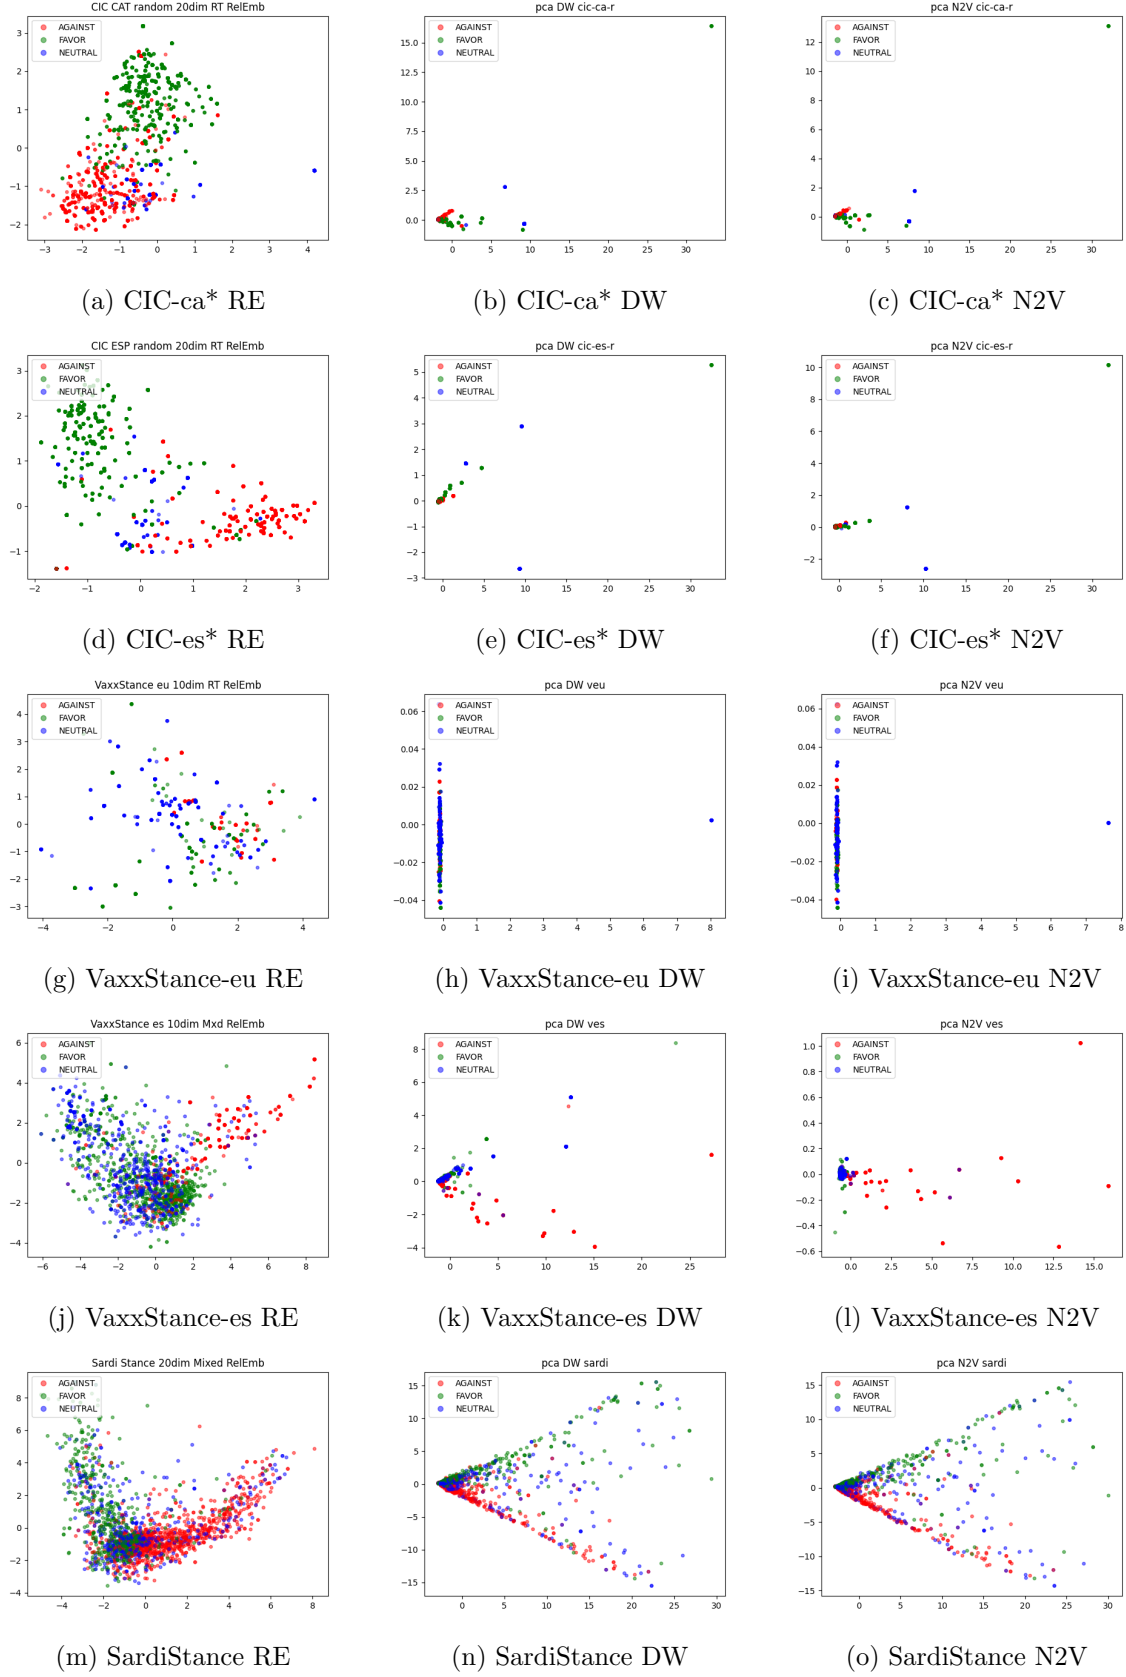


Figure 6: Comparison between Relational Embedding (RE), Deep Walk (DW) and Node2Vec (N2V) representations of training set users (PCA dimension reduction to 2).

B Prompts

B.1 Few-Shot Prompt

Few-shot prompt with 6 examples developed by Taranukhin, Shwartz, and Milios (2024).

Question: Consider the tweet in a conversation about the target, what could the tweet's point of view be towards the target?

The options are:

- against
- favor
- neutral

tweet: <I'm sick of celebrities who think being a well-known actor makes them an authority on anything else. #robertredford #UN>

target: Liberal Values

stance: against

tweet: <I believe in a world where people are free to move and choose where they want to live>

target: Immigration

stance: favor

tweet: <I love the way the sun sets every day. #Nature #Beauty>

target: Taxes

stance: neutral

tweet: <If a woman chooses to pursue a career instead of staying at home, is she any less of a mother?>

target: Conservative Party

stance: against

tweet: <We need to make sure that mentally unstable people can't become killers #protect #US>

target: Gun Control

stance: favor

tweet: <There is no shortcut to success, there's only hard work and dedication #Success #SuccessMantra>

target: Open Borders

stance: neutral

B.2 Stance Reasoner Prompt

The Stance Reasoner prompt containing 6 examples and reasoning chains (Taranukhin, Shwartz, and Milios, 2024).

Question: Consider the tweet in a conversation about the target, what could the tweet's point of view be towards the target?

The options are:

- against
- favor
- neutral

tweet: <I'm sick of celebrities who think being a well-known actor makes them an authority on anything else. #robertredford #UN>

target: Liberal Values

reasoning: the author is implying that celebrities should not be seen as authorities on political issues, which is often associated with liberal values such as Robert Redford who is a climate change activist -> the author is against liberal values

stance: against

tweet: <I believe in a world where people are free to move and choose where they want to live>

target: Immigration

reasoning: the author is expressing a belief in a world with more freedom of movement -> the author is in favor of immigration.

