On the Impact of Emotions on the Detection of False Information

Paolo Rosso
Universitat Politècnica de València, Spain
prosso@dsic.upv.es

Bilal Ghanem
University of Alberta, Canada
bilalhgm@gmail.com

Anastasia Giachanou
Utrecht University, The Netherlands
a.giachanou@uu.nl

ABSTRACT

A great amount of fake news are propagated in online social media, with the aim, usually, to deceive users and formulate specific opinions. The threat is even greater when the purpose is political or ideological and they are used during electoral campaigns. Bots play a key role in disseminating these false claims. False information is intentionally written to trigger emotions to the readers in an attempt to be believed and be disseminated in social media. Therefore, in order to discriminate credible from non credible information, we believe that it is important to take into account these emotional signals. In this paper we describe the way that emotional features have been integrated in deep learning models in order to detect if and when emotions are evoked in fake news.

CCS CONCEPTS

• **Computing methodologies**; • **Neural networks; Natural language processing;**;

KEYWORDS

Fake News, False Information, Credibility of Claims, Emotions

ACM Reference Format:
Paolo Rosso, Bilal Ghanem, and Anastasia Giachanou. 2021. On the Impact of Emotions on the Detection of False Information. In *2021 International Symposium on Electrical, Electronics and Information Engineering (ISEEIE 2021), February 19–21, 2021, Seoul, Republic of Korea.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3459104.3459150

1 INTRODUCTION

Emotions play an important role in our life and taking them into account may help also when processing texts. For instance, the authors of [20] studied the impact of emotions on author profiling, concretely identifying age and gender. The proposed emotion-based graph model obtained state-of-the-art results. Despite the large variety of models proposed in the literature, the impact that emotions could have on the detection of false claims and false information in general, has not been explored much [27]. In this paper our **objective** is to investigate whether emotions may be effective or not in detecting fake news. Concretely, we aim at answering the following **research question**: Do emotional features help detecting false information and profiling fake news spreaders?

Figure 1: A fake news tweet that may potentially spread emotions such as disgust, anger, and surprise.

In figure 1, we show an example of a fake news tweet that went viral during the US 2016 elections. The tweet talks about a well-known false claim whose target was Hillary Clinton. As can be noticed, the author of this tweet tries to affect her reputation negatively by making her a criminal because behind a child sex ring.

The rest of the paper is structured as follows. In Section 2 we describe the related work. Section 3 summarises how we addressed the problem of verifying the credibility of claims in news taking into account emotions signals. Section 4 describes our emotionally-infused model that is used to detect false information both in online news articles as well as in Twitter. In Section 5 we propose a model that takes into account how affective information changes in fake news. In Section 6 we present some preliminary results addressing fake news detection from an author profile perspective with one of the models presented. Finally, in Section 7 we draw some conclusions and discuss future work.

2 RELATED WORK

In this section we cite some of the previous works that addressed the problem of the detection of false information. In [29] false information was categorized into eight different types: fabricated information, hoaxes, rumors, clickbaits, propaganda, conspiracy theories, biased or one-sided information, and satire. In [1] clickbait detection was addressed employing several features, among them the existence of hyperbolic words: i.e., words with a high positive or negative sentiment (e.g. terrifying, awe-inspiring, etc). The approach achieved an F1 of 0.93. In [10] the authors proposed a set of content-based features, including readability features, stylistic features, etc. These features were fed to a Support Vector Machine (SVM) obtaining 91% of accuracy on differentiating satire from real news. A lower accuracy of 78% was instead obtained for discriminatind between fake and real news. In [16], the authors proposed FakeNews-Detector, a model whose representation was based on word unigrams and bigrams, psycholinguistic, readability, punctuation (e.g. existence of special chars), and syntax features (features based on rules from the dependency tree). These features were fed

