Journal of the Brazilian Computer Society, 2026, 32:1, doi: 10.5753/jbcs.2026.5687

A survey of social media stance detection using non-textual features

Laís Carraro Leme Cavalheiro [**University of São Paulo** | ***laiscarraro@usp.br***]
Ivandré Paraboni [**University of São Paulo** | ***ivandre@usp.br***]

School of Arts, Sciences, and Humanities. Av Arlindo Bettio, 1000, São Paulo, SP, 03828-000, Brazil.

Received: 14 March 2025 • **Accepted:** 25 August 2025 • **Published:** 25 March 2026

Abstract. Stance detection is known as the computational task of estimating an individual's attitude towards a given target topic, which is often of a political or moral nature. In traditional NLP fashion, models of this kind have relied mainly on learning features extracted from social media text. However, social media may provide many other types of non-content information in conjunction with text, such as friends networks, interactions with other users, etc. These knowledge sources, despite being potentially useful for stance prediction, remain relatively little discussed in existing surveys of the field. To fill this gap in the literature, this article presents a survey of stance detection research focusing on the use of network-related features and on how these are combined with more standard text models.

Keywords: Natural Language Processing, Stance, Social media, Network features.

1 Introduction

In the analysis of human discourse, *stance* may be defined as a person's attitude, viewpoint, or judgement toward a target topic [Mohammad *et al.*, 2016]. The following sentences, for example, show favourable (ex.1) and opposing (ex.2) stances towards the same target 'legalisation of abortion':

ex1. *A woman's body belongs to her, and she should be able to decide whether she wants to have a child or not.*
ex2. *It is not up to us to decide whether a child that has not even been born yet will live or die.*

Political and moral stances, among others, are abundantly found in online debates and social media in general. This observation has led to the development of a wide range of computational approaches to automatically detect stance from social media data with multiple practical applications. In social sciences, for example, stance detection may help estimate the proportion of votes in a population in a future election [Magdy *et al.*, 2016] or, more generally, may identify dominant trends and opinions on polarising issues [Zhang *et al.*, 2024].

From a computational perspective, and particularly so in the NLP field, stance detection generally involves estimating whether an individual (e.g., a social media user) is against, in favour, or neutral towards a given target [Zhang *et al.*, 2024]. In traditional NLP fashion, this is usually implemented by extracting learning features from *text data* (i.e., social media posts) [Mohammad *et al.*, 2016; Cignarella *et al.*, 2020; Pavan and Paraboni, 2022].

Online social networks may, however, provide many other types of network-related information in addition to various forms of user-produced content. These include, for instance, friends and followers network structures, interactions with other users, publication timestamps, and others, all potentially useful sources of information for stance prediction in their own right. Accordingly, a growing number of stance detection models have been introduced that combine text and non-textual features [Dutta *et al.*, 2022], in most cases motivated by the principle of *homophily* [McPherson *et al.*, 2001], that is, the observation that individuals tend to form connections with others who share similar interests, backgrounds, or values. Thus, for instance, individuals may gather in online communities to discuss a common interest or may interact more frequently with those with whom they share similar opinions.

Despite seemingly positive results, however, stance detection models based on non-textual information are generally less common than purely text-based approaches. This difference is also noticeable in the availability of surveys in the field, which are primarily devoted to text-based stance detection.

The lack of surveys devoted to the use of non-textual features for stance detection is the main motivation of the current work. More specifically, this article presents a survey of stance detection research focusing on the use of *network-related features* (namely, excluding image, sound, and other content-based features that are the domain of multimodal approaches) and on how these are combined with standard text models. The present study, which we believe to be the first of its kind, aims to address the following research questions.

q1. What kinds of textual and non-textual features are employed for stance detection?
q2. How are textual and non-textual features combined into a single stance detection model?
q3. What stance datasets are prevalent in the field?
q4. What kinds of task definition are supported?
q5. What computational methods are implemented?

By shedding light on these issues, we expect to provide a broad view of the stance detection task beyond the standard realm of (text-based) NLP and, in doing so, to present a more structured summary of its current practices and opportunities.

The rest of this article is structured as follows. Section 2 briefly discusses existing surveys in stance detection. Section 3 presents the survey method. Section 4 describes our main findings in light of the above research questions q1–q5. Finally, Section 5 draws conclusions from the present work and discusses its limitations.

2 Related work

Existing surveys of stance detection methods are few and are largely devoted to the more standard text-based task, along the lines of Mohammad *et al*. [2016] and others. For instance, the recent survey in Zhang *et al*. [2024] is mainly focused on text, including only a side discussion of multimodality. Similarly, the work in Khiabani and Zubiaga [2024] is focused solely on text-based methods.

Some of the existing surveys focus on other tasks related to stance detection or on specific applications. For instance, the studies in Wang *et al*. [2019] and Zubiaga *et al*. [2018a] address the related issue of aspect-level stance detection in product reviews. Similarly, the surveys in Hardalov *et al*. [2022] and Akhtar *et al*. [2018] focus on the identification of rumours and disinformation on social media.

The exception seems to be the survey in Aldayel and Magdy [2021], which is, to our knowledge, the only one to explicitly mention the use of so-called network features. In the survey, these are divided into two categories, namely user behaviour data and user meta-data attributes, but the discussion is kept brief. In addition to that, the survey only covers studies that have been published up to the year 2019.

3 Approach

We carried out a systematic review of methods of stance detection that included the use of non-textual learning features. Table 1 details the sources and queries taken into account.

We retrieved an initial set of 765 studies. From this set, we only kept the studies that met all the inclusion criteria (i1..i3) and none of the exclusion criteria (e1..e3) as follows.

i1. The selected studies should be fully available from a public scientific database.
i2. The selected studies should address the issue of stance detection from social media data.
i3. The selected studies should present a survey of stance detection, computational models, or datasets for the task, even when not using text-based features.

e1. Studies published before 2016 are discarded.
e2. Studies not related to NLP or social media analysis are discarded.
e3. Studies that discuss purely textual models or purely textual datasets are discarded.

After applying the above filters, we obtained 51 studies on stance detection that use network-related information, hereby called 'non-textual' features, and which were published between 2016 and 2024. The yearly distribution of the articles is illustrated in Figure 1.

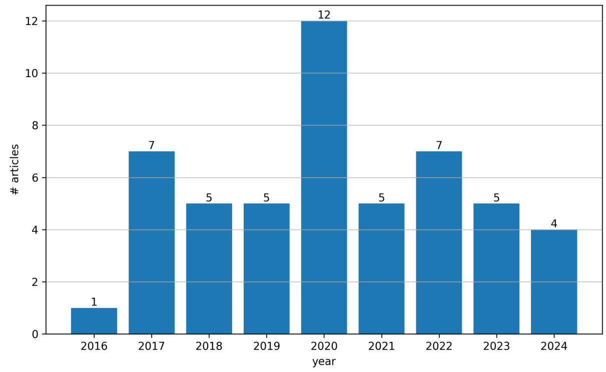

Figure 1. Articles per year.

In order to investigate our research questions q1–q5 described in the previous section, we attempted to identify in each study the following attributes: the kinds of textual and non-textual features under consideration (q1), the methods for combining them into a single model (q2), dataset information (q3), task definitions (q4) and main computational methods (q5). For ease of discussion, the possible values to be assigned to each attribute are summarised as follows and will be discussed further in the next sections.

- *Non-textual features*: *int.*: interaction features such as mentions, retweets, replies, etc.; *pref.*: social media preferences such as favourites, participation in communities, etc.; *con.*: connections to other users through friendship, follower/following relationships, etc.

- *Textual features*: *text*: underlying text representation model such as bag-of-words, word embeddings, etc.; *surf*: surface or syntactic features such as text length, hashtag counts, URL counts, word classes, etc.; *emo*: emotion- and cognition-related information sentiment, psycholinguistics-motivated features, etc.; *demo*: demographics inferred from text such as location, age, gender, political affiliation, etc.

- *Combination (Comb.)*: the method for combining textual and non-textual features, if any. *concat*: intermediate fusion through feature concatenation; *sim*: intermediate fusion based on similarity-based metrics; *max*: intermediate fusion through function maximisation; *vote*: late fusion by majority voting; or *other*.

