Exploiting Affective-based Information for Profiling Ironic Users on Twitter

Delia Irazú Hernández-Farías, Manuel Montes-y-Gómez

Laboratorio de Tecnologías del Lenguaje, Instituto Nacional de Astrofísica, Óptica y Electrónica, Mexico

Abstract

This paper describes our approach to addressing the Profiling Irony and Stereotype Spreaders on Twitter (IROSTEREO) 2022 shared task. For classifying the users, we consider the number of ironic tweets they posted. The ironic content in a given tweet is determined according to an irony detection model which mainly relies on affective information. Different traditional classifiers were evaluated being the Random Forest the one with the highest performance. According to the official results, our system obtained a 0.933 accuracy rate. Additionally, a sub-task on stance detection focused in ironic users was also organized. For participating in this sub-task, we evaluated the same set of features than in the main task, obtaining the first place in the official ranking.

Keywords

irony detection, author profiling, stance detection, affective information

1. Introduction

Nowadays, social media platforms have become one of the main communication channels. People share ideas, information, opinions, and judgments about a wide variety of topics from general events to personal experiences. Such content has been used for different research purposes ranging from sentiment analysis to author profiling. Considering the latter, different aspects like gender, age, personality, native language, political ideology, among others [1]. Recently, interest in profiling authors who use social media to achieve particular communication purposes has grown. Some shared tasks have been organized to profile users spreading *fake news* [2] in 2020, *hate speech* [3] in 2021, and this year *irony and stereotypes* [4]. IROSTEREO aims to determine whether or not an author spreads Irony and Stereotypes considering a set of tweets posted by him/her. Then, the authors must be classified as ironic or not depending on the number of ironic tweets he/she has. In addition, a subtask on Stance Detection is also proposed aiming to identify if an ironic author is in favour or against the target of a given tweet.

Most of the time, people only have an intuitive definition of what irony is. Thus, dealing with this kind of figurative language device from a computational linguistics perspective is an ongoing and challenging task. For some natural language processing areas like sentiment analysis and human-computer interaction, irony detection is a very related task that could help avoid misinterpreting ironic statements as literal.

CLEF 2022 – Conference and Labs of the Evaluation Forum, September 5–8, 2022, Bologna, Italy
dirazuhf@inaoep.mx (D. I. Hernández-Farías); mmontes@inaoep.mx (M. Montes-y-Gómez)
https://ccc.inaoep.mx/~mmontesg/ (M. Montes-y-Gómez)
0000-0003-2133-8716 (D. I. Hernández-Farías); 0000-0002-7601-501X (M. Montes-y-Gómez)

Irony is a concept difficult to define which serves to express opinions in an indirect way. It has been considered a linguistic device where the speaker intends to communicate the opposite meaning of what is literally said [5]. *Irony* serves to express an evaluative judgment towards a particular target [6] and/or to reveal the speaker's position (approval or disapproval) on the result of something [7]. In addition, when *Irony* involves a negative evaluation towards a particular target it is considered as *Sarcasm* [8, 9].

Data from different social media platforms such as Amazon reviews, Reddit, and mainly Twitter has been used for research purposes, being Twitter the most widely exploited [10]. Irony detection has been addressed as a text classification task by using different perspectives like textual-based features [11, 12, 13], information regarding the context surrounding the comments [14, 15, 16], and deep learning-based methods like word-embeddings, convolutional neural networks, and transformers [17, 18, 19, 20, 21]. Some works have also considered irony detection as a class imbalance problem [22, 23, 24] due to the inherent data skew on the presence of irony in social media. Research considering irony and profiling information is scarce. In [25], a Spanish dataset collected from Facebook labeled with emotions, irony, and the author's gender is described.

2. Our proposal

Ironic utterances are very related to the expression of feelings, emotions, and evaluations (often in an indirect way) towards a particular target. Research has been done on the role of affect in the presence of irony [6, 26]. We are interested in assessing the role of affective-based information in profiling ironic users. For determining whether or not a given user can be profiled as being ironic, we decided to take into account the number of ironic tweets published by her/him. First, we classify each tweet per user as ironic or not by using *emotIDM* (for more details see Section 2.1). In this model, a wide set of resources covering different facets of affect from sentiment to finer-grained emotions is exploited for identifying the presence of irony. Then, depending on the number of ironic tweets automatically identified, a given user is labeled as "ironic" or not. In the following paragraphs, we describe in detail the proposed approach.