ISEEIE 2021, February 19–21, 2021, Seoul, Republic of Korea

ACM ISBN 978-1-4503-8983-9/21/02…$15.00
https://doi.org/10.1145/3459104.3459150

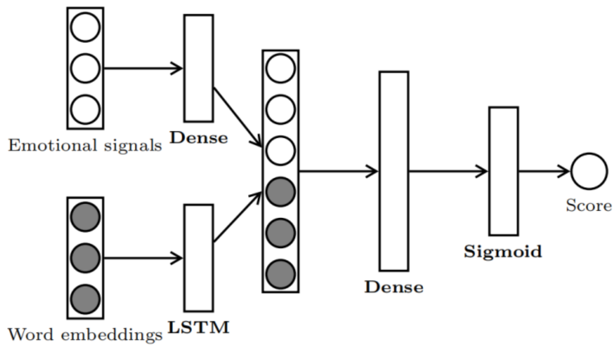

Figure 2: EmoCred neural network architecture for credibility assessment.

to a SVM obtaining weighted F1 values of 0.74 and 0.73 on two different datasets.

3 CREDIBILITY OF CLAIMS IN NEWS

In [6] we proposed the EmoCred model which incorporated emotional signals into a Long Short Term Memory (LSTM) deep learning architecture in order to differentiate between credible and non-credible claims. The hypothesis behind is that credible and non-credible claims trigger different emotions to the readers.

3.1 Dataset

For the dataset we use data from PolitiFact[1] that is a fact checking website where the credibility of different claims is investigated. We use two different Politifact datasets presented in two different studies. The first dataset[2] (Politifact- 1) was presented by Popat et al. [17] and contains 3,568 claims of which 1,867 are credible and 1,701 are not credible claims. The second dataset (Politifact-2) was presented by Rashkin et al. [23] and consists of 2,575 training, 712 development and 1,074 test statements. There are six different credibility ratings: true, mostly true, half true, mostly false, false and pants-on-fire. As it was previously done in [17, 23], we combined true, mostly true and half true labels into the true class label, and the rest as false, addressing the problem as a binary classification.

3.2 Model

EmoCred is a model based on an LSTM architecture whose input are word embeddings from the text of the claims together with a vector of emotional signals. Three are the approaches that we considered for generating the emotional signals from the claims: (i) *emoLexi*, a lexicon-based approach that considers the number of emotional words that appear in the claim, (ii) *emoInt*, an approach that uses an emotional intensity lexicon to calculate the emotional intensity expressed in the claim, and (iii) *emoReact*, a deep learning-based approach that predicts the level of emotional intensity that can be triggered to the users. Figure 2 gives an overview of the architecture of our model.

[1] https://www.politifact.com/
[2] https://www.mpi-inf.mpg.de/dl-cred-analysis/

Table 1: Performance results of EmoCred approach when using different approaches for generating the emotional signals. A star (*) indicates statistically significant improvement over the LSTM-text approach.

Dataset	Method	Accuracy	F1-score
	LSTM-text	0.551	0.549
Politifact-1	EmoCred-emoLexi	0.608	0.602*
	EmoCred-emoInt	0.604	0.602*
	EmoCred-emoReact	0.617	0.617*
	LSTM-text	0.597	0.567
Politifact-2	EmoCred-emoLexi	0.621	0.606*
	EmoCred-emoInt	0.628	0.586
	EmoCred-emoReact	0.619	0.601*

3.3 Experiments and Results

In order to calculate the emotional signals in the claims with the *emoLexi* approach, we use the following emotional lexicons: EmoLex [14], SentiSense [2] and EmoSenticNet [18], whereas for the *emoInt* approach, we use the NRC Affect Intensity Lexicon [12]. *emoReact* detects the emotional reactions signals using a dataset crawled by Facebook as explained in [7]. The dataset was aimed for determining the emotional triggers of news posts and contains 26,560 news posts that span from April 2016 to September 2017 crawled from New York Time Facebook page together with the actual number of emotional reactions that they triggered. For each of the claims, we predict the probability to trigger any of the five different intensities (very low, low, average, high, very high) of five different emotional reactions (love, joy, sadness, surprise, anger). We use the pre-trained GloVe Wikipedia 6B word embeddings [15] to initialise the word embeddings.