- *Media*: the genre of the text or online social media type from which the train / test data are extracted. *fo*: online discussion forums; *nw*: news; *tw*: Twitter/X; *rd*: Reddit; *fb*: Facebook; *yt*: YouTube.

- *Language (Lang.)*: language of the text data, if any. *ar*: Arabic; *en*: English; *it*: Italian; *es*: Spanish; *fr*: French; *pt*: Portuguese.

- *Dataset*: corpus or dataset taken as train/test data. *SemEval*: the SemEval-2016 stance corpus in Mohammad *et al*. [2016]; *Rumour*: the RumourEval-2017 corpus in Derczynski *et al*. [2017]; *PHEME*: the

Table 1. Article sources and initial queries

Source	Queries	Articles
IEEE	Topics: "learning (artificial intelligence)", "social networking (online)", "pattern classification","text analysis", "feature extraction", "neural networks"	452
ACM	[Title: "stance"] AND [Abstract: "stance"] AND [Publication Date: (01/01/2016 TO 12/07/2024)]	178
USP theses	Contains "stance" in title and abstract	0
CAPES journals	peer-reviewed, English or Portuguese, areas "Technology" and "Computer Science" except "Engineering" and "Robotics"	184
ACL Anthology	Contains "stance" in any field	275

PHEME corpus in Ma *et al*. [2018]; *ConRef*: the ConRef-STANCE-ita corpus in Lai *et al*. [2019].

- *Size*: total number N of positive and negative instances in the dataset, clustered as 1..7 size categories. *1*: $<$ 5k; *2*: 5k $<$ N $<$ 10k ; *3*: 10k $<$ N $<$ 50k ; *4*: 50k $<$ N $<$ 100k; *5*: 100k $<$ N $<$ 500k ; *6*: 500k $<$ N $<$ 1mi ; *7*: $>$ 1mi.

- *Task*: computational task definition. *stance*: standard stance classification or target-oriented stance detection [Mohammad *et al*., 2016]; *rumour*: rumour detection, a 4-class (deny, support, comment and query) variation of the standard stance classification task, introduced in Derczynski *et al*. [2017]; *fakenews*: use of stance information as a feature to determine whether a given piece of information is false [Jamialahmadi *et al*., 2022]; *analysis*: descriptive studies of group or community stance on social media, as in Liu *et al*. [2023]; *dataset*: creation or stance corpora or other datasets for stance detection, as in Abeysinghe *et al*. [2022].

- *Type*: the type of stance target under consideration. *target*: target-based stance detection as in the SemEval tradition [Mohammad *et al*., 2016]; *claim*: claim-based or text-to-text stance detection/classification, as in, e.g., Allaway and McKeown [2020]; Zhao and Caragea [2024].

- *Methods*: computational or machine learning methods, if any. *svm*: support vector machine; *lr*: logistic regression; *nb*: Naive Bayes; *rf*: random forest; *lstm*: long short-term memory networks; *dt*: decision tree induction; *cnn*: convolutional neural network; *mlp*: multilayer perceptron; *rnn*: recurrent neural network; *dnn*: other deep neural network; *xgb*: extreme gradient boosting; *cd*: community detection; *graph*: graph-based methods; *bert*: bidirectional encoder representations for transformers; or *other*.

4 Survey results

Table 2 summarises the results of our review of the literature, which are discussed in the next sections. Information not disclosed in the source articles is marked as '-'.

4.1 Overview

Since the focus of the present survey is on the use of non-textual features in stance detection, we took the feature taxonomy in Chancellor and De Choudhury [2020] as a starting point to categorise the surveyed studies as summarised in Table 3.

Details concerning the use of textual and non-textual features in the surveyed studies are provided in the following sections. With regard to the primary focus of this survey – the use of non-textual features – we notice that interactions, preferences, and connections frequently overlap, thereby preventing a clear-cut categorisation of existing approaches. Furthermore, since the selection of features is largely constrained by the type of information available on a given social media platform (for instance, Facebook provides 'friendship' relations, whereas Reddit does not), feature choices should not be interpreted as an indication that one approach is inherently superior to another. Differences between models do not necessarily indicate advantages or disadvantages and may simply reflect the constraints of a given social media platform. As a result, the discussion that follows should be regarded as a portrayal of existing practices in the field rather than as a comparative analysis.

4.2 Textual and non-textual feature use (q1)

Question q1 mainly intends to identify the kinds of non-textual feature used in stance detection and, when relevant, the textual features included in these models. The two issues are discussed individually as follows.

4.2.1 Non-textual features

Non-textual features – on the left side of Table 3 – are divided according to the classification in Aldayel and Magdy [2019].

Table 2. Survey results

Study	Non-text	Text	Comb.	Media	Lang.	Dataset	Size	Task	Type	Methods
Igarashi *et al.* [2016]	int.	surf emo txt	concat	tw	en	SemEval	1	**stance**	target	lr
Darwish *et al.* [2017]	int.	surf txt	sim	tw	ar en	other	3	**stance**	target	svm
Lai *et al.* [2017]	int.	surf emo	concat	tw	en	SemEval	1	**stance**	target	nb
Bahuleyan and Vechtomova [2017]	int.	surf emo	concat	tw	en	Rumour	2	rumour	claim	xgb
Addawood *et al.* [2017]	con. int.	surf txt emo	concat	tw	en	other	6	**stance**	target	svm dt nb
Dong *et al.* [2017]	int.	surf emo	max	nw fo	en	other	7	**stance**	claim	other
Veyseh *et al.* [2017]	con. int. pref.	surf txt	concat	tw	en	Rumour	2	rumour	claim	cnn lstm
Sadiq *et al.* [2017]	con. int.	surf emo demo	other	tw	en	other	1	**stance**	target	rf
Fraisier *et al.* [2018]	con. int.	surf demo	sim	tw	en fr	other	7	**stance**	target	cd
Chen and Ku [2018]	pref.	emo	max	fb	en	other	3	**stance**	target	other
Akhtar *et al.* [2018]	int. pref.	surf txt emo demo	concat	tw	en	Rumour	2	rumour	claim	svm nb dt mlp
Zubiaga *et al.* [2018b]	int. pref.	surf txt	concat	tw	en	PHEME	1	rumour	claim	lstm other
Al-Ayyoub *et al.* [2018]	int.	emo	-	tw	ar	other	7	analysis	claim	graph
Lynn *et al.* [2019]	con.	demo	concat	tw	en	SemEval	1	**stance**	target	lr
Aldayel and Magdy [2019]	con. int. pref.	surf txt	concat	tw	en	SemEval	1	**stance**	target	svm
Lai *et al.* [2019]	con. int.	surf	concat	tw	it	ConRef	1	**stance**	target	svm
Xuan and Xia [2019]	con. int. pref.	surf emo	concat	tw	en	Rumour	2	rumour	claim	svm rf nb lr dt
Chang *et al.* [2019]	con.	surf txt demo	concat	tw	en	other	2	**stance**	target	mlp
Islam *et al.* [2019]	con.	txt	concat	tw	en	PHEME	2	rumour	claim	rnn
Lai *et al.* [2020a]	con. int. pref.	surf txt emo	concat	tw	in o	SemEval	3	**stance**	target	svm lr nb
Lai *et al.* [2020b]	int. pref.	surf txt emo	concat	tw	en	other	2	**stance**	target	svm
Vanta and Aono [2020]	con. int. pref.	surf	concat	tw rd	en	Rumour	2	rumour	claim	lstm
Graells-Garrido *et al.* [2020]	int.	surf txt demo	concat	tw	es	other	7	**stance**	target	xgb
Tyagi *et al.* [2020]	int.	surf txt emo	other	tw	en	other	7	analysis	target	-
Yang *et al.* [2020]	int.	txt	max	tw	en	other	2	**stance**	target	other
Ebeling *et al.* [2020]	con. int. pref.	surf demo	-	tw	pt	other	5	analysis	target	-
Vilella *et al.* [2020]	int. pref.	surf txt demo	-	tw	it	other	7	analysis	target	cd
Kumar *et al.* [2020]	int.	txt	-	tw	en	other	-	analysis	target	graph
Blackburn *et al.* [2020]	int. pref.	surf txt emo	-	tw	en	other	3	dataset	claim	-
Li and Scarton [2020]	con. int.	surf txt	concat	tw	en	Rumour	3	rumour	claim	lr rf mlp bert
Rezayi *et al.* [2021]	con. pref.	txt	concat	tw	en	PHEME	1	fakenews	claim	dnn
Jia *et al.* [2021]	int. pref.	txt emo	concat	tw	it	ConRef	1	**stance**	target	svm lstm
Masood and Abbasi [2021]	con. int. pref.	txt emo demo	concat	tw	en	other	5	rumour	claim	svm rf nb lr
Tyagi *et al.* [2021]	int.	emo	-	tw	en	other	7	analysis	target	-
Geiss *et al.* [2022]	int. pref.	txt demo	other	rd	en	other	1	**stance**	target	svm lr rf
Dutta *et al.* [2022]	con.	surf txt	vote	tw	en	other	5	**stance**	target	cnn lstm
Abeysinghe *et al.* [2022]	int. pref.	surf txt	-	tw	en	other	7	dataset	claim	-
Luo *et al.* [2022]	con. int. pref.	surf txt emo	concat	tw	en	other	1	rumour	claim	mlp
Rochert *et al.* [2022]	int. pref.	txt	other	yt	en	other	5	analysis	target	bert
Aldayel and Magdy [2022]	con. int.	-	concat	tw	en	SemEval	1	**stance**	target	svm
de Oliveira [2022]	int.	txt	other	tw	pt	other	7	**stance**	target	bert
Abdine *et al.* [2022]	int.	-	other	tw	fr	other	7	**stance**	target	cd
Williams and Carley [2023]	int.	txt	other	tw	es	other	7	**stance**	target	bert graph
Pougué-Biyong *et al.* [2023]	int. pref.	-	concat	tw	en	other	6	**stance**	target	lr graph
dos Santos and Goya [2023]	int.	txt	sim	tw	pt	other	3	**stance**	target	bert
Liu *et al.* [2023]	int.	txt	other	tw	fr	other	7	analysis	claim	graph
Cavalheiro *et al.* [2023]	con. int.	txt	vote	tw	pt	other	4	**stance**	target	bert lr
de Vinco *et al.* [2024]	int.	txt	other	rd	en	other	5	**stance**	target	llm
Kuo *et al.* [2024]	int.	txt emo	other	fb	o	other	1	**stance**	target	bert graph
Sutter *et al.* [2024]	int.	txt	sim	tw	en	ConRef	1	**stance**	target	gnn bert
Penzo *et al.* [2024]	int.	txt	other	o	en	other	5	**stance**	claim	bert mlp