2.1. emotIDM

In order to determine whether or not a given tweet is ironic, we decided to take advantage of an irony detection model which relies mainly on affective information, i.e., emotIDM [27]. This model was evaluated over a set of corpora in the state-of-the-art achieving higher results than in the literature and validating the usefulness of affect-related information for detecting ironic content in tweets. emotIDM comprises a total of 78 features distributed in three different groups for representing a tweet:

(i) *Structural.* Aspects like punctuation marks, length of words and length of chars, part-of-speech labels, Twitter marks (i.e., hashtags, mentions, etc.), semantic similarity between the words composing a given tweet, among others are considered.
(ii) *Sentiment.* Different facets of sentiment are considered from an overall value in terms of how many *positive* and *negative* words a given tweet contains, to a polarity degree in

numerical terms depending on the words composing a tweet. A wide range of English lexical resources were exploited like: AFINN [28], Hu&Liu [29], and SentiWordNet [30], among others.

(iii) *Emotions.* With the intention of considering as much information regarding emotions as possible, the main theories in the nature of emotions are comprised in emotIDM: *a) Categorical model* where the emotions are associated to labels such as "anger", "fear", "joy", "surprise", "disgust", etc. by means of lexical resources like EmoLex [31], EmoSenticNet [32], and LIWC [33]; and *b) Dimensional model* where emotions are associated to its position in a space of independent dimensions like "activation", "pleasantness", "imagery", etc., by using ANEW [34], Dictionary of Affect in Language [35], and SenticNet [36].

2.2. Ironying degree

To classify a given user as *ironic* or *non-ironic*, we decided to consider an *ironying degree*, i.e., the number of ironic tweets posted by him/her. The tweets are labeled according to emotIDM. For doing so, we evaluate two different approaches. The first approach (denoted as *majority*) was to assign a user as ironic considering the number of ironic (henceforth *numIronic*) and non-ironic tweets (henceforth *numNonIronic*) he/she has, then the final decision was made according to:

if *numIronic* > *numNonIronic*:
    the user is labeled as *userIronic*
otherwise:
    the user is labeled as *userNonIronic*

In the second approach, we decided to calculate the *ironying degree* as a threshold (denoted a *iD*) which was determined according to the average of how many *ironic* tweets each user in the training set has. By taking advantage of the *iD* values and the *numIronic* from each user, the following criterion[1] was used to classify the users:

if *numIronic* > *iD*:
    the user is labeled as *userIronic*
otherwise:
    the user is labeled as *userNonIronic*

3. Results

3.1. Profiling ironic users

Participating teams in IROSTEREO were provided with training and test subsets of data. The former is composed by 420 users equally distributed in the two classes: ironic and non-ironic. For each user a total of 200 tweets are available. Concerning to the test partition, the aim is to

[1]It is important to mention that we also evaluate other criteria including also the standard deviation obtaining lower results.

classify a total of 180 users by using the same amount of tweets than in the training data. The official evaluation metric is the *Accuracy*.

As mentioned before, in order to determine whether or not a given user is ironic, we rely on the labels assigned by emotIDM for the tweets posted by her. Then, assessing its performance for doing this task is very important. In this sense, binary classification experiments with a 5-fold cross-validation setting were carried out. The Scikit-learn implementation of traditional classifiers such as Support Vector Machine (SVM), Decision Tree (DT), k-Nearest Neighbors (kNN, with values of 3, 5, and 7 for k), and Random Forest (RF) with default parameters was used. For experimental purposes, the official training data was splitted into two subsets: *train* and *val* in order to determine the *ironying degree*.

First, we decided to assess the performance of emotIDM to categorize the tweets as a standard irony detection setting. We evaluated different configurations with respect to the group of features used, for participating in the shared task we only consider a total of 60 features concerning to the *Sentiment* and *Emotions* described in Section 2.1. In these experiments, all tweets were merged as a single dataset without distinction of belonging to its corresponding author. We assume that all tweets belonging to authors labeled as ironic/non-ironic have the same class at tweet level. Table 1 shows the obtained results in terms of Accuracy for both partitions *train* and *val*. As it can be observed, the results are almost the same among classifiers concerning the data partition. The best classification rate was obtained by the RF in both cases. It is interesting to note that, unlike [27] the highest classification rate was not obtained with DT, however, RF was not considered in the evaluation setting described. Notwithstanding, in [37] the best classification performance was achieved by RF for identifying ironic tweets. Both approaches also exploited affective information for detecting irony on Twitter.