Table 1 summarises the results of our experiments. From the results, we observe that EmoCred outperforms the LSTM baseline by a large margin, that is, incorporating emotional signals into LSTM allows for significantly improving the results. On Politifact-1, the best performance was achieved considering the emotional reactions: emoReact significantly outperformed LSTM-text by 12.39% in terms of F1-score. This is quite interesting because *emoReact* was trained on different data (crawled from Facebook). Although the model was trained on data from a different domain, it seems that still the emotional features are very helpful for the credibility assessment task. In case of *emoInt* and *emoLexi*, we observe that the two approaches obtain a similar performance. Both emoLexi and emoInt manage to significantly outperform LSTM-text by 9.65%. On Politifact-2, a simlar behaviour could be observed: EmoCred obtained the best results with *emoLexi*, that significantly outperformed the LSTM baseline by 6.5%. Our results show that the emotional signals are effective for the credibility assessment task.

4 FALSE INFORMATION IN NEWS AND TWITTER

In [5] we study the different types of false information in news articles and Twitter. Following, we describe the two datasets and the emotionally-infused model that we proposed.

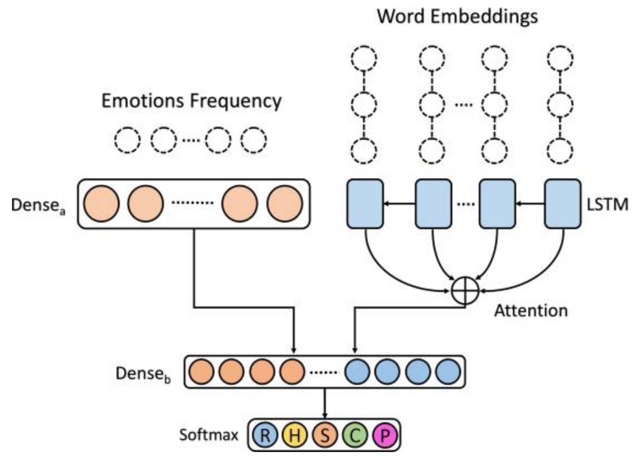

Figure 3: Emotionally-infused neural network architecture for false information detection. RHSCP in the softmax layer stands for real, hoax, satire, clickbait, and propaganda respectively.

4.1 Dataset

News Articles. Our dataset of news articles was used in [23]. This dataset was built from two different sources, for the trusted news (real news) the authors sampled news articles from the English Gigaword corpus; for the false news, they collected articles from seven different unreliable news sites. These articles include hoaxes, propagandas, and satires but not clickbaits. In order to be able to analyse also clickbaits, we included part of the Stop- Clickbait dataset [1], that was originally collected from the following two sources: Wikinews articles' headlines and other online sites that are known to publish clickbaits. The satire, hoax, and propaganda news articles are considerably long (some of them reach the length of 5,000 words). In total we have 31,550 news articles.

Twitter. For this dataset, we rely on a list of several Twitter accounts for each type of false information [26], where tweets were collected from suspicious Twitter accounts that were previously annotated. For the real news, we considered also other 32 Twitter accounts from [11]. The final Twitter dataset is composed of 152,026 tweets.

4.2 Model

We choose an LSTM [9] in order to take the sequence of words as input and predict the false information type. The input of our network is based on word embedding (content-based) and emotional features. See Figure 3 for the architecture of the Emotionally-Infused Neural (EIN) network.

Regarding the emotional features, we consider several emotion resources to increase the coverage of the emotion words in texts as well to have a wider range of emotions in the analysis. Concretely, we use EmoSenticNet [18], EmoLex [13], SentiSense [2], LIWC [25] and Empath [4].

Table 2: Results of the proposed model (EIN) vs. the baselines.

Method	Precision$_{macro}$	Recall$_{macro}$	F1$_{macro}$
News Articles			
BOW+SVM	0.72	0.71	0.71
W2V+LR	0.70	0.70	0.70
LSTM	0.77	0.74	0.74
EIN	0.79	0.80	0.79
Twitter			
BOW+SVM	0.60	0.56	0.57
W2V+LR	0.49	0.35	0.36
LSTM	0.65	0.54	0.56
EIN	0.61	0.59	0.60