In this classification, *interaction* features (int.) represent direct exchanges between social media users (mentions to other individuals, responses, etc.); *preference* (pref.) features represent users' content preferences and participation in social media groups; and *connection* (con.) features model the social media relationships between individuals, being either bidirectional (when both individuals are linked to each other, as in, e.g., Facebook 'friendship' relations), or unidirectional (when only one individual is connected to the other, but not the other way round, as in, e.g., Twitter/x 'follow' relations.)

As in the 'Non-text' column in Table 2, Figure 2 illustrates the main types of non-textual features found in the survey.

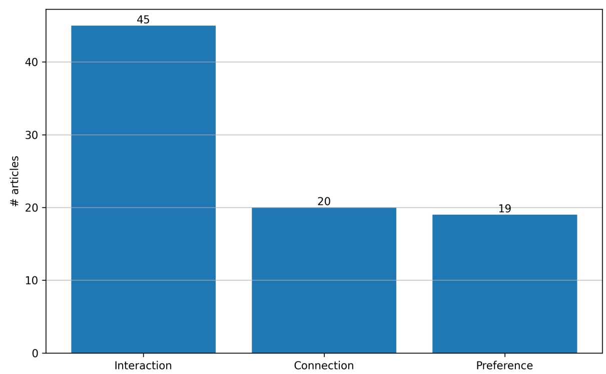

Figure 2. Non-textual feature usage.

Figure 2 shows that the most common source of non-textual features for the task was found to be interaction-related (int.) information. To some extent, this was expected, as exchanges among social media users are likely to be an expression of ho-

Table 3. Textual and non-textual features for stance detection

Non-textual		Textual	
Interactions (int.)	mentions reposts / retweets replies	Surface (surf)	punctuation emojis part-of-speech lemmas dependencies hashtags URLs
Preferences (pref.)	favourites community affiliation	Text models (text)	BoW word embeddings BERT
Connections (con.)	unidirectional (follow / following) bidirectional (friendship)	Emotion and Cognition (emo)	sentiment psycholinguistics info
		Demographics inference (demo)	age, gender, location

mophily and, accordingly, are generally perceived as strong stance predictor candidates. In other words, the users with whom an individual keeps conversation (about the target topic or others) are often taken as learning features for stance detection.

Interaction features are followed by user connections (con.) such as social media friendship relations, which are also influenced by homophily. Interestingly, we also notice that most studies model non-textual features simply as atomic properties (or simple counts) such as the number of friends or mentions to other users (e.g., Lai *et al*. [2017]; Bahuleyan and Vechtomova [2017]), and that few studies build full relational models using, e.g., friendship graphs [Espinosa *et al*., 2020].

4.2.2 Textual features

Textual features – on the right side of Table 3 – are divided into four categories: *surface* features (surf) represent text surface form (e.g., text length, hashtag counts, URLs, etc.), or syntactical information (e.g., syntactic dependencies, grammatical classes, etc.); the *text models* (txt) category encompasses word- and sentence-related information provided by an underlying text representation (e.g, words frequencies in bag-of-words model, static or contextual word embeddings, etc.) in general-purpose models such as BERT [Devlin *et al*., 2019] or more genre-specific alternatives such as RetweetBERT [de Oliveira, 2022] or BERTabaporu [Costa *et al*., 2023]; *emotion and cognition* features (emo) encompass the kinds of information computed with the aid of affective or psycholinguistics dictionaries such as LIWC [Pennebaker *et al*., 2001], as in Ferraccioli *et al*. [2020]; and the *demographics* inference category (demo) comprises those features that, although user-related (e.g., age, gender, etc.), are inferred automatically from text using author profiling or similar methods (e.g., Flores *et al*. [2022]; Onikoyi *et al*. [2023]; dos Santos and Paraboni [2022]).

As in the 'Textual' column in the previous Table 2, Figure 3 illustrates the main types of textual features found in the survey.

Figure 3 shows that nearly all studies model some form of text feature or text representation (e.g., dos Santos and Goya [2023]; de Vinco *et al*. [2024]; Penzo *et al*. [2024]). Among these, the use of full text models (e.g., BoW and word embeddings) is the most common approach [Kumar *et al*., 2020; Rochert *et al*., 2022]. This is followed by the use of surface features in general (e.g., hashtags, URLs, punctuation,

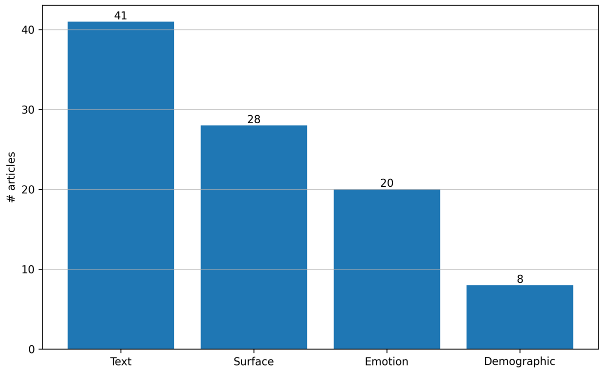

Figure 3. Textual feature usage.

etc.) as in Lai *et al*. [2019]; Vanta and Aono [2020]. We also notice that, although less common in recent years, the use of emotion and cognition-related features (sentiment dictionaries, psycholinguistics-motivated features, etc.) appears in a considerable number of studies (e.g., Al-Ayyoub *et al*. [2018]; Tyagi *et al*. [2021]). On the other hand, few studies make use of large language models, which may be explained by the observation that the focus of many of these studies is the use of non-textual features, and less so the text component of the model.