Table 1
Obtained results of classifying tweets as ironic or not with emotIDM

	SVM	DT	RF	3NN	5NN	7NN
train	0.634	0.599	**0.668**	0.59	0.597	0.604
val	0.633	0.597	**0.669**	0.595	0.60	0.606

In our approach, the criterion to classify users as ironic or not is the *ironying degree*. As mentioned before, one way to obtain such a value (denoted as *iD*) is to determine to how many *ironic* tweets each user has in the *train* partition according to the classification obtained from emotIDM. Table 2 shows the obtained values. As it can be observed, the *iD* value is very similar across the classifiers and it represents practically half of the available tweets per user.

Table 2
Threshold values obtained with each classifier

SVM	DT	RF	3NN	5NN	7NN
98.148	99.99	103.651	101.51	101.88	102.121

Finally, considering only those users in the *val* subset, we labeled all the tweets for each user and calculate her/his *ironying degree* obtained by either criteria: *majority* and *iD* to determine whether or not each user is ironic. Table 3 shows the obtained results of classifying the users in the *val* subset. An improvement in terms of Accuracy was observed in all classifiers when the *iD* is used as a decision criterion with respect to use the difference between ironic and non-ironic tweets.

Table 3
Obtained results of labeling users as ironic or non-ironic according to the number of ironic tweets determined by emotIDM on the *val* subset during the development phase.

Criterion	SVM	DT	RF	3NN	5NN	7NN
majority	0.716	0.874	0.853	0.819	0.822	0.807
iD	0.88	0.89	0.9	0.9	0.89	0.895

For participating in the shared task we chose to exploit emotIDM together with a Random Forest to determine an *iD* threshold. Then, for each user in the official *test* set provided we labeled each tweet with emotIDM and calculate his/her *ironying degree*. According to the official results, we ranked in the 32^{th} position up to 64^{th} with an accuracy rate of 0.933. It is important to mention that, the difference in comparison with the best ranked approach is of 0.0611. Besides, the proposed method achieves higher results than three out of the four baselines considered by the organizers.

3.1.1. Analysis of the results

We have the intuition that most of the mislabeled users are due to the errors provoked by the emotIDM model. Concerning the official training data, we decided to analyze the obtained results in the *val* partition. First, we identified that for some users labeled as *nonironic* (being *ironic* in the golden label) the errors could be due to very small differences in terms of the number of ironic tweets determined by emotIDM. The same phenomenon occurs in the opposite way. With respect to the correctly classified users, in the case of the *ironic* ones, we observed that there are some cases where the number or ironic tweets identified represents more than the 60% of the available samples. Probably, a higher threshold could help to improve the performance in identifying ironic users. In the case of the *nonironic*, some instances having less than 10% of ironic tweets were identified. We hypothesise that, enriching the subset of emotIDM used for participating in the shared task with features capturing stylistic cues could help to capture ironic profiles.

3.2. Pilot experiments on Stance Detection

In order to participate in the subtask dedicated to identify if an ironic author is *INFAVOR* or *AGAINST* the target of a given tweet, we decided to assess the performance of emotIDM for this challenging task. Affective-based information for dealing with this task has been already evaluated in the state-of-the-art [38]. The criterion to determine the stance of an author is similar to the *majority* one used in the profiling task, the difference is that we classify each tweet as *INFAVOR* or *AGAINST* instead of in terms of irony, and then we count how many tweets are for each class, the majority one is then assigned. Organizers provided with 140 users (for each of them a total of 200 tweets are available) distributed as 94 labeled as *AGAINST* and 46 as *INFAVOR*.