stance: favor

tweet: <I love the way the sun sets every day. #Nature #Beauty>

target: Taxes

reasoning: the author is in favor of nature and beauty -> the author is neutral towards taxes

stance: neutral

tweet: <If a woman chooses to pursue a career instead of staying at home, is she any less of a mother?>

target: Conservative Party

reasoning: the author is questioning traditional gender roles, which are often supported by the conservative party -> the author is against the conservative party

stance: against

tweet: <We need to make sure that mentally unstable people can't become killers #protect #US>

target: Gun Control

reasoning: the author is advocating for measures to prevent mentally unstable people from accessing guns -> the author is in favor of gun control.

stance: favor

tweet: <There is no shortcut to success, there's only hard work and dedication #Success #SuccessMantra>

target: Open Borders

reasoning: the author is in favor of hard work and dedication -> the author is neutral towards open borders

stance: neutral

B.3 Translated tweets for Few-Shot and Stance Reasoner Prompts

We retain the original prompts while translating the language of the 'tweet' into the specific target language.

B.3.1 Basque

```
tweet: <Nazkatuta nago aktore ezagunak izateak beste edozein gaitan aditu
bihurtzen dituela uste duten famatuekin. #robertredford #UN>
tweet: <Jendea aske mugitzeko eta bizi nahi duen lekua aukeratzeko munduan
sinesten dut>
tweet: <Eguzkia egunero nola sartzen den maite dut. #Natura #Edertasuna>
tweet: <Emakume batek karrera bat jarraitzea aukeratzen badu etxean geratu
beharrean, horrek ama txarragoa bihurtzen al du?>
tweet: <Buruko gaixotasunak dituzten pertsonak hiltzaile bihurtu ez daitezen
ziurtatu behar dugu #babestu #AEB>
tweet: <Ez dago arrakastara daraman lasterbiderik, lana eta dedikazioa baino ez
daude #Arrakasta #ArrakastarenGakoa>
```

B.3.2 Catalan

```
tweet: <Etic fart dels famosos que es pensen que ser un actor conegut els
converteix en una autoritat en qualsevol altra cosa. #robertredford #UN>
tweet: <Crec en un món on la gent sigui lliure de moure's i escollir on vol viure>
tweet: <M'encanta com es pon el sol cada dia. #Natura #Bellesa>
tweet: <Si una dona tria seguir una carrera professional en lloc de quedar-se a
casa, és menys mare per això?>
tweet: <Hem d'assegurar-nos que les persones mentalment inestables no puguin
convertir-se en assassins #protegir #EUA>
tweet: <No hi ha dreceres cap a l'èxit, només treball dur i dedicació #Èxit
#MantraDelÈxit>
```

B.3.3 Italian

```
tweet: <Sono stufo dei personaggi famosi che pensano che essere un attore
conosciuto li renda esperti di qualsiasi altro argomento. #robertredford #UN>
tweet: <Credo in un mondo dove le persone siano libere di muoversi e scegliere
dove vogliono vivere>
tweet: <Amo il modo in cui il sole tramonta ogni giorno. #Natura #Bellezza>
tweet: <Se una donna sceglie di perseguire una carriera invece di rimanere a casa,
è meno madre per questo?>
tweet: <Dobbiamo assicurarci che le persone mentalmente instabili non possano
diventare assassini #proteggere #USA>
tweet: <Non esistono scorciatoie per il successo, ci sono solo duro lavoro e
dedizione #Successo #MantraDelSuccesso>
```

B.3.4 Spanish

```
tweet: <Estoy harto de los famosos que creen que ser un actor conocido los
convierte en autoridad sobre cualquier otro tema. #robertredford #UN>
tweet: <Creo en un mundo donde las personas sean libres de moverse y elegir dónde
quieren vivir>
tweet: <Me encanta cómo se pone el sol cada día. #Naturaleza #Belleza>
tweet: <Si una mujer elige seguir una carrera en lugar de quedarse en casa, ¿es
```

```
menos madre por ello?>  
tweet: <Necesitamos asegurarnos de que las personas mentalmente inestables no  
puedan convertirse en asesinos #proteger #EEUU>  
tweet: <No hay atajos hacia el éxito, solo trabajo duro y dedicación #Éxito  
#MantraDelÉxito>
```