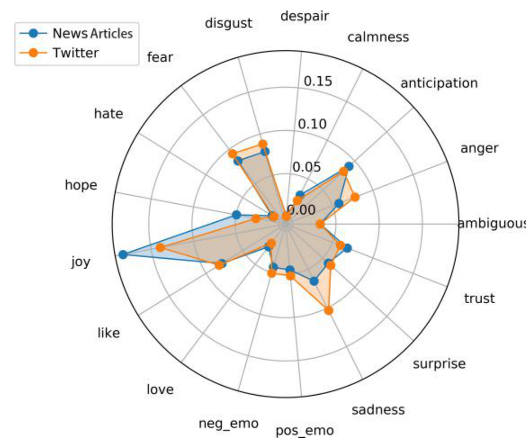

Figure 4: Best ranked features according to Information Gain.

4.3 Experiments and Results

To validate the performance of our proposed model, we compared it to a set of baselines: (i) **BOW-SVM.** It is based on a bag-of-words representation with a SVM classifier. We test different classifiers, and we choose SVM since it gives the highest result in the 10-fold-Cross validation;

(ii) **W2V-LR.** It is based on word embeddings where for each input document we extract an average word embedding vector by taking the mean of the embeddings for the document's words. Similarly, we test different classifiers and the Logistic Regression classifier shows the best performance;

(iii) **LSTM.** The last baseline is the same as our neural architecture but without emotional features: an LSTM layer followed by attention and dense layers.

Table 2 summarises the performance of the proposed model and compares it to those obtained by the several baselines. We report macro-precision, recall, and F1. LSTM obtained the best performance on the news articles dataset, whereas in Twitter there is a different scenario: the BOW-SVM base- line shows a higher performance with respect to LSTM. EIN results outperform the baselines with a large margin, especially in the news articles dataset (around

3% in Twitter and 5% in news articles). On Twitter the difference between EIN and the best baseline is lower. Comparing the results obtained by both EIN and LSTM, it is possible to appreciate that emotional features allows for an improvement.

In Figure 4, we show the importance of the emotions in each dataset. The figure shows that the emotion "joy" is the important emotion to discriminate true from false information in both datasets, followed by "sadness" in Twitter data.

5 FAKE NEWS AT FRAGMENT LEVEL

We hypothesise that fake news articles may evoke more exaggerated emotions at the beginning of their text than in the rest. Based on that, we propose a fake news detection system that takes into account the way emotional information flows in news articles. To evaluate the performance of our model, we compare our model to several state-of-the-art models.

5.1 Dataset

We built our dataset in two parts, training and test. For the training part, we built a list of fake news Websites using a set of news Websites' lists that were annotated from a factuality perspective (fake or real content): 560 domains from *OpenSources.co* (OS), 548 domains from *MediaBiasFactCheck.com* (MBFC), and 227 domains from PolitiFact[3] lists. We use Websites' domains that were annotated in a consistent way across the three lists; e.g. we discard domains that are annotated as *fake* in OS list but real in MBFC list. The final list contains 85 domains. Our approach is to project the domain-level label onto the content of those domains. Thus, we sample randomly a maximum of 100 articles per domain. On the other hand, for the test part, we use *leadstories.com*, a fact checking Website in which expert journalists annotated online news at article-level. We scrape all the available articles in the Website, and in total we obtain around 700 fake articles. For the real class, we sample around 1000 articles from the training part. Finally, we postprocess all the articles by discarding very short articles (less than 30 words).

5.2 Model

Giving a news article, we split it into N segments based on equal number of words. We extract a vector of length 23 features out of each segment. These features are *emotions* [13], *sentiment* [13], *morality* [8], *imageability*[4], and hyperbolic words [1]. We feed the article's vectors to a Bidi- rectional Gated Recurrent Unit (Bi-GRU) neural network to learn the flow of these features in the article. The topic of a news article is important to be associated with the extracted emotional information: e.g. a fake news article about Islam or Black people likely triggers fear and a negative sentiment. On the other hand, a fake news that is in favor of a politician will trigger more positive emotions and some overstatements. We use a Convolutional Neural Network (CNN) to extract important topic words from the text fragments. After applying CNN on the text fragments, we concatenate the output of each text fragment with its emotional feature vector to learn their joint interaction. After this, we feed them to a dense layer and then we apply a dot product operation between the output of this step with the output of the Bi-GRU layer. Finally, we apply dense and softmax layers.