Attempting to compensate class imbalanced towards the *AGAINST* class, we decided to apply the random over-sampling (ROS) implementation in Scikit-learn with default parameters. The same set of classifiers mentioned before was used. Given the fact that, emotIDM have not been evaluated before for stance detection, we experimented with all the features at the same time (*allFeatures*), and by separating them by groups as mentioned in Section 2.1. Table 4 shows the obtained results over the training data considering two data settings *Original* and *ROS*, and with two groups of features showing the best performance. These experiments were performed in a five fold-cross validation setting, in each fold, ROS was applied only for the training partition while the test partition was left untouched. The official metric in this sub-task is the Macro F-score. As it can be observed, the *kNN* classifier reaches the highest results. Besides, the use of ROS has a positive impact in the performance with both configurations.

Table 4
Obtained results in Macro F-score of labeling ironic users in terms of their stance. Bold numbers represent the best performing configurations.

Data	SVM	DT	RF	3NN	5NN	7NN
			allFeatures			
Original	0.39	0.39	0.39	0.39	0.39	0.39
ROS	0.38	0.38	0.38	0.48	**0.5**	**0.54**
			Structural			
Original	0.38	0.38	0.38	0.38	0.38	0.38
ROS	0.44	0.44	0.38	**0.51**	**0.55**	0.46

For participating in the shared task, we submitted the four configurations showing the best performance during the development phase. Table 5 shows the official results obtained in the shared task. Considering the official ranking with the configuration composed by "*Structural + 3NN*" we ranked in the first position among the 14 submissions in the shared task. One more time, our best-ranking proposal achieves higher results than two of the three the baselines established by the organizers. Interestingly, the *Structural* subset of features was not used in the task concerning ironic users. The second best result in the shared task was achieved by using all the features in emotIDM, which could serve as a starting point to further investigate the usefulness of affective-based information for stance detection.

Table 5
Official results in the stance detection sub-task.

Configuration	ACC	F1-Macro	F1-InFavour	F1-Against
allFeatures + 7NN	0.65	0.5433	0.3226	0.764
Structural + 5NN	0.6333	0.4876	0.2143	0.7609
Structural + 3NN	0.7833	**0.6248**	0.381	0.8687
allFeatures + 5NN	0.7	0.5807	0.3571	0.8043

4. Conclusions

In this paper, we describe our participation in IROSTEREO 2022. We propose an approach to classify users as ironic or not depending on the amount of potentially ironic tweets he/she posted. To identify the presence of irony, we took advantage of a wide range of lexical resources comprising different facets of affective information. Our system reached a 0.933 accuracy rate according to the official results. We consider that our approach obtained competitive results despite its simplicity, which could validate the usefulness of considering the role of affect for detecting irony in social media. Our model reaches a higher performance than most of the baselines, one of them using deep learning approach. With respect to the stance detection subtask, we ranked at the first position according to the official results. In both subtasks, the baseline outperforming our proposal is the LDSE (Low-Dimensionality Statistical Embedding) which considers the probability of distribution of the occurrence of terms for text representation. Our approach is not directly using term-based information for text representation. As future work, it could be interesting to enhance the proposed approach with a deep learning-based classification schema. Moreover, further analysis of the correctly classified instances could be interesting, particularly in those where the stereotypes were spreading. Furthermore, an analysis of the performance of emotIDM for dealing with stance detection is also an interesting research direction.

References

[1] M. Álvarez Carmona, E. Villatoro Tello, M. Montes y Gómez, L. Villaseñor Pineda, Author Profiling in Social Media with Multimodal Information, Computación y Sistemas 24 (2020) 1289 – 1304.

[2] F. Rangel, A. Giachanou, B. Ghanem, P. Rosso, Overview of the 8th Author Profiling Task at PAN 2020: Profiling Fake News Spreaders on Twitter, in: L. Cappellato, C. Eickhoff, N. Ferro, A. Névéol (Eds.), CLEF 2020 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2020.

[3] J. Bevendorff, B. Chulvi, G. L. De La Peña Sarracén, M. Kestemont, E. Manjavacas, I. Markov, M. Mayerl, M. Potthast, F. Rangel, P. Rosso, E. Stamatatos, B. Stein, M. Wiegmann, M. Wolska, E. Zangerle, Overview of PAN 2021: Authorship Verification, Profiling Hate Speech Spreaders on Twitter, and Style Change Detection, in: Experimental IR Meets Multilinguality, Multimodality, and Interaction: 12th International Conference of the CLEF Association, CLEF 2021, Springer-Verlag, Berlin, Heidelberg, 2021, p. 419–431.

