



HAL
open science

An Approach Utilizing Linguistic Features for Fake News Detection

Dimitrios Panagiotis Kasseropoulos, Christos Tjortjis

► **To cite this version:**

Dimitrios Panagiotis Kasseropoulos, Christos Tjortjis. An Approach Utilizing Linguistic Features for Fake News Detection. 17th IFIP International Conference on Artificial Intelligence Applications and Innovations (AIAI), Jun 2021, Hersonissos, Crete, Greece. pp.646-658, 10.1007/978-3-030-79150-6_51 . hal-03287679

HAL Id: hal-03287679

<https://inria.hal.science/hal-03287679v1>

Submitted on 15 Jul 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

An Approach Utilizing Linguistic Features for Fake News Detection

Dimitrios Panagiotis Kasseropoulos¹ and Christos Tjortjis^{1[0000-0001-8263-9024]}

¹The Data Mining and Analytics Research Group, School of Science & Technology,
International Hellenic University, Thessaloniki, Greece
c.tjortjis@ihu.edu.gr

Abstract

Easy propagation and access to information on the web has the potential to become a serious issue when it comes to disinformation. The term “fake news” describes the intentional propagation of news with the intention to mislead and harm the public and has gained more attention recently. This paper proposes a style-based Machine Learning (ML) approach, which relies on the textual information from news, such as manually extracted lexical features e.g. part of speech counts, and evaluates the performance of several ML algorithms. We identified a subset of the best performing linguistic features, using information-based metrics, which tend to agree with the literature. We also, combined Named Entity Recognition (NER) functionality with the Frequent Pattern (FP) Growth association rule algorithm to gain a deeper perspective of the named entities used in the two classes. Both methods reinforce the claim that fake and real news have limited differences in content, setting limitations to style-based methods. Results showed that convolutional neural networks resulted in the best accuracy, outperforming the rest of the algorithms.

Keywords: Fake news, Social media, Machine Learning (ML), Natural Language Processing (NLP), Association Rule (AR) Mining, Data Mining

1 Introduction

Social media are an important part of our everyday lives, changing the way we interact with other people [8, 15]. One of the aspects that could be affected is the way we receive and publish information. Easy access to high-speed internet, tools that made website deployment easier, and the growing popularity of many microblog websites made publishing and receiving news information accessible to everyone, anytime. Although the large number of informative online sources increased the variety of aspects available, many of them have low quality, making filtering a necessity [21]. These conditions have also created a trend/danger called fake news.

Social media gave the opportunity to news to have an alternative way of reaching the public rapidly, but at the same time, they also benefit disinformation propagation. Studies have shown that fake, extremely one-side (hyperpartisan) and emotional news tend to spread far more rapidly than traditional news [13, 24].

Some factors that benefit disinformation on the web are the difficulty of accessing trustworthy information and the lack of trust in the traditional informative means. A fake story can have serious impacts on society if a significant volume of people believes it. Finally, during the COVID-19 era, fake news is on the rise, making the work of

health professionals more difficult in an already critical situation, endangering this way the public [10].

The rapid propagation of information on the web makes quick detection of fake news crucial. That is why the new technologies of Machine Learning (ML) and artificial intelligence have been utilized widely in the last