[3] https://www.politifact.com/article/2017/apr/20/politifactsguide-fake-news-websites-and-what-they/

[4] https://github.com/ytsvetko/metaphor/imageability

Table 3: Results on the collected dataset. A star (*) indicates a statistically significant improvement of *FakeFlow* over the referred model using the Mc- Nemar test.

Method	Precision	Recall	F1*macro*
Majority Class	0.35	0.59	0.37
Horne & Adali, 2017 [10]	0.75	0.78	0.80
BERT CLS	0.84	0.82	0.82
FakeNewsDetector	0.86	0.86	0.86
BERT LSTM	0.89	0.89	0.89
LSTM	0.86	0.91	0.90
CNN	0.89	0.89	0.91
Rashkin et el., 2017 [23]	0.92	0.92	0.92
EIN	0.94	0.93	0.93*
HAN	0.94	0.94	0.93*
FakeFlow	0.93	0.97	**0.96**

Figure 5: The flow of Fear emotion in fake (►) and real (●) news articles in a percentage.

5.3 Experiments and Results

Baselines. To evaluate our model, we propose different baselines: (i) LSTM and CNN; (ii) two BERT-based [3] baselines, one using the CLS BERT token and the other using the output of the BERT layer of each token in another LSTM; (iii) state-of-the-art models: [10], FakeNewsDetector, [23], and our model that was proposed in Section 4; (iv) HAN [28] which is a Hierarchical Attention Networks model that has two levels of attention mechanisms: word and sentence-level attentions. The model splits a document into sentences (splits on dots), and starts learning the sentences' representation from words.

Results. Table 3 presents the results of our model comparing them with those obtained by the baselines. For the N parameter of our model, we tested several numbers of fragments on the validation set and it turns out that 10 gives the highest macro F1 value. The results demonstrate the effectiveness of our proposed model comparing to the rest of the baselines. Anyway, it has to be said that the dataset has been compiled from multiple sources, sampling positive class instances from different data sources than the negative class instances. As future work, it will be important to investigate if the resulting classifiers recognise the class label or merely the source distribution.

Finally, in order to illustrate the emotions distribution across the documents segments, in Figure 5 we show how emotions such as fear flow in real and fake news articles. The figure shows that fake news start with a higher average value of fear that finally decreases.

Table 4: Statistics of the PAN-AP-20 dataset for the shared task on profiling fake news spreaders on Twitter.

Language	Training	Test	Total
English	300	200	500
Spanish	300	200	500

6 PROFILING FAKE NEWS SPREADERS ON TWITTER

In 2020 we addressed the problem of fake news detection from the author profiling perspective, with the aim of profiling those users that have shared some fake news in the past. In order to prevent fake news from being propagated among online users, it is important to identify possible fake news spreaders on Twitter. This should help for their early detection and, therefore, for preventing their further dissemination. A shared task on *Profiling fake news spreaders on Twitter* was organised at the PAN lab[5].

6.1 Dataset

In order to discriminate authors that have shared some fake news in the past from those that, to the best of our knowledge, have never done it, we built a dataset of fake and real news spreaders. The dataset consists of 500 authors for each of the two languages (English and Spanish). The dataset for each language is balanced, with 250 fake news spreaders and 250 real news spreaders. For each author we retrieved her last 100 tweets via the Twitter API. Table 4 presents the statistics of the dataset.

6.2 Models

We compared the Emotionally-Infused Neural (EIN) net- work[6] described in Sec. 4 with: (i) an LSTM that uses fast-Text[7] embeddings to represent texts; (ii) a Support Vector Machine (SVM) with char n-grams (size 2-6); (iii) an SVM with Low Dimensionality Statistical Embeddings (LDSE)[22] to represent texts; (iv) a Neural Network (NN) with word-grams (size 1-3); and (v) a Random prediction. We represent each author in the dataset by concatenating her tweets into one document and then we feed this document to the above models.