[4] O.-B. Reynier, C. Berta, R. Francisco, R. Paolo, F. Elisabetta, Profiling Irony and Stereotype Spreaders on Twitter (IROSTEREO) at PAN 2022, in: CLEF 2022 Labs and Workshops, Notebook Papers, CEUR-WS.org, 2022.

[5] H. P. Grice, Logic and Conversation, in: P. Cole, J. L. Morgan (Eds.), Syntax and Semantics: Vol. 3: Speech Acts, Academic Press, 1975, pp. 41–58.

[6] L. Alba-Juez, S. Attardo, The Evaluative Palette of Verbal Irony, in: G. Thompson, L. Alba-Juez (Eds.), Evaluation in Context, John Benjamins Publishing Company, Amsterdam/Philadelphia, 2014, pp. 93–116.

[7] A. Utsumi, Verbal Irony as Implicit Display of Ironic Environment: Distinguishing Ironic Utterances from Nonirony, Journal of Pragmatics 32 (2000) 1777 – 1806.

[8] S. Attardo, Irony as Relevant Inappropriateness, in: H. Colston, R. Gibbs (Eds.), Irony in language and thought: A cognitive science reader, Lawrence Erlbaum, 2007, pp. 135–172.

[9] S. McDonald, Neuropsychological Studies of Sarcasm, in: H. Colston, R. Gibbs (Eds.), Irony in language and thought: A cognitive science reader, Lawrence Erlbaum, 2007, pp. 217–230.

[10] A. Joshi, P. Bhattacharyya, M. J. Carman, Automatic Sarcasm Detection: A Survey, CoRR abs/1602.03426 (2016).

[11] F. Barbieri, H. Saggion, F. Ronzano, Modelling Sarcasm in Twitter, A Novel Approach, in: Proc. of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis, ACL, USA, 2014, pp. 50–58.

[12] A. Reyes, P. Rosso, T. Veale, A Multidimensional Approach for Detecting Irony in Twitter, Language Resources and Evaluation 47 (2013) 239–268.

[13] E. Riloff, A. Qadir, P. Surve, L. D. Silva, N. Gilbert, R. Huang, Sarcasm as Contrast between a Positive Sentiment and Negative Situation, in: Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, ACL, USA, 2013, pp. 704–714.

[14] D. Bamman, N. A. Smith, Contextualized Sarcasm Detection on Twitter, in: Proceedings of the Ninth International Conference on Web and Social Media, 2015, pp. 574–577.

[15] A. Khattri, A. Joshi, P. Bhattacharyya, M. Carman, Your Sentiment Precedes You: Using an author's Historical Tweets to Predict Sarcasm, in: Proceedings of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis, ACL, Portugal, 2015, pp. 25–30.

[16] B. C. Wallace, D. K. Choe, E. Charniak, Sparse, Contextually Informed Models for Irony Detection: Exploiting User Communities, Entities and Sentiment, in: Proc. of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing, ACL, China, 2015, pp. 1035–1044.

[17] A. Joshi, V. Tripathi, K. Patel, P. Bhattacharyya, M. J. Carman, Are Word Embedding-based Features Useful for Sarcasm Detection?, in: Proc. of the 2016 Conference on Empirical Methods in Natural Language Processing, Austin, Texas, USA, 2016, pp. 1006–1011.

[18] D. Nozza, E. Fersini, E. Messina, Unsupervised Irony Detection: A Probabilistic Model with Word Embeddings, in: Proc. of the 8th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management, 2016, pp. 68–76.

[19] S. Poria, E. Cambria, D. Hazarika, P. Vij, A Deeper Look into Sarcastic Tweets Using Deep Convolutional Neural Networks, in: Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics, Japan, 2016, p. 1601–1612.

[20] R. A. Potamias, G. Siolas, A. G. Stafylopatis, A transformer-based approach to irony and sarcasm detection, Neural Computing and Applications 32 (2020) 17309–17320.

[21] A. Agrawal, A. K. Jha, A. Jaiswal, V. Kumar, Irony detection using transformers, in: 2020 International Conference on Computing and Data Science (CDS), 2020, pp. 165–168.

