IRIS AperTO

UNIVERSITÀ DEGLI STUDI DI TORINO

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Friends and Enemies of Clinton and Trump: Using Context for Detecting Stance in Political Tweets

This is the author's manuscript

Original Citation:

Availability:

This version is available http://hdl.handle.net/2318/1652427 since 2017-11-21T17:01:08Z

Publisher:

Springer

Published version:

DOI:10.1007/978-3-319-62434-1_13

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

13 October 2026

Friends and Enemies of Clinton and Trump: Using Context for Detecting Stance in Political Tweets

Mirko Lai[1,2], Delia Irazú Hernández Farías[1,2], Viviana Patti[1] and Paolo Rosso[2]

[1] Dipartimento di Informatica, Università degli Studi di Torino, Torino, Italy
[2] Pattern Recognition and Human Language Technologies Research Center, Universitat Politécnica de Valéncia, Valencia, Spain

Abstract. Stance detection, the task of identifying the speaker's opinion towards a particular target, has attracted the attention of researchers. This paper describes a novel approach for detecting stance in Twitter. We define a set of features in order to consider the context surrounding a target of interest with the final aim of training a model for predicting the stance towards the mentioned targets. In particular, we are interested in investigating political debates in social media. For this reason we evaluated our approach focusing on two targets of the SemEval-2016 Task 6 on Detecting stance in tweets, which are related to the political campaign for the 2016 U.S. presidential elections: Hillary Clinton vs. Donald Trump. For the sake of comparison with the state of the art, we evaluated our model against the dataset released in the SemEval-2016 Task 6 shared task competition. Our results outperform the best ones obtained by participating teams, and show that information about enemies and friends of politicians help in detecting stance towards them.

1 Introduction

Social media provide a way for expressing opinions about different topics. From this kind of user-generated content it is possible to discover relevant information under several perspectives. A wide range of research has been carried out in order to exploit the vast amount of data generated in social media. One of the most interesting research areas concerns to investigate how people expose their feelings, evaluations, attitudes and emotions. These kinds of aspects are the subject of interest of Sentiment Analysis (SA) [1].

Determining the subjective value of a piece of text is the most general task of SA. Recently, the interest on studying finer-grained and different facets of sentiment in texts has derived in areas such as *Aspect based sentiment analysis* [2] and *Stance Detection* (SD) [3], which is the focus of our work. Identifying the speaker's opinion towards a particular target is the main goal of SD. It is not enough to recognize whether or not a text is positive/negative/neutral but it is necessary to infer the point of view of the tweeter towards a particular target.

Stance detection could not only provide useful information for improving the performance of SA but it could also help to better understand the way in which

people communicate ideas in order to highlight their point of view towards a particular target entity. This is particularly interesting when the target entity is controversial issue (e.g., political reforms [4, 5]) or a polarizing person (e.g., candidates in political elections). Therefore, detecting stance in social media could become a helpful tool for various sectors of society, such as journalism, companies and government, having politics as an especially good application domain. Several efforts have been made in order to investigate different aspects related to social media and politics [6]. We are interested in political debates in social media, particularly in the interaction between polarized communities. We consider that being able to detect stance in user-generated content could provide useful insights to discover novel information about social network structures. Political debate texts coming from social media where people discuss their different points of view offer an attractive information source.

This year, for the first time a shared task on stance detection in tweets was organized [3]. Two of the targets considered in order to evaluate stance detection systems were: Hillary Clinton and Donald Trump[3]. Both targets have been the focus of different research, for instance in [7] the authors studied their speeches during the 2016 political campaign. In such way, studying these targets is an attracting topic of research due to the impact of the use of social media during the political campaign for the 2016 U.S. Presidential elections.

Our approach to detect stance in tweets relies mainly on the context of the targets of interest: Hillary Clinton and Donald Trump. Besides, we also took advantage of widely used features in SA.

The paper is organized as follows. Section 2 introduces the first shared task on Twitter stance detection. Section 3 describes our method to detect stance by exploiting different features. Section 4 describes the evaluation and results. Finally, Section 5 draws some conclusions.