6.3 Experiments and Results

In Table 5 we present the results. Whereas for the task of false information detection the EIN model obtained better results than those of the baselines (see Section 4), for profiling fake news spreaders its performance, although better than the one of a simple LSTM without emotional features, was not better than those obtained by classical classifiers based on n-grams. Therefore, apart from the tweets containing fake news (and that have been removed from the dataset), the rest of tweets of fake news spreaders do not seem to use emotions differently than the other authors. The lower performance of EIN and LSTM is also likely due to the small size of the training data that does not allow deep neural models to generalise, and in line with in previous author profiling shared tasks [21]. The results of Table 5 will be also compared with the state-of-the-art results of the teams that participated in the shared task [19].

[5]https://pan.webis.de/clef20/pan20-web/author-profiling.

[6]For Spanish, we use the Spanish Emotion Lexicon [24] to extract emotions from tweets.

[7]https://fasttext.cc/docs/en/crawl-vectors.html

7 CONCLUSIONS AND FUTURE WORK

In this paper we showed how effective emotion-based deep learning models may be for discriminating credible from non credible claims, and for detecting false information, in news (also at fragment level) and in Twitter. Finally, we applied one of the proposed models, the emotionally-infused neural network, to address the problem of profiling fake news spreaders in Twitter. As future work, we aim at investigating a more effective way to consider emotional information also for detecting fake news from the author profile perspective.

ACKNOWLEDGMENTS

The work of Paolo Rosso was partially funded by the Spanish MICINN under the research project MISMIS-FAKEnHATE on Misinformation and Miscommunication in social media: FAKE news and HATE speech (PGC2018-096212-B-C31), and by the Generalitat Valenciana under the project PROM- ETEO/2019/121 (DeepPattern).

Anastasia Giachanou was supported by the SNSF Early Postdoc Mobility grant P2TIP2 181441 under the project *Early Fake News Detection on Social Media, Switzerland.*

REFERENCES

[1] Chakraborty, B. Paranjape, S. Kakarla, and N. Ganguly. Stop clickbait: Detecting and preventing clickbaits in online news media. In *Proceedings of the 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 9–16. IEEE Press, 2016.

[2] J. C. de Albornoz, L. Plaza, and P. Gerv´as. Sentisense: An easily scalable concept-based affective lexicon for sentiment analysis. In *LREC*, pages 3562–3567, 2012.

[3] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019.

[4] E. Fast, B. Chen, and M. S. Bernstein. Empath: Understanding topic signals in large-scale text. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 4647–4657. ACM, 2016.

[5] Ghanem, P. Rosso, and F. Rangel. An emotional analysis of false information in social media and news articles. *ACM Transactions on Internet Technology (TOIT)*, 20(2):1–18, 2020.

[6] A. Giachanou, P. Rosso, and F. Crestani. Leveraging emotional signals for credibility detection. In Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'19, page 877–880, New York, NY, USA, 2019. Association for Computing Machinery.

[7] Giachanou, P. Rosso, I. Mele, and F. Crestani. Emotional influence prediction of news posts. In *Proceedings of the 12th International AAAI Conference on Web and Social Media*, ICWSM '18, pages 592–595, 2018.

[8] J. Graham, J. Haidt, and B. A. Nosek. Liberals and Conservatives Rely on Different Sets of Moral Foundations. *Journal of personality and social psychology*, 96(5):1029, 2009.

[9] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[10] D. Horne and S. Adali. This Just In: Fake News Packs a Lot in Title, Uses Simpler, Repetitive Content in Text Body, More Similar to Satire than Real News. *In Eleventh International AAAI Conference on Web and Social Media*, 2017.

[11] A.Karduni, R. Wesslen, S. Santhanam, I. Cho, S. Volkova, D. Arendt, S. Shaikh, and W. Dou. Can you verifi this? studying uncertainty and decision-making about misinformation using visual analytics. 2018.

[12] S. M. Mohammad. Word affect intensities. In *Proceedings of the 11th International Conference on Language Resources and Evaluation*, pages 174–183, 2018.

[13] S. M. Mohammad and P. D. Turney. Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon. In *Proceedings of the*

Table 5: Results on the PAN-AP-20 dataset.