4.3 Feature combination (q2)

Question q2 seeks to identify the computational strategies for combining textual and non-textual features in existing stance detection models, which are present in most studies[1]. In multimodal machine learning, strategies of this kind are usually divided into early, intermediate, and late fusion [Boulahia *et al*., 2021]. Early fusion consists of integrating different modalities of raw data into a single representation before training. Intermediate fusion, or feature-level fusion, computes and then unifies features from raw data. Late fusion combines decisions (e.g., class probabilities or labels) from single-modality architectures to produce a final decision.

As in the 'Comb.' column in the previous Table 2, Figure 4 illustrates the main feature combination strategies found in the survey.

Figure 4 shows that none of the existing work makes use of early fusion methods. This is to be expected, as models of this kind would most likely require large datasets, which are not common in the field of multimodal stance detection. Most

[1]The exceptions are those that do not propose a novel stance detection model, as in the case of dataset creation, etc.

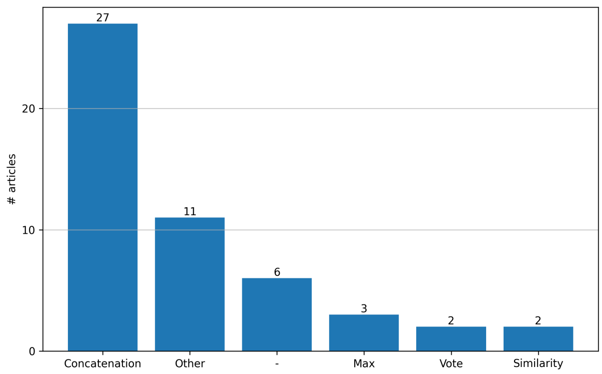

Figure 4. Feature combination strategies.

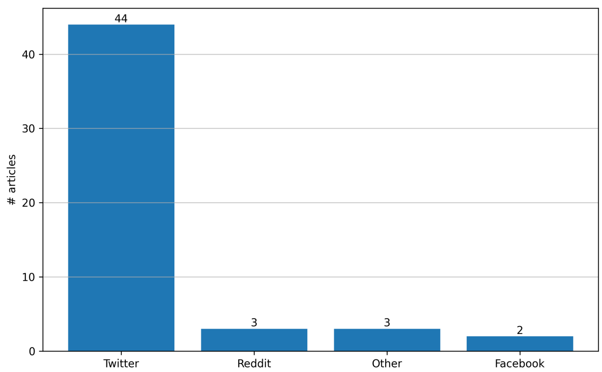

Figure 5. Social media types.

of the existing work resorts instead to simple feature concatenation, which is an instance of intermediate or feature-level fusion. This may be surprising given the significant differences between the two feature types (e.g., word embedding vectors and friend counts), but it may once again be explained by the observation that these studies tend to focus more on the use of non-textual features and pay less attention to text and its aggregation to the stance detection model even though contents (e.g., text) is arguably the single most important stance predictor [Cavalheiro *et al*., 2023].

Interestingly, there are few instances of late fusion methods in the literature, and only the work in Cavalheiro *et al*. [2023] makes use of dedicated classifiers for textual and non-textual features. The only two other studies that implement an instance of late fusion [Islam *et al*., 2019; Espinosa *et al*., 2020] do so by combining results from multiple and often redundant classifiers that may share attributes even between modalities. Methods of this kind, particularly when using stack ensemble architectures [Wolpert, 1992], have been shown to obtain positive results in other social media analysis tasks (e.g., Custódio and Paraboni [2021]; de Souza *et al*. [2022]), but the issue remains little explored in the present context.

4.4 Stance detection datasets (q3)

Question q3 intends to identify the main datasets employed in stance detection according to the type, language, and size of social media. As in the 'Media' column in Table 2, Figure 5 illustrates the types of social media found in the survey.

Figure 5 shows that Twitter is by far the most popular choice (e.g., Luo *et al*. [2022]). This may be explained by the ease of collection that this platform used to afford before being relaunched as X.

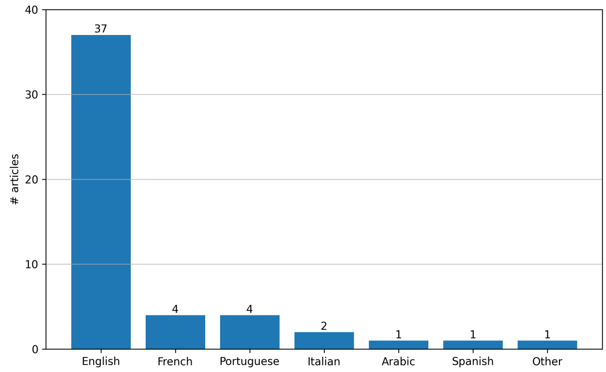

Figure 6. Target languages.

As in the 'Lang.' column in Table 2, Figure 6 illustrates the target languages found in the survey.

Figure 6 shows that, as elsewhere in NLP, the language of the (text) datasets is mostly English (e.g., Aldayel and Magdy [2022]; Rochert *et al*. [2022]), followed by other European languages. In particular, the use of the Italian language driven by the EvalIta2020 stance detection shared task in Cignarella *et al*. [2020] is noteworthy.

Regarding the use of standard datasets or benchmarks in the field, the 'Dataset' column in Table 2 shows a number of studies focused on well-known and relatively small datasets such as SemEval2016 [Mohammad *et al*., 2016], RumourEval2017 [Derczynski *et al*., 2017], PHEME [Ma *et al*., 2018] and ConRef-STANCE-ita [Lai *et al*., 2019]. However, most studies choose other resources that are in many cases developed from scratch. This seemingly lack of standards may be motivated by the need to use larger datasets and more non-textual features since in some cases (most noticeably in the case of SemEval2016), non-textual features tend to be limited to mentions to other users, which may be processed simply as textual features. Moreover, since stance detection tends to be culture-dependent (e.g., stance data on the Sardines movement in Italy, as in Cignarella *et al*. [2020], may be less relevant to other countries and languages), it seems only natural that different research projects will focus on more local issues.

Finally, the 'Size' column in Table 2 shows that most studies are based on relatively small datasets with fewer than 5k instances, and that nearly half use fewer than 10k instances.

4.5 Task definitions (q4)

Question q4 identifies the task definitions found in the literature, which includes both standard stance detection and, in some cases, related tasks. As in the 'Task' column in Table 2, Figure 7 illustrates the task definitions found in the survey.

Figure 7 shows that studies related to stance detection generally focus on stance classification proper (e.g., Aldayel and Magdy [2022]; Dutta *et al*. [2022]), but may also address other tasks such as rumour detection (e.g., Masood and Abbasi [2021]; Luo *et al*. [2022]). The latter is mainly driven by the RumourEval rumour and stance shared tasks series in Derczynski *et al*. [2017]; Gorrell *et al*. [2019].

The 'Target' column in Table 2 mirrors a well-established division in stance detection research. In the tradition of SemEval2016 [Mohammad *et al*., 2016], most studies implement some form of *target-oriented* stance detection by considering a fixed, predefined set of target topics for stance

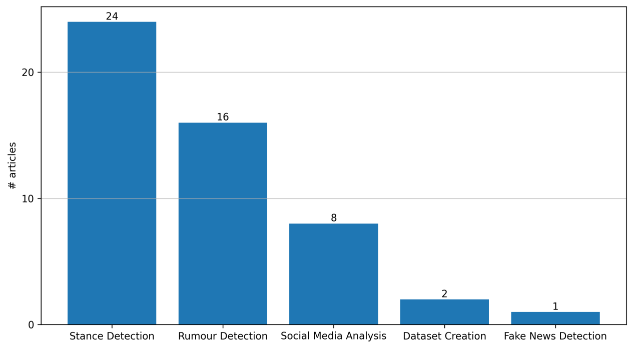

Figure 7. Task definitions.

detection. Others, particularly when focused on zero-shot stance detection, do away with the notion of a target and instead focus on assessing whether one piece of text supports a given claim, in so-called claim-based stance detection (e.g., Luo *et al*. [2022]; Abeysinghe *et al*. [2022]). For further details on zero-shot and claim-based stance detection, see, e.g., Allaway and McKeown [2020]; Zhao and Caragea [2024]; Pavan and Paraboni [2024]; Pereira *et al*. [2026].