2 Detecting Stance on Tweets

The SemEval-2016 Task 6: Detecting Stance in Tweets[4] was the first shared task on detecting stance from tweets. Mohammad et. al in [3] describe the task as: *Given a tweet text and a target entity (person, organization, movement, policy, etc.), automatic natural language systems must determine whether the tweeter is in favor of the target, against the given target, or whether inference is likely.*

Let us to introduce the following example[5]:
Support #independent #BernieSanders because he's not a liar. #POTUS #libcrib #democrats #tlot #republicans #WakeUpAmerica #SemST

The target of interest is "Hillary Clinton". Here, the tweeter expresses a positive opinion towards an adversary of the target. Consequently the annotator

[3] They are the candidates who won the Party Presidential Primaries for the Democratic and Republican parties, respectively.
[4] http://alt.qcri.org/semeval2016/task6/
[5] This tweet was extracted from the training set of SemEval-2016 Task 6.

inferred that the tweeter expresses a negative opinion towards the target. As can be noticed, this tweet does not contain any explicit clue to find the target.

For evaluating the task, the organizers annotated near to 5,000 English tweets for stance towards six commonly known targets in the United States: "Atheism", "Climate Change is a Real Concern", "Feminism Movement", "Hillary Clinton", "Legalization of Abortion", and "Donald Trump" (Stance Dataset, henceforth). A set of hashtags widely used by people when tweeting about these targets was compiled; then it was used to retrieve tweets according three categories: in-favor hashtags, against hashtags and stance-ambiguous hashtags. The tweets were manually annotated by crowdsourcing. More details about the Stance Dataset can be found in [3].

The participants in the SemEval-2016 Task 6 were required to classify tweet-target pairs into exactly one of three classes: *Favor*: It can be inferred from the tweet that the tweeter supports the target (e.g., directly or indirectly by supporting someone/something, by opposing or criticizing someone/something opposed to the target, or by echoing the stance of somebody else); *Against*: It can be inferred from the tweet that the tweeter is against the target (e.g., directly or indirectly by opposing or criticizing someone/something, by supporting someone/something opposed to the target, or by echoing the stance of somebody else); and *Neither*: None of the above.

The SemEval-2016 Task 6 was divided into two subtasks:

- Task A. Supervised Framework. The participating systems were asked to perform stance detection towards the following targets: "Atheism", "Climate Change is a Real Concern", "Feminism Movement", "Hillary Clinton", and "Legalization of Abortion". For evaluation the organizer provided a training (2,914 tweets) and test (1,249 tweets) sets.
- Task B. Weakly Supervised Framework. The task was detecting stance towards one target "Donald Trump" in 707 tweets. For this task the participants were not provided with any training data about this target.

Nineteen teams participated in Task A while only nine competed in Task B. It is important to highlight that only two systems were evaluated specifically on Task B. Figure 1 shows a brief summary of the systems. Further information about the systems in the task can be found in [8][6].

Both tasks were addressed in similar ways. Most teams exploited standard text classification features such as n-grams and word embedding vectors. Besides, some SA features from well-known lexical resources, such as *EmoLex* [9], *MPQA* [10], *Hu and Liu* [11] and *NRC Hashtag* [12], were used to detect stance in tweets. Furthermore, some teams decided to take advantage of additional data by harvesting Twitter using stance-bearing hashtags in order to have more stance tweets. It is important to highlight that the best system in Task A (MITRE) did use this alternative. A similar approach was adopted by the three best ranked systems on Task B (pkudblab, LitisMind, and INF-UFRGS). For what concerns

[6] Notice that not all the reports describing systems and approaches of teams participating at SemEval-2016 Task 6 are available in [8].

to Task B, in order to deal with the lack of training data, some systems attempted to generalize the supervised data from task A in different ways such as defining rules or by exploiting multi-stage classifiers.

3 Our approach

We are proposing a supervised approach for stance detection [7]. Our work is focused on detecting stance towards Hillary Clinton and Donald Trump that are currently contesting the political campaign for the 2016 U.S. Presidential election. An important aspect to mention concerns to the fact that when the Stance Dataset was built the two targets were still participating to the Party Presidential Primaries for the Democratic and Republican parties, respectively. We address the stance detection in tweets, casting it as a classification task. A set of features that comprises different aspects was exploited. The most novel one refers to the extraction of context-related information regarding to the target of interest. Our hypothesis is that domain knowledge could provide useful information to improve the performance of SD systems. For instance, in order to correctly identify stance in a tweet as the one mentioned in Section 2, it is needed to recognize that *Bernie Sander* was an adversary of Hillary Clinton during the Party Presidential Primaries of the Democratic party. Attempting to capture information related to domain knowledge, we define two concepts: "enemies" and "friends". These concepts are used for denoting the entities related to the target. By using the terms "enemies" and "friends", we are trying to infer that when a tweeter is against an "enemy"/"friend" of the target, then the tweeter is in favor/against towards the target and, on the other hand, when a tweeter is in favor towards an "enemy"/"friend" of the target, then the tweeter is against/in favor towards the target. Figure 1 shows an example of the relationships between the "friends" and "enemies" according to their political party, in this case the target of interest is Hillary Clinton.
Three groups of features were considered: sentiment, structural, and context-based.