Lang.	Method	$F1_{fake\_news}$	Accuracy
English	Random	0.51	0.51
	NN + word ngrams	0.73	0.69
	SVM + char ngrams	0.70	0.68
	LDSE	**0.74**	**0.75**
	LSTM	0.55	0.56
	EIN	0.67	0.64
Spanish	Random	0.50	0.50
	NN + word ngrams	0.76	0.70
	SVM + char ngrams	**0.78**	**0.79**
	LDSE	0.77	**0.79**
	LSTM	0.60	0.60
	EIN	0.64	0.64

NAACL HLT 2010 workshop on computational approaches to analysis and generation of emotion in text, pages 26–34. Association for Computational Linguistics, 2010.

[14] S. M. Mohammad and P. D. Turney. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465, 2013.

[15] J. Pennington, R. Socher, and C. Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language*, EMNLP '14, pages 1532–1543, 2014.

[16] V. P´erez-Rosas, B. Kleinberg, A. Lefevre, and R. Mihalcea. Automatic Detection of Fake News. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3391–3401, Santa Fe, New Mexico, USA, Aug. 2018. Association for Computational Linguistics.

[17] K. Popat, S. Mukherjee, A. Yates, and G. Weikum. Declare: Debunking fake news and false claims using evidence-aware deep learning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language*, EMNLP '18, pages 22–32, 2018.

[18] S. Poria, A. Gelbukh, A. Hussain, N. Howard, D. Das, and S. Bandyopadhyay. Enhanced senticnet with affective labels for concept-based opinion mining. *IEEE Intelligent Systems*, 28(2):31–38, 2013.

[19] F. Rangel, A. Giachanou, B. Ghanem, and P. Rosso. Overview of the 8th Author Profiling Task at PAN 2020: Profiling Fake News Spreaders on Twitter (In Press). In L. Cappellato, C. Eickhoff, N. Ferro, and A. N´ev´eol, editors, *CLEF 2020 Labs and Workshops, Notebook Papers*. CEUR-WS.org, Sept. 2020.

[20] F. Rangel and P. Rosso. On the impact of emotions on author profiling. *Information processing & management*, 52(1):73–92, 2016.

[21] F. Rangel and P. Rosso. Overview of the 7th Author Profiling Task at PAN 2019: Bots and Gender Profiling. In *CLEF 2019 labs and workshops, notebook papers. CEUR Workshop Proceedings. CEUR-WS.org.*, volume 2380, 2019.

[22] F. Rangel, P. Rosso, and M. Franco-Salvador. A Low Dimensionality Representation for Language Variety Identification. In *In 17th International Conference on Intelligent Text Processing and Computational Linguistics, CICLing'16. Springer-Verlag, LNCS(9624)*, pages 156–169, 2018.

[23] H. Rashkin, E. Choi, J. Y. Jang, S. Volkova, and Y. Choi. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2931–2937, 2017.

[24] G. Sidorov, S. Miranda-Jim´enez, F. Viveros-Jim´enez, A. Gelbukh, N. Castro-S´anchez, F. Vel´asquez, I. Dıaz-Rangel, S. Su´arez-Guerra, A. Trevino, and J. Gordon. Empirical study of opinion mining in spanish tweets. micai 2012. *Lect Notes Comput Sci*, 7629:1–14, 2012.

[25] Y. R. Tausczik and J. W. Pennebaker. The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of language and social psychology*, 29(1):24–54, 2010.

[26] S. Volkova, K. Shaffer, J. Y. Jang, and N. Hodas. Separating facts from fiction: Linguistic models to classify suspicious and trusted news posts on twitter. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 647–653, 2017.

[27] S. Vosoughi, D. Roy, and S. Aral. The Spread of True and False News Online. *Science*, 359(6380):1146–1151, 2018.

[28] Z. Yang, D. Yang, C. Dyer, X. He, A. Smola, and E. Hovy. Hierarchical Attention Networks for Document Classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*, pages 1480–1489, 2016.

[29] S. Zannettou, M. Sirivianos, J. Blackburn, and N. Kourtellis. The web of false information: Rumors, fake news, hoaxes, clickbait, and various other shenanigans. *arXiv preprint arXiv:1804.03461*, 2018.