4.6 Computational methods (q5)

Question q5 describes the computational methods – particularly focused on machine learning – found in the articles surveyed. As illustrated in the 'Methods' column in Table 2, these are not unlike those observed in other NLP applications. The most common choices are (perhaps surprisingly) Support Vector Machine (SVM) (e.g., as in Aldayel and Magdy [2022]; Geiss *et al*. [2022]), Logistic Regression (LR) [Masood and Abbasi, 2021; Geiss *et al*., 2022] and Naive Bayes (NB) [Lai *et al*., 2020a; Masood and Abbasi, 2021]. In contrast, neural models that may be considered the SOTA in NLP are much less frequent. As discussed in the previous sections, this also may be explained by the more general focus on non-textual features, in which case the use of certain neural approaches (e.g., for sequence labelling as in the case of text data in Veyseh *et al*. [2017]; Zubiaga *et al*. [2018b]; Vanta and Aono [2020]; Jia *et al*. [2021]; Dutta *et al*. [2022]) may be unhelpful or unnecessarily costly.

Regarding the textual component of existing models, by contrast, this scenario has changed more recently with the use of Bidirectional Encoder Representations from Transformers (BERT) architecture [Devlin *et al*., 2019], a now well-established standard in text classification. BERT has played a central role in a number of more recent studies, such as in Li and Scarton [2020]; Rochert *et al*. [2022], but current practices still lack behind those observed in purely text-based NLP tasks, which tend to be more focused on the use of much larger language models.

4.7 Summary

The surveyed studies comprise a wide range of approaches to stance detection using textual or non-textual information, or both. In these studies, text features are nearly always taken into account and, generally speaking, outperform other knowledge sources for the task, including network-related features, by a wide margin. This should not come as a surprise since user-produced content is semantically richer than, e.g., user connections, making text a primary knowledge source for social media stance detection. In other words, non-textual features are generally taken as being *complementary* to the information provided by the text, but rarely considered in their own right. Table 4 summarises a number of pros and cons of these two perspectives, which are discussed in more detail below.

In addition to generally obtaining higher accuracy, text-based models are arguably trained from more easily accessible data (i.e., social media text publications as opposed to network-related information) and are less dependent on a particular social media type. On the other hand, models of this kind are obviously sensitive to the ambiguity of natural language, and since their input consists of an explicit text stance, they may be less suitable for identifying user's stances that have not been explicitly undisclosed on social media [Cavalheiro *et al*., 2023].

In addition to that, models that rely on network-related information tend to require more complex data collection procedures and are often specific to a particular type of social media. For instance, a model built from Facebook friendship relations does not easily generalise to other social media such as X or Reddit, in which these relations are not available. Some of these weaknesses may be naturally mitigated by hybrid models that combine both textual and non-textual features, as indeed found in the vast majority of the surveyed studies.

5 Final remarks

This work presented a survey on stance detection studies using non-textual features extracted from social media. Studies of this kind, which are relatively few compared to the more common text-based NLP approaches, show that purely text-based models, if enriched with non-textual information, may obtain significant performance gains, as in Darwish *et al*. [2017]; Dong *et al*. [2017]; Chen and Ku [2018]; Lai *et al*. [2019]; Lynn *et al*. [2019]; Aldayel and Magdy [2019]; Espinosa *et al*. [2020]; Jia *et al*. [2021]; Rezayi *et al*. [2021]; Dutta *et al*. [2022]; Geiss *et al*. [2022]. Based on the present investigation, in what follows we highlight a number of points that indicate opportunities and directions for future research in the field.

5.1 The role of textual data

Although the focus of the literature review was mainly on the use of non-textual features, we notice that textual (or more generally content-based) features are generally considered to be the main knowledge source in stance detection, and that models that do not rely on text data remain rare. More importantly, reliance on textual features gives rise to the question of how to perform the task when text is unavailable if, for instance, we need to estimate the stance of a social media user who does not discuss the intended target publicly. In these situations, the design of non-textual models might be a way forward[2].

[2]For an introduction to this issue, see, e.g., Cavalheiro *et al*. [2023].

Table 4. Textual versus non-textual feature usage

Criteria	Non-textual	Textual
Task accuracy	low	high
Train data requirements	large network datasets	explicit text statements
Ambiguity issues	minimal	sensitive to sarcasm, irony, etc.
Dependency on social media type	high	moderate
Language dependency	no	yes
Suitability for user stance detection	high	limited
Existing work	few	many

5.2 The need for zero-shot learning

Also related to the issue of lack of knowledge for the task, we notice that all of the non-textual models identified in the present survey follow the standard *in-domain* approach to stance detection [Gorrell *et al*., 2019], that is, they are trained and tested using data about the same target topic. This scenario is considerably behind current practices in text-based stance detection, in which the focus has nowadays shifted considerably toward *zero-shot* learning so as to favour models that are tested on topics not seen during training [Allaway and McKeown, 2022]. This shift, which is largely motivated by the observations that creating training data for every possible topic of interest is unfeasible and that practical stance detection models should not have to rely on the availability of labelled training data, arguably makes zero-shot learning the standard approach for text-based stance detection, but the issue has yet to be investigated in non-textual settings.

5.3 Non-textual feature representation

The survey has revealed that interaction-related information (e.g., the users with whom an individual establishes a conversation) is the most frequently used non-textual feature in stance detection. This is consistent with the expected effects of homophily [McPherson *et al*., 2001], and further highlights the relevance of non-textual features for stance detection and social media analysis in general. However, even taking the need for non-textual features for granted, the way in which these features should be represented remains less clear. In particular, we have found that many existing studies model non-textual information simply as atomic properties (e.g., by simply counting the number of social media friends of an individual, etc.), which have in many cases obtained minimal or negative influence on the results [Akhtar *et al*., 2018; Xuan and Xia, 2019; Masood and Abbasi, 2021; Rezayi *et al*., 2021]. This may not be entirely surprising given that, for instance, the number of friends of an individual may not be correlated with their stance towards a target. On the contrary, the number of friends and similar user counts may be actually influenced by many other factors such as the time since the creation of their social media account, their frequency of use, etc., which may bear little or no relation to their stance towards a particular topic.

As an alternative to simple user or connection counts, the use of relational information (e.g., modelling full networks of friends, followers, etc. as graphs or similar structures) has been found to be a potentially more useful approach to stance detection [Lai *et al*., 2019; Lynn *et al*., 2019; Aldayel and Magdy, 2019; Lai *et al*., 2020b; Espinosa *et al*., 2020; Dutta *et al*., 2022]. Studies of this kind are, however, still relatively few, which may be explained by the added challenges concerning data collection and processing potentially large network structures. It is not uncommon, for instance, for a Twitter/X user to have several thousand friends, which may escalate significantly even in small datasets, or when considering the need for building additional network structures of followers, interactions, and others. In this regard, methods analogous to those seen in text processing may represent an opportunity to improve current stance detection models. These include, for instance, the use of graph neural networks [Sutter *et al*., 2024], node embeddings [Grover and Leskovec, 2016], etc.

5.4 Fusion methods

Regardless of how non-textual information is represented, however, one question that may still deserve further investigation is how these should be combined with standard textual features. In most of the studies reviewed, intermediate fusion is the method of choice [Aldayel and Magdy, 2022; Luo *et al*., 2022; Masood and Abbasi, 2021] but, as discussed in Boulahia *et al*. [2021], late fusion methods may require less training data, which may be particularly useful in stance detection for resource-poor languages and topics.