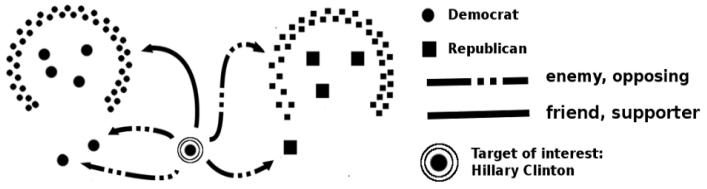

Fig. 1. Diagram of relationships between friends and enemies of Hillary Clinton

[7] https://github.com/mirkolai/Friends-and-Enemies-of-Clinton-and-Trump

Table 1. Brief description of the participating systems at SemEval-2016 Task 6

System	Description
MITRE [13] Task A	Overall approach: Recurrent neural networks. External resources: Words embeddings with the word2vect skip-gram method. Near to 300,000 tweets containing hashtags related to the targets.
pkudblab [14] Tasks A and B	Overall approach: Convolutional neural network. External resources: Words embeddings using the Google News dataset.
TakeLab [15] Task A	Overall approach: An ensemble of learning algorithms (such as SVM, random forest) fine-tuned using a genetic algorithm. External resources: Word features, word embeddings, frequency of emoticons, uppercase characters, among others.
ECNU [16] Tasks A and B	Overall approach: A pipeline-based procedure involving relevance and orientation detection. External resources: N-grams, topic features and sentiment lexicon features (such as Hu&Liu and MPQA, among others).
CU-GWU [17] Task A	Overall approach: Classification using SVM External resources: N-grams, Stanford's SA system and LIWC.
IUCL-RF [18] Task A	Overall approach: Classification algorithms (SVM, random forest, gradient boosting decision trees) and an ensemble classifier (TiMBL). External resources: Bag-of-Words and word vectors.
DeepStance [19] Task A	Overall approach: A set of naive bayes classifiers using deep learning. External resources: More than 1.5 million of tweets were added by using representative hashtag for target-stance pairs.
UWB [20] Tasks A and B	Overall approach: Maximum entropy classifier. External resources: N-grams, PoS labels, General Inquirer. Additional tweets were gathered based on frequent hashtags in the training set.
IDI@NTNU [21] Task A	Overall approach: A soft voting classifier approach (naive bayes and logistic regression). External resources: Word vectors, n-grams, char-grams, negation, punctuation marks, elongated words, among others.
Tohoku [22] Task A	Overall approach: Two methods: a feature based approach and a neural network based approach. External resources: Bag-of-Words, PoS labels, SentiWordNet. Additional Twitter data was gathered from target words.
ltl.uni-due [23] Tasks A and B	Overall approach: Multidimensional classification problem External resources: N-grams, punctuation marks, negation, nouns.
JU_NLP [24] Task A	Overall approach: Classification using SVM External resources: N-Gram and sentiment analysis resources such as: SentiWordNet, EmoLex and NRC Hashtag Emotion Lexicon.
nldsucsc [25] Task A	Overall approach: Classification using SVM, J48 and naive bayes. External resources: N-grams, PoS labels, LIWC. Additional tweets were gathered based on frequent hashtags in the training set.
INF_UFRGS [26] Task B	Overall approach: Set of rules together with SVM. External resources: N-grams.
USFD [27] Task B	Overall approach: Classification using logistic regression. External resources: Bag-of-words autoencoder. Additional tweets were gathered by using two keywords per target.

Sentiment-based Features

We shared the idea that stance detection is strongly related to sentiment analysis [3, 16]. As far as we know, there are not sentiment analysis lexica retrieved specifically in the political domain[8]; thus, in order to take advantage of sentiment features it is possible to exploit the wide range of resources available for English. We used a set of four lexica to cover different facets of affect ranging from prior polarity of words to fine-grained emotional information:

- **AFINN**. It is an affective lexicon of 2,477 English words manually labeled with a polarity value between -5 to +5. AFINN was collected by Finn Årup Nielsen [28]. We consider one feature from AFINN: the sum of the polarity of the words present in each tweet.