[22] P. Liu, W. Chen, G. Ou, T. Wang, D. Yang, K. Lei, Sarcasm Detection in Social Media Based on Imbalanced Classification, in: F. Li, G. Li, S.-w. Hwang, B. Yao, Z. Zhang (Eds.), Proceedings of the Web-Age Information Management: 15th International Conference, Springer International Publishing, China, 2014, pp. 459–471.

[23] G. Abercrombie, D. Hovy, Putting Sarcasm Detection into Context: The Effects of Class Imbalance and Manual Labelling on Supervised Machine Classification of Twitter Conversations, in: Proceedings of the ACL 2016 Student Research Workshop, ACL, Germany, 2016, pp. 107–113.

[24] D. I. Hernández Farías, R. Prati, F. Herrera, P. Rosso, Irony detection in twitter with imbalanced class distributions, Journal of Intelligent & Fuzzy Systems 39 (2020) 2147–2163.

[25] F. Rangel, D. I. H. Farías, P. Rosso, A. Reyes, Emotions and Irony per Gender in Facebook, in: Proceedings of the Workshop on Emotion, Social Signals, Sentiment & Linked Open Data (ES3LOD), LREC-2014, 2014.

[26] D. Wilson, D. Sperber, On Verbal Irony, Lingua 87 (1992) 53–76.

[27] D. I. Hernández Farías, V. Patti, P. Rosso, Irony detection in Twitter: The Role of Affective Content, ACM Trans. Internet Technol. 16 (2016) 19:1–19:24.

[28] F. Å. Nielsen, A New ANEW: Evaluation of a Word List for Sentiment Analysis in Microblogs, in: Proceedings of the ESWC2011 Workshop on 'Making Sense of Microposts': Big things come in small packages, volume 718 of *CEUR Workshop Proceedings*, CEUR-WS.org, Greece, 2011, pp. 93–98.

[29] M. Hu, B. Liu, Mining and Summarizing Customer Reviews, in: Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '04, ACM, USA, 2004, pp. 168–177.

[30] S. Baccianella, A. Esuli, F. Sebastiani, SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining, in: Proceedings of the Seventh Conference on International Language Resources and Evaluation, Malta, 2010, pp. 2200–2204.

[31] S. M. Mohammad, P. D. Turney, Crowdsourcing a Word–emotion Association Lexicon, Computational Intelligence 29 (2013) 436–465.

[32] S. Poria, A. Gelbukh, A. Hussain, N. Howard, D. Das, S. Bandyopadhyay, Enhanced SenticNet with Affective Labels for Concept-Based Opinion Mining, IEEE Intelligent Systems 28 (2013) 31–38.

[33] J. W. Pennebaker, M. E. Francis, R. J. Booth, Linguistic Inquiry and Word Count: LIWC 2001, Mahway: Lawrence Erlbaum Associates 71 (2001) 2–23.

[34] M. M. Bradley, P. J. Lang, Affective Norms for English Words (ANEW): Instruction Manual and Affective Ratings, Technical Report, Center for Research in Psychophysiology, University of Florida, Florida, 1999.

[35] C. Whissell, Using the Revised Dictionary of Affect in Language to Quantify the Emotional Undertones of Samples of Natural Languages, Psychological Reports 2 (2009) 509–521.

[36] E. Cambria, D. Olsher, D. Rajagopal, SenticNet 3: A Common and Common-Sense Knowledge Base for Cognition-Driven Sentiment Analysis, in: Proceedings of AAAI Conference on Artificial Intelligence, AAAI, Canada, 2014, pp. 1515–1521.

[37] E. Sulis, D. I. Hernández Farías, P. Rosso, V. Patti, G. Ruffo, Figurative Messages and Affect in Twitter: Differences between #irony, #sarcasm and #not, Knowledge-Based Systems 108 (2016) 132 – 143.

[38] M. Lai, D. I. Hernández Farías, V. Patti, P. Rosso, Friends and Enemies of Clinton and Trump: Using Context for Detecting Stance in Political Tweets, in: G. Sidorov, O. Herrera-Alcántara (Eds.), Advances in Computational Intelligence, Springer International Publishing, Cham, 2017, pp. 155–168.