5.5 Resources for non-textual stance research

With respect to the datasets taken as the basis for existing work, Twitter/X is clearly the most popular choice but, given the present difficulties in sharing and reusing Twitter/X data, it remains to be seen whether this trend will continue or whether it may be replaced by Reddit (e.g., as in Vanta and Aono [2020]; Geiss *et al*. [2022]) or others (e.g., Dong *et al*. [2017]; Chen and Ku [2018]; Rochert *et al*. [2022]). We also notice that most of these datasets are devoted to the English language and, as pointed out in Ebeling *et al*. [2020], remain largely unavailable for reuse for reasons of, e.g., data protection or confidentiality. However, given that stance detection corpora and models are often focused on events that only affect specific cultures (e.g., elections and other political-related events in a particular country or region), and since network dynamics may vary considerably across cultures, expanding the stance detection task to cover other languages and topics (even when written text is not taken into account) represented in new datasets is still very much in demand to obtain a full picture of the global social media landscape.

5.6 Ethical considerations

Beyond open research questions, the present survey also highlights some ethical issues. Although existing datasets for stance detection are generally anonymised, we notice that recovering the original posts from social media text – and the corresponding authorship information – is in most cases straightforward, and this may be particularly problematic given that stance corpora often cover sensitive issues of political or moral nature [Mohammad *et al*., 2016]. In contrast, the situation regarding the use of non-textual information (e.g., network connections) seems more manageable as usernames and other critical information may be in principle replaced by anonymised tokens with little loss of informativeness (e.g., Pereira *et al*. [2026]).

5.7 Limitations

Finally, it should be acknowledged that the present survey has a number of limitations. In particular, the current choice of queries may have excluded a number of studies in related tasks (e.g., sentiment analysis, opinion mining, etc.) that could fit the present analysis. It would be interesting to investigate how these tasks compare to each other and which methods or features they may share. As in the case of the above observations, this issue is left as a suggestion for future work.

Declarations

Authors' Contributions

LCLC contributed to conceptualisation, investigation, methodology, and the writing of the original draft. IP contributed to conceptualisation, project supervision, and writing review and editing. LCLC is the main contributor and writer of this manuscript. Both authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Availability of data and materials

This survey does not involve data or code that will be shared.

References

Abdine, H., Guo, Y., Rennard, V., and Vazirgiannis, M. (2022). Political communities on Twitter: Case study of the 2022 French presidential election. In *LREC 2022 workshop on Natural Language Processing for Political Sciences*, pages 62–71. Available at:`https://aclanthology.org/2022.politicalnlp-1.9/`.

Abeysinghe, B., Vulupala, G. R., and Sunderraman, R. (2022). Misinformation in social media platforms and web articles: a dataset to infer user stance. In *IEEE 16th International Conference on Semantic Computing (ICSC)*, pages 269–273. IEEE. DOI: 10.1109/icsc52841.2022.00051.

Addawood, A., Schneider, J., and Bashir, M. (2017). Stance classification of twitter debates: The encryption debate as a use case. In *8th international conference on Social Media & Society*, pages 1–10. DOI: 10.1145/3097286.3097288.

Akhtar, M. S., Ekbal, A., Narayan, S., and Singh, V. (2018). No, that never happened!! investigating rumors on twitter. *IEEE Intelligent Systems*, 33(5):8–15. DOI: 10.1109/mis.2018.2877279.

Al-Ayyoub, M., Rabab'ah, A., Jararweh, Y., Al-Kabi, M. N., and Gupta, B. (2018). Studying the controversy in online crowds' interactions. *Applied Soft Computing*, pages 557–563. DOI: 10.1016/j.asoc.2017.03.022.

Aldayel, A. and Magdy, W. (2019). Your stance is exposed! analysing possible factors for stance detection on social media. In *ACM proceedings on Human-Computer Interaction*, volume 3, pages 1–20. ACM New York. DOI: 10.1145/3359307.

Aldayel, A. and Magdy, W. (2021). Stance detection on social media: State of the art and trends. *Information Processing & Management*, 58(4):102597. DOI: 10.1016/j.ipm.2021.102597.

Aldayel, A. and Magdy, W. (2022). Characterizing the role of bots' in polarized stance on social media. *Social Network Analysis and Mining*, 12(1):30–30. DOI: 10.1007/s13278-022-00858-z.

Allaway, E. and McKeown, K. R. (2020). Zero-shot stance detection: A dataset and model using generalized topic representations. In *EMNLP-2020 proceedings*, pages 8913–8931, Online. Assoc. for Computational Linguistics. DOI: 10.18653/v1/2020.emnlp-main.717.

Allaway, E. and McKeown, K. R. (2022). Zero-shot stance detection: Paradigms and challenges. *Frontiers in Artificial Intelligence*, 5:1070429. DOI: 10.3389/frai.2022.1070429.

Bahuleyan, H. and Vechtomova, O. (2017). UWaterloo at SemEval-2017 task 8: Detecting stance towards rumours with topic independent features. In *11th intl. workshop on semantic evaluation (SemEval-2017)*, pages 461–464. DOI: 10.18653/v1/S17-2080.

Blackburn, M., Yu, N., Berrie, J., Gordon, B., Longfellow, D., Tirrell, W., and Williams, M. (2020). Corpus development for studying online disinformation campaign: A narrative + stance approach. In *First International Workshop on Social Threats in Online Conversations: Understanding and Management*, pages 41–47. Avaialble at:`https://aclanthology.org/2020.stoc-1.7/`.

Boulahia, S. Y., Amamra, A., Madi, M. R., and Daikh, S. (2021). Early, intermediate and late fusion strategies for robust deep learning-based multimodal action recognition. *Mach. Vision Appl.*, 32(6). DOI: 10.1007/s00138-021-01249-8.

Cavalheiro, L. C. L., Pavan, M. C., and Paraboni, I. (2023). Stance prediction from multimodal social media data. In *14th International Conference on Recent Advances in Natural Language Processing*, pages 242–248. DOI: 10.26615/978-954-452-092-$2_0$27.

Chancellor, S. and De Choudhury, M. (2020). Methods in predictive techniques for mental health status on social media: a critical review. *NPJ digital medicine*, 3(1):43. DOI: 10.1038/s41746-020-0233-7.

Chang, W., Li, J., and Lee, C. (2019). Learning semantic-preserving space using user profile and multimodal media content from political social network. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3990–3994. IEEE. DOI: 10.1109/icassp.2019.8682596.

Chen, W. and Ku, L. (2018). We like, we post: A joint user-post approach for facebook post stance labeling. *IEEE Transactions on Knowledge and Data Engineering*, 30:2013–2023. DOI: 10.1109/tkde.2018.2810875.

Cignarella, A. T., Lai, M., Bosco, C., Patti, V., Paolo, R., *et al*. (2020). Sardistance@ evalita2020: Overview of the task on stance detection in italian tweets. In *7th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020)*. CEURWS.org. DOI: 10.4000/books.aaccademia.7084.

Costa, P. B., Pavan, M. C., Santos, W. R., Silva, S. C., and Paraboni, I. (2023). BERTabaporu: Assessing a genre-specific language model for Portuguese NLP. In Mitkov, R. and Angelova, G., editors, *14th International Conference on Recent Advances in Natural Language Processing*, pages 217–223, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria. Available at:`https://aclanthology.org/2023.ranlp-1.24/`.

Custódio, J. E. and Paraboni, I. (2021). Stacked authorship attribution of digital texts. *Expert Systems with Applications*, 176:114866. DOI: 10.1016/j.eswa.2021.114866.

Darwish, K., Magdy, W., and Zanouda, T. (2017). Improved stance prediction in a user similarity feature space. In *IEEE/ACM International conference on advances in social networks analysis and mining 2017*, pages 145–148. DOI: 10.1145/3110025.3110112.

de Oliveira, R. L. (2022). Detecção de posicionamento em tweets sobre Covid-19 no Brasil utilizando métodos de aprendizagem de máquina. Master's thesis, Universidade Federal de Pernambuco. Availablet at:https://bdtd.ibict.br/vufind/Record/UFPE$_6$6b0c6f870fb3ecfe35dfc7b1c1ad1f79

de Souza, V. B., Nobre, J. C., and Becker, K. (2022). DAC stacking: A deep learning ensemble to classify anxiety, depression, and their comorbidity from Reddit texts. *IEEE Journal of Biomedical and Health Informatics*, 26(7):3303–3311. DOI: 10.1109/jbhi.2022.3151589.

de Vinco, D., Antelmi, A., Spagnuolo, C., and Aiello, L. M. (2024). Deciphering Conversational Networks: Stance Detection via Hypergraphs and LLMs. In *Companion Publication of the 16th ACM Web Science Conference*, page 3–4. DOI: 10.1145/3630744.3658418.