- **Hu&Liu (HL)**. It includes about 6,800 positive and negative words. We calculate the difference between the positive and negative words in a tweet as a feature.

- **LIWC**. The Linguistic Inquiry and Word Counts (LIWC) [29] is a dictionary that contains about 4,500 entries distributed in 64 categories that can be further used to analyse psycholinguistic features in texts. We calculate the difference between PosEmo (with 405 entries) and NegEmo (with 500 entries) categories in a tweet as a feature.

- **DAL**. The Dictionary of Affect in Language (DAL) contains 8,742 English words; it was developed by Whissell [30]. Each word is rated in a three-point scale into three dimensions: Pleasantness (It refers to the degree of pleasure produced by words), Activation (It refers the degree of response that humans have under an emotional state) and Imagery (It refers to how difficult to form a mental picture of a given word is). We consider six features, i.e. the sum and the mean of the rates of the words present in the tweet for each one of the three dimensions.

Structural Features

We also explore structural characteristics of tweets because we believe that could be useful to detect stance. We experimented with several kinds of structural features, however only the most relevant ones were included in the final approach:

- **Hashtags**. The frequency of hashtags present in each tweet.

- **Mentions**. The frequency of screen names (often called mentions) in each tweet.

- **Punctuation marks (punct_marks)**. We consider a set of 6 different features: the frequency of exclamation marks, question marks, periods, commas, semicolons, and finally the sum of all the punctuation marks mentioned before.

[8] For example, the term *vote* is strongly related to politics, but it is not present in commonly used SA lexica such as: AFINN, Hu&Liu, and LIWC.

Context-based Features

Our hypothesis is that the context-based features should capture some domain-related information. An overall perspective of the context surrounding a target can be acquired by the relationships that exist between the target and other entities in its domain. As mentioned before we are interested in investigating Political debates: for this reason we selected as targets of interest politicians such as Hillary Clinton and Donald Trump. We manually created a list of entities related to the Party Presidential Primaries for the Democratic and Republican parties from Wikipedia[9]. We exploited 6 types of context-based features considering different kinds of relationships between the target and the entities around the target:

- **Target of interest mentioned by name (targetByName)**: This feature captures the presence of the target of interest in the tweet in hand. *#Stop**Hillary**2016 **HillaryClinton** if there was a woman with integrity and honesty I would vote for such as woman president, NO.* The list of tokens used to check the presence of the target of interest are: *hillaryclinton*, *hillary*, *clinton*, and *hill* for Hillary Clinton; while for Donald Trump are *realdonaldtrump*, *donald*, and *trump*.

- **Target of interest mentioned by pronoun (targetByPronoun)**: This feature allows to identify those cases when the target of interest is mentioned by using a pronoun. In the following example, knowing that the target of the tweet is Hillary Clinton, it is possible to exploit the pronoun "she" to capture the presence of the target in hand.
  *HomeOfUncleSam ScotsFyre RWNutjob1 SA_Hartdegen **She**'s too old to understand the internet...that **she** can be fact checked.*
  Two pronouns were considered for each one of the targets of interest: *she* and *her* for Hillary Clinton, while *he* and *his* for Donald Trump.

- **Target's party (targetParty)**: As people involved in politics, our targets belong to a political party. Using this feature we identify if the stance against (or in favor) towards the target of interest was expressed mentioning the name of the party instead of the target. In the following example the tweeter expresses a negative opinion toward Hillary Clinton party.
  *It's a miracle, suddenly #**Democrats** don't mind having someone who voted for war.*
  In this case we consider the tokens *dem*, *democratic*, *democrat*, *democrats*, *progressive* in order to check the entity party for Hillary Clinton, while we consider the tokens *republican*, *republicans*, and *conservative* for Donald Trump.

- **Party colleague opposite (targetPartyColleagues)**: We also considered the case where the party colleagues of the target of interest are mentioned

[9] Articles: *Democratic Party presidential primaries, 2016* and *Republican Party presidential candidates, 2016*

to express an opinion towards it. We use the name and the surnames of the candidates for the Party Presidential Primaries for both Democratic and Republican parties. In the example, Hillary Clinton's party colleagues are mentioned. *msnbc* ***Lawrence*** *JoeBiden Sen****Sanders*** *we love Joe and* ***Bernie****–but they ARE too OLD–they would end up a #OneTerm President #SemST*
The list of names used for Hillary Clinton is: *bernie*, *sanders*, *martin*, *o'malley*, *lincoln*, *chafee*, *webb*, *lawrence*, and *lessig*; while for Donald Trump is: *ted*, *cruz*, *marco*, *rubio*, *john*, *kasich*, *ben*, *carson*, *jeb*, *bush*, *rand*, *paul*, *mike*, *huckabee*, *carly*, *fiorina*, *chris*, *christie*, *rick*, *santorum*, *gilmore*, *rick*, *perry*, *scott*, *walker*, *bobby*, *jindal*, *lindsey*, *graham*, *george*, *pataki*.

– **Target's oppositors party (targetsOppositors)**: This feature captures the presence of oppositors belonging to the rival party of target of interest's. In the following example a positive opinion is expressed towards two candidates from the Republican party. Thus, the tweet is against Hillary Clinton.
*PhilGlutting megadreamin Thank you so much for RT and FAV!!! #WakeUpAmerica #****Rubio****2016 #****Cruz****2016 #SemST*
We use the Donald Trump's tokens lists targetParty and targetPartyColleagues in order to create Hillary Clinton's targetsOppositors tokens list, while we use Hillary Clinton's tokens lists targetParty and targetPartyColleagues in order to create Donald Trump's targetsOppositors tokens list.

– **Nobody (nobody)**: This feature allows to catch those cases where any of the above described entities are mentioned in a tweet. In the following example the term *Ambassador* refers to Chris Stevens, who served as the U.S. Ambassador to Libya and who was killed at Bengasi in 2012. The diplomat is related to Hillary Clinton in a situation not related with the election campaign[10].
I don't want to be appointed to an Ambassador post.
The example also shows how difficult is to infer the stance without a deep knowledge of the context.

After the evaluation of participating systems, the organizers of Semeval-2016 Task 6 annotated the Stance Datataset for sentiment and target in order to explore the relationship between sentiment and stance [31, 3][11]. In particular, tweets were manually annotated by using two additional labels: *Sentiment* and *Opinion Towards*, used to mark the overall sentiment polarity of the tweet and information about the fact that opinion is expressed directly towards the target, respectively:

– **Sentiment**. It can be positive, negative, neutral or none.
– **Opinion_target**. It can take three different values: (1) if a tweet expresses an opinion about the target; (2) if a tweet expresses an opinion related to an aspect of the target or related to something that is not the target; and (3) if there is not opinion expressed.

[10] https://en.wikipedia.org/wiki/J._Christopher_Stevens
[11] Notice that this is the first publicly available Twitter dataset annotated with both stance and sentiment.

We decided to exploit such new labels, by enriching our model with corresponding **labeled-based features**, with the aim to experiment with both context and sentiment information provided by human annotators.

4 Evaluation

We experimented with a set of tweets belonging to Hillary Clinton and Donald Trump from the Stance Dataset, the Table 2 shows the distribution of tweets annotated with stance in the training and the test set for our targets of interest.

Table 2. Distribution of stance in training and test set

Targets	% Instances in training				% Instances in test			
	Total	Against	Favor	None	Total	Against	Favor	None
Hillary Clinton	689	57.1	17.1	25.8	295	58.3	15.3	26.4
Donald Trump	-	-	-	-	707	42.3	20.9	36.8

We evaluated our approach by using the same measure defined in [3] in order to compare our results with those participating in the task. We trained a Gaussian Naive Bayes classifier [32] implemented in Scikit-learn Python library[12] to built a model for identifying stance in tweets.

We adopted two experimental setting: a) **experiment1**. It means to the use of the Sentiment-based, Structural and Context-based features; b) **experiment2**. It refers to the use of all the features described in Section 3 including the labeled-based ones. Besides, we experimented using different feature combinations in order to identify which kinds of features could be more relevant for stance detection.

Table 3. Best features combination for Hillary Clinton, and the respective results for Donald Trump with experiment1 setting

Feature set	Hillary Clinton			Donald Trump		
	F_{avg}	$F_{against}$	F_{favor}	F_{avg}	$F_{against}$	F_{favor}
mention punct_marks AFINN LIWC HL context_based	63.75	71.95	55.56	53.46	50.29	56.63
punct_marks AFINN LIWC HL context_based	62.70	71.47	53.93	52.76	49.61	55.91
hashtag punct_marks AFINN LIWC HL DAL context_based	62.3	70.43	54.17	50.44	47.69	53.19

Tables 3 and 4 present the best results obtained for Hillary Clinton in the experiment1 and experiment2, respectively. Moreover, those obtained by using