Derczynski, L., Bontcheva, K., Liakata, M., Procter, R., Hoi, G. W. S., and Zubiaga, A. (2017). Semeval-2017 task 8: Rumoureval: Determining rumour veracity and support for rumours. *arXiv:1704.05972*. DOI: 10.18653/v1/s17-2006.

Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2019). BERT: pre-training of deep bidirectional transformers for language understanding. In Burstein, J., Doran, C., and Solorio, T., editors, *Conference of the North American Chapter of the Association for Computational Linguistics*, pages 4171–4186. Association for Computational Linguistics. DOI: 10.48550/arXiv.1810.04805.

Dong, R., Sun, Y., Wang, L., Gu, Y., and Zhong, Y. (2017). Weakly-guided user stance prediction via joint modeling of content and social interaction. In *2017 ACM on Conference on Information and Knowledge Management*, pages 1249–1258. DOI: 10.1145/3132847.3133020.

dos Santos, P. D. and Goya, D. H. (2023). Detecção de Posicionamento e Rotulação Automática de Usuários do Twitter: o caso da CPI da Covid-19. *iSys-Brazilian Journal of Information Systems*, 16:15–1. DOI: 10.5753/brasnam.2022.223212.

dos Santos, V. G. and Paraboni, I. (2022). Myers-briggs personality classification from social media text using pre-trained language models. *Journal of Universal Computer Science*, 28(4):378–395. DOI: 10.3897/jucs.70941.

Dutta, S., Caur, S., Chakrabarti, S., and Chakraborty, T. (2022). Semi-supervised stance detection of tweets via distant network supervision. In *15th ACM International conf. on Web search and data mining*, pages 241–251. DOI: 10.1145/3488560.3498511.

Ebeling, R., Sáenz, C. A. C., Nobre, J., and Becker, K. (2020). Quarenteners vs. chloroquiners: A framework to analyze how political polarization affects the behavior of groups. In *2020 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, pages 203–210. IEEE. DOI: 10.1109/wi-iat50758.2020.00031.

Espinosa, M. S., Agerri, R., Rodrigo, A., and Centeno, R. (2020). Deepreading@ sardistance: Combining textual, social and emotional features. In *7th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020)*. CEURWS.org. DOI: 10.4000/books.aaccademia.7129.

Ferraccioli, F., Sciandra, A., Da Pont, M., Girardi, P., Solari, D., and Finos, L. (2020). Textwiller@ sardistance, haspeede2: Text or con-text? a smart use of social network data in predicting polarization. In *7th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020)*. CEURWS.org. DOI: 10.4000/books.aaccademia.7152.

Flores, A. M., Pavan, M. C., and Paraboni, I. (2022). User profiling and satisfaction inference in public information access services. *Journal of Intelligent Information Systems*, 58(1):67–89. DOI: 10.1007/s10844-021-00661-w.

Fraisier, O., Cabanac, G., Pitarch, Y., Besancon, R., and Boughanem, M. (2018). Stance classification through proximity-based community detection. In *29th on Hypertext and Social Media*, pages 220–228. DOI: 10.1145/3209542.3209549.

Geiss, H.-J., Sakketou, F., and Flek, L. (2022). OK boomer: Probing the socio-demographic divide in echo chambers. In *Tenth International Workshop on Natural Language Processing for Social Media*, pages 83–105. DOI: 10.18653/v1/2022.socialnlp-1.8.

Gorrell, G., Kochkina, E., Liakata, M., Aker, A., Zubiaga, A., Bontcheva, K., and Derczynski, L. (2019). SemEval-2019 task 7: RumourEval, determining rumour veracity and support for rumours. In *13th International Workshop on Semantic Evaluation*, pages 845–854, Minneapolis, Minnesota, USA. Association for Computational Linguistics. DOI: 10.18653/v1/S19-2147.

Graells-Garrido, E., Baeza-Yates, R., and Lalmas, M. (2020). Every colour you are: Stance prediction and turnaround in controversial issues. In *12th ACM Conference on Web Science*, pages 174–183. DOI: 10.1145/3394231.3397907.

Grover, A. and Leskovec, J. (2016). node2vec: Scalable Feature Learning for Networks. In *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 855–864, San Francisco, USA. Association for Computing Machinery. DOI: 10.1145/2939672.2939754.

Hardalov, M., Arora, A., Nakov, P., and Augenstein, I. (2022). A survey on stance detection for mis- and disinformation identification. *arXiv:2103.00242*. DOI: 10.18653/v1/2022.findings-naacl.94.

Igarashi, Y., Komatsu, H., Kobayashi, S., Okazaki, N., and Inui, K. (2016). Tohoku at SemEval-2016 task 6: Feature-based model versus convolutional neural network for stance detection. In *10th international workshop on sematic evaluation (SemEval-2016)*, pages 401–407. DOI: 10.18653/v1/S16-1065.

Islam, M. R., Muthiah, S., and Ramakrishnan, N. (2019). Rumorsleuth: Joint detection of rumor veracity and user stance. In *2019 IEEE/ACM intl. conf. on advances in social networks analysis and mining*, pages 131–136. DOI: 10.1145/3341161.3342916.

Jamialahmadi, S., Sahebi, I., Sabermahani, M. M., Shariatpanahi, S. P., Dadlani, A., and Maham, B. (2022). Rumor stance classification in online social networks: the state-of-the-art, prospects, and future challenges. *IEEE Access*, 10:113131–113148. DOI: 10.1109/access.2022.3216835.

Jia, P., Du, Y., Lyu, B., and Hu, R. (2021). Stance detection using multi-head attention based bidirectional gru. In *7th International Conference on Computer and Communications (ICCC)*, pages 625–630. IEEE. DOI: 10.1109/iccc54389.2021.9674443.

Khiabani, P. J. and Zubiaga, A. (2024). Cross-target stance detection: A survey of techniques, datasets, and challenges. *arXiv:2409.13594*. DOI: 10.1016/j.eswa.2025.127790.

Kumar, A., Mishra, D., and Das, B. (2020). Twitter as a mirror - perspectives of common men and key personalities. In *2020 International Conference on Emerging Trends in Information Technology and Engineering (ic-ETITE)*, pages 1–5. IEEE. DOI: 10.1109/ic-etite47903.2020.436.

Kuo, K.-H., Wang, M.-H., Kao, H.-Y., and Dai, Y.-C. (2024). Advancing stance detection of political fan pages: A multimodal approach. In *ACM on Web Conference*, page 702–705. DOI: 10.1145/3589335.3651467.

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

Lai, M., Farías, D. I. H., Patti, V., and Rosso, P. (2017). Friends and enemies of clinton and trump: Using context for detecting stance in political tweets. In *15th Mexican International Conference on Artificial Intelligence, MICAI 2016*, pages 155–168, Cancún, Mexico. Springer. DOI: 10.1007/978-3-319-62434-$1_1$3.

Lai, M., Patti, V., Ruffo, G., and Rosso, P. (2020b). Brexit: Leave or remain? the role of user's community and diachronic evolution on stance detection. *Journal of Intelligent & Fuzzy Systems*, 39(2):2341–2352. DOI: 10.3233/jifs-179895.

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

Li, Y. and Scarton, C. (2020). Revisiting rumour stance classification: Dealing with imbalanced data. In *3rd international workshop on rumours and deception in social media (RDSM)*, pages 38–44. Avaialble at:`https://aclanthology.org/2020.rdsm-1.4/`.

Liu, X., Wang, R., Sun, D., Li, J., Youn, C., Lyu, Y., Zhan, J., Wu, D., Xu, X., Liu, M., Lei, X., Xu, Z., Zhang, Y., Li, Z., Yang, Q., and Abdelzaher, T. (2023). Influence pathway discovery on social media. In *2023 IEEE 9th International Conference on Collaboration and Internet Computing (CIC)*, pages 105–109. DOI: 10.1109/cic58953.2023.00023.