[12] http://scikit-learn.org/

Table 4. Feature set for Hillary Clinton and the respective results for Donald Trump with experiment2

Feature set	Hillary Clinton			Donald Trump		
	F_{avg}	$F_{against}$	F_{favor}	F_{avg}	$F_{against}$	F_{favor}
hashtag mention context-based labeled-based	71.21	77.17	65.26	69.59	61.99	77.19
hashtag context-based labeled-based	71.02	76.77	65.26	70.40	62.77	78.48
hashtag mention LIWC context-based labeled-based	70.98	78.23	63.73	70.20	63.06	77.35

the same set of features for Donald Trump are shown. From the results can be noted that the F1-score in "against" class is higher than in "favor". Interestingly, the opposite happens for Donald Trump. The results in Table 4 are higher than those from Table 3. Table 5 shows the best results for Donald Trump using for both experiment1 and experiment2.

Table 5. Best feature set for Donald Trump using experiment2 and experiment1 setting

Feature set	Donald Trump		
	F_{avg}	$F_{against}$	F_{favor}
* LIWC HL context_based labeld_based	74.49	69.26	79.72
mention punc_marks HL context_based	55.51	50	61.02

The * indicate the use of features belonging exclusively to experiment2.

Table 6. Results of task A and B

	Task A: Hillary Clinton		Task B: Donald Trump	
	F_{avg}	Ranking	F_{avg}	Ranking
experiment1 for Hillary Clinton	**63.75**	3	53.46	2
experiment1 for Donald Trump	61.25	4	**55.51**	2
experiment2 for Hillary Clinton	**71.21**	1	69.59	1
experiment2 for Donald Trump	68.29	1	**74.49**	1
Systems in the official competition				
INF-UFRGS	-	-	42.32	3
LitisMind	42.08	17	44.66	2
pkudbblab	64.41	2	**56.28**	1
PKULCWM	62.26	3	-	-
TakeLab	**67.12**	1	-	-

As can be noted the context-based features seem to be so relevant for both targets. Besides, it is important to highlight that the best result for each target was not achieved by the same set of features. This is maybe not surprising,

if we consider the different political campaign marketing strategies of the two candidates, which can influence also the communication of candidates' oppositors and supporters, both in terms of language register used and addressed topics. For the sake of comparison with the state of the art, we present the results obtained by the three best ranked systems at SemEval-2016 Task 6. We only include the results concerning to Hillary Clinton and Donald Trump. Both the F-measure average and the rank position of each system are included in the Table 6. We also show our best results for the two targets using both experimental settings as well as the position in the official ranking in the shared task.

Our approach achieves strongly competitive results. We ranked in the first position for both Task A and Task B using the experiment2 setting considering Hillary Clinton and Donald Trump results. For what concerns to the experiment1 we ranked in the third position for Task A and the second one for Task B. The obtained results outperform the baselines proposed in [3][13]. Besides, our outcomes outrank those obtained by submissions from all teams participating in the shared task (both task A and B). Overall, the results for Hillary Clinton are higher than those for Donald Trump. This was in someway expected, due to the lack of a training set of tweets concerning the target *Donald Trump*.

5 Conclusions

In this paper we have shown that including context-related information is crucial in order to improve the performance of stance detection systems. Experiments confirms that stance detection is highly dependent on the domain knowledge of the target in hand. Our approach relies on the presence of entities related to a target in order to try to extract the opinion expressed towards it. Besides, our proposal allows to infer the stance in both cases when the target is explicitly mentioned and also when it is not. The results obtained by exploiting context-related features outperforms those from the best ranked systems in the SemEval-2016 Task 6.

Let us highlight that we are not using either n-grams or any word-based representation, but our approach mainly relies on the context of the target in hand. We plan to investigate the performance of our approach in different domains. Exploiting semantic resources in order to catch additional context information is also an interesting line for future research. Also user's information and her social network structure could be useful. For what concerns to the sentiment-related features, overall results confirm that these kinds of features help in identifying the stance towards a particular target. We exploited different sentiment-related features, ranging from those extracted from affective resources to manually assigned polarity labels.

A further interesting matter of future work could be explore also the stance w.r.t. different aspects of a political target entity. This means to perform a sort

[13] The authors experimented with n-grams, char-grams and majority class to establish the baselines for the task.

of aspect-based sentiment analysis in a political domain, e.g., a tweeter can be in favor of Hillary for aspects related to "Health", but not for other aspects.

Finally, we think that it could be also interesting to investigate how to fruitfully combine information about stance and information about the presence of figurative devices in tweets, such as irony and sarcasm [33, 34], since the use of such devices is very frequent in political debates also in social media and detecting irony and sarcasm have been considered as one of the biggest challenges for sentiment analysis.

Acknowledgments

The National Council for Science and Technology (CONACyT Mexico) has funded the research work of Delia Irazú Hernández Farías (218109/313683). The work of Paolo Rosso has been partially funded by the SomEMBED TIN2015-71147-C2-1-P MINECO research project and by the Generalitat Valenciana under the grant ALMAMATER (PrometeoII/2014/030). The work of Viviana Patti was partially carried out at the Universitat Politècnica de València within the framework of a fellowship of the University of Turin co-funded by Fondazione CRT (World Wide Style Program 2).

References

1. Liu, B.: Sentiment analysis and opinion mining. Synthesis Lectures on Human Language Technologies **5** (2012) 1–167
2. Pontiki, M., Galanis, D., Papageorgiou, H., Androutsopoulos, I., Manandhar, S., AL-Smadi, M., Al-Ayyoub, M., Zhao, Y., Qin, B., De Clercq, O., Hoste, V., Apidianaki, M., Tannier, X., Loukachevitch, N., Kotelnikov, E., Bel, N., Jiménez-Zafra, S.M., Eryiğit, G.: Semeval-2016 task 5: Aspect based sentiment analysis. In: Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016), ACL (2016) 19–30
3. Mohammad, S., Kiritchenko, S., Sobhani, P., Zhu, X., Cherry, C.: SemEval-2016 Task 6: Detecting Stance in Tweets. In: Bethard et al. [8]. (2016) 31–41
4. Bosco, C., Lai, M., Patti, V., Virone, D.: Tweeting and being ironic in the debate about a political reform: the french annotated corpus twitter-mariagepourtous. In: Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016), ELRA (2016) 1619–1626
5. Marco Stranisci, Cristina Bosco, Delia Irazú Hernández Farías, Viviana Patti: Annotating Sentiment and Irony in the Online Italian Political Debate on #labuonascuola. In: Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016), ELRA (2016) 2892–2899
6. Maldonado, M., Sierra, V.: Twitter Predicting the 2012 US Presidential Election?: Lessons Learned from an Unconscious Value Co-Creation Platform. Journal of Organizational and End User Computing (JOEUC) **28** (2016) 10–30
7. Schumacher, E., Eskenazi, M.: A readability analysis of campaign speeches from the 2016 US presidential campaign. CoRR **abs/1603.05739** (2016)

8. Steven Bethard, Marine Carpuat, Daniel Cer, David Jurgens, Preslav Nakov, Torsten Zesch, eds.: Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016). Association for Computational Linguistics, San Diego, California (2016)
9. Mohammad, S.M., Turney, P.D.: Crowdsourcing a word–emotion association lexicon. Computational Intelligence **29** (2013) 436–465
10. Wilson, T., Wiebe, J., Hoffmann, P.: Recognizing Contextual Polarity in Phrase-level Sentiment Analysis. In: Proceedings of the Conference on HLT and Empirical Methods in Natural Language Processing, ACL (2005) 347–354
11. Hu, M., Liu, B.: Mining and Summarizing Customer Reviews. In: Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '04, Seattle, WA, USA, ACM (2004) 168–177
12. Mohammad, S., Kiritchenko, S., Zhu, X.: NRC-Canada: Building the State-of-the-Art in Sentiment Analysis of Tweets. In: Proceedings of the seventh international workshop on Semantic Evaluation Exercises (SemEval-2013), Atlanta, Georgia, USA (2013)
13. Zarrella, G., Marsh, A.: MITRE at SemEval-2016 Task 6: Transfer Learning for Stance Detection. In: Bethard et al. [8]. (2016) 470–475
14. Wei, W., Zhang, X., Liu, X., Chen, W., Wang, T.: pkudblab at SemEval-2016 Task 6 : A Specific Convolutional Neural Network System for Effective Stance detection. In: Bethard et al. [8]. (2016) 396–400
15. Tutek, M., Sekulic, I., Gombar, P., Paljak, I., Culinovic, F., Boltuzic, F., Karan, M., Alagić, D., Šnajder, J.: TakeLab at SemEval-2016 Task 6: Stance Classification in Tweets Using a Genetic Algorithm Based Ensemble. In: Bethard et al. [8]. (2016) 476–480
16. Zhang, Z., Lan, M.: ECNU at SemEval 2016 Task 6: Relevant or Not? Supportive or Not? A Two-step Learning System for Automatic Detecting Stance in Tweets. In: Bethard et al. [8]. (2016) 463–469
17. Elfardy, H., Diab, M.: CU-GWU Perspective at SemEval-2016 Task 6: Ideological Stance Detection in Informal Text. In: Bethard et al. [8]. (2016) 446–451
18. Liu, C., Li, W., Demarest, B., Chen, Y., Couture, S., Dakota, D., Haduong, N., Kaufman, N., Lamont, A., Pancholi, M., Steimel, K., Kübler, S.: IUCL at SemEval-2016 task 6: An Ensemble Model for Stance Detection in Twitter. In: Bethard et al. [8]. (2016) 406–412
19. Vijayaraghavan, P., Sysoev, I., Vosoughi, S., Roy, D.: Deepstance at SemEval-2016 Task 6: Detecting Stance in Tweets Using Character and Word-Level CNNs. In: Bethard et al. [8]. (2016) 425–431
20. Krejzl, P., Steinberger, J.: UWB at SemEval-2016 Task 6: Stance Detection. In: Bethard et al. [8]. (2016) 420–424
21. Bøhler, H., Asla, P., Marsi, E., Sætre, R.: IDI@NTNU at SemEval-2016 Task 6: Detecting Stance in Tweets Using Shallow Features and GloVe Vectors for Word Representation. In: Bethard et al. [8]. (2016) 457–462
22. Igarashi, Y., Komatsu, H., Kobayashi, S., Okazaki, N., Inui, K.: Tohoku at SemEval-2016 Task 6: Feature-based Model versus Convolutional Neural Network for Stance Detection. In: Bethard et al. [8]. (2016) 413–419
23. Wojatzki, M., Zesch, T.: ltl.uni-due at SemEval-2016 Task 6: Stance Detection in Social Media Using Stacked Classifiers. In: Bethard et al. [8]. (2016) 440–445
24. Patra, B.G., Das, D., Bandyopadhyay, S.: JU_NLP at SemEval-2016 Task 6: Detecting Stance in Tweets using Support Vector Machines. In: Bethard et al. [8]. (2016) 452–456

25. Misra, A., Ecker, B., Handleman, T., Hahn, N., Walker, M.: NLDS-UCSC at SemEval-2016 Task 6: A Semi-Supervised Approach to Detecting Stance in Tweets. In: Bethard et al. [8]. (2016) 432–439
26. Dias, M., Becker, K.: INF-UFRGS-OPINION-MINING at SemEval-2016 Task 6: Automatic Generation of a Training Corpus for Unsupervised Identification of Stance in Tweets. In: Bethard et al. [8]. (2016) 390–395
27. Augenstein, I., Vlachos, A., Bontcheva, K.: USFD at SemEval-2016 Task 6: Any-Target Stance Detection on Twitter with Autoencoders. In: Bethard et al. [8]. (2016) 401–405
28. Finn Årup Nielsen: A new ANEW: Evaluation of a word list for sentiment analysis in microblogs. CoRR **abs/1103.2903** (2011)
29. Pennebaker, J.W., Francis, M.E., Booth, R.J.: Linguistic inquiry and word count: LIWC 2001. Lawrence Erlbaum Associates, Mahwah,NJ (2001)
30. Whissell, C.: Using the Revised Dictionary of Affect in Language to Quantify the Emotional Undertones of Samples of Natural Language. Psychological Reports **105** (2009) 509–521
31. Mohammad, S., Kiritchenko, S., Sobhani, P., Zhu, X., Cherry, C.: A dataset for detecting stance in tweets. In: Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016), ELRA (2016)
32. Chan, T.F., Golub, G.H., LeVeque, R.J.: Updating Formulae and a Pairwise Algorithm for Computing Sample Variances. Technical report, Stanford University, Stanford, CA, USA (1979)
33. Hernández Farías, D.I., Patti, V., Rosso, P.: Irony detection in twitter: The role of affective content. ACM Trans. Internet Technol. **16** (2016) 19:1–19:24
34. Sulis, E., Hernández Farías, D.I., Rosso, P., Patti, V., Ruffo, G.: Figurative messages and affect in twitter: Differences between #irony, #sarcasm and #not. Knowledge-Based Systems **108** (2016) 132 – 143 New Avenues in Knowledge Bases for Natural Language Processing.