Luo, Y., Ma, J., and Yeo, C. K. (2022). Identification of rumour stances by considering network topology and social media comments. *Journal of Information Science*, 48(1):118–130. DOI: 10.1177/0165551520944352.

Lynn, V., Giorgi, S., Balasubramanian, N., and Schwartz, H. A. (2019). Tweet classification without the tweet: An empirical examination of user versus document attributes. In *3rd workshop on natural language processing and computational social science*, pages 18–28. DOI: 10.18653/v1/w19-2103.

Ma, J., Gao, W., and Wong, K.-F. (2018). Detect rumor and stance jointly by neural multi-task learning. In *Companion proceedings of the the web conference 2018*, pages 585–593. DOI: 10.1145/3184558.3188729.

Magdy, W., Darwish, K., Abokhodair, N., Rahimi, A., and Baldwin, T. (2016). #isisisnotislam or #deportallmuslims? predicting unspoken views. In *8th ACM Conference on Web Science*, pages 95–106, Hannover, Germany. Association for Computing Machinery. DOI: 10.1145/2908131.2908150.

Masood, M. A. and Abbasi, R. A. (2021). Using graph embedding and machine learning to identify rebels on twitter. *Journal of Infometrics*, 15(1):101121. DOI: 10.1016/j.joi.2020.101121.

McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444. DOI: 10.1146/annurev.soc.27.1.415.

Mohammad, S., Kiritchenko, S., Sobhani, P., Zhu, X., and Cherry, C. (2016). Semeval-2016 task 6: Detecting stance in tweets. In *10th international workshop on semantic evaluation (SemEval-2016)*, pages 31–41, San Diego, California. Association for Computational Linguistics. DOI: 10.18653/v1/s16-1003.

Onikoyi, B., Nnamoko, N., and Korkontzelos, I. (2023). Gender prediction with descriptive textual data using a machine learning approach. *Natural Language Processing Journal*, 4:100018. DOI: 10.1016/j.nlp.2023.100018.

Pavan, M. C. and Paraboni, I. (2022). Cross-target stance classification as domain adaptation. In *Advances in Computational Intelligence - MICAI 2022 - Lecture Notes in Artificial Intelligence vol 13612*, pages 15–25, Cham. Springer Nature Switzerland. DOI: 10.1007/978-3-031-19493-1_2.

Pavan, M. C. and Paraboni, I. (2024). A benchmark for portuguese zero-shot stance detection. *Journal of the Brazilian Computer Society*, 30(1):469–479. DOI: 10.5753/jbcs.2024.3932.

Pennebaker, J. W., Francis, M. E., and Booth, R. J. (2001). *Inquiry and Word Count: LIWC*. Lawrence Erlbaum, Mahwah, NJ. Book.

Penzo, N., Longa, A., Lepri, B., Tonelli, S., and Guerini, M. (2024). Putting context in context: the impact of discussion structure on text classification. In *18th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1793–1811. DOI: 10.18653/v1/2024.eacl-long.108.

Pereira, C., Pavan, M., Yoon, S., Ramos, R., Costa, P., Cavalheiro, L., and Paraboni, I. (2026). UstanceBR: a social media language resource for stance prediction. *Language Resources and Evaluation*, 60(14). DOI: 10.1007/s10579-025-09896-3.

Pougué-Biyong, J., Gupta, A., Haghighi, A., and El-Kishky, A. (2023). Learning stance embeddings from signed social graphs. In *16th ACM International Conference on Web Search and Data Mining*, page 177–185. DOI: 10.48550/arXiv.2201.11675.

Rezayi, S., Soleymani, S., Arabnia, H. R., and Li, S. (2021). Socially aware multimodal deep neural networks for fake news classification. In *IEEE 4th international conference on multimedia information processing and retrieval (MIPR)*, pages 253–259. IEEE. DOI: 10.1109/mipr51284.2021.00048.

Rochert, D., Neubaum, G., Ross, B., and Stieglitz, S. (2022). Caught in a networked collusion? homogeneity in conspiracy-related discussion networks on youtube. *Information Systems*, 103:101866. DOI: 10.1016/j.is.2021.101866.

Sadiq, S., Yan, Y., Taylor, A., Shyu, M. L., Chen, S. C., and Feaster, D. (2017). Aafa: Associative affinity factor analysis for bot detection and stance classification in twitter. In *2017 IEEE International Conference on Information Reuse and Integratiin (IRI)*, pages 356–365. DOI: 10.1109/iri.2017.25.

Sutter, M., Gourru, A., Trabelsi, A., and Largeron, C. (2024). Unsupervised stance detection for social media discussions: A generic baseline. In *18th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1782–1792. DOI: 10.18653/v1/2024.eacl-long.107.

Tyagi, A., Uyheng, J., and Carley, K. M. (2020). Affective polarization in online climate change discourse on twitter. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 443–447. IEEE. DOI: 10.1109/asonam49781.2020.9381419.

Tyagi, A., Uyheng, J., and Carley, K. M. (2021). Heated conversations in a warming world: affective polarization in online climate change discourse follows real-world climate anomalies. *Social Network Analysis and Mining*, 11:1–12. DOI: 10.1007/s13278-021-00792-6.

Vanta, T. and Aono, M. (2020). Stance classification and rumor analysis: Using new dialog-act features and augmenting input tweets. In *7th International Conference on Advance Informatics: Concepts, Theory and Applications (ICAICTA)*, pages 1–6. IEEE. DOI: 10.1109/icaicta49861.2020.9429036.

Veyseh, A. P. B., Ebrahimi, J., Dou, D., and Lowd, D. (2017). A temporal attentional model for rumor stance classification. In *2017 ACM on Conference on Information and Knowledge Management*, pages 2335–2338. DOI: 10.1145/3132847.3133116.

Vilella, S., Lai, M., Paolotti, D., and Ruffo, G. (2020). Immigration as a divisive topic: Clusters and content diffusion in the italian twitter debate. *Future internet*, 12(10):173. DOI: 10.3390/fi12100173.

Wang, R., Zhou, D., Jiang, M., Si, J., and Yang, Y. (2019). A survey on opinion mining: From stance to product aspect. *IEEE Access*, 7:41101–41124. DOI: 10.1109/access.2019.2906754.

Williams, E. M. and Carley, K. M. (2023). Tspa: Efficient target-stance detection on twitter. In *2022 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, page 242–246. DOI: 10.1109/asonam55673.2022.10068608.

Wolpert, D. H. (1992). Stacked generalization. *Neural networks*, 5(2):241–259. DOI: 10.1016/s0893-6080(05)80023-1.

Xuan, K. and Xia, R. (2019). Rumor stance classification via machine learning with text, user and propagation features. In *International Conference on Data Mining Workshops (ICDMW)*, pages 560–566. IEEE. DOI: 10.1109/icdmw.2019.00085.

Yang, C., Li, J., Wang, R., Yao, S., Shao, H., Liu, D., Liu, S., Wang, T., and Abdelzaher, T. F. (2020). Hierarchical overlapping belief estimation by structured matrix factorization. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 81–88. IEEE. DOI: 10.1109/asonam49781.2020.9381477.

Zhang, B., Dai, G., Niu, F., Yin, N., Fan, X., Wang, S., Cao, X., and Huang, H. (2024). A survey of stance detection on social media: New directions and perspectives. *arXiv:2409.15690*. DOI: 10.48550/arxiv.2409.15690.

Zhao, C. and Caragea, C. (2024). EZ-STANCE: A large dataset for English zero-shot stance detection. In *62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15697–15714, Bangkok, Thailand. Association for Computational Linguistics. DOI: 10.18653/v1/2024.acl-long.838.

Zubiaga, A., Aker, A., Bontcheva, K., Liakata, M., and Procter, R. (2018a). Detection and resolution of rumours in social media: A survey. *Acm Computing Surveys (Csur)*, 51(2):1–36. DOI: 10.1145/3161603.

Zubiaga, A., Kochkina, E., Liakata, M., Procter, R., Lukasik, M., Bontcheva, K., Cohn, T., and Augenstein, I. (2018b). Discourse-aware rumour stance classification in social media using sequential classifiers. *Information Processing & Management*, 54(2):273–290. DOI: 10.1016/j.ipm.2017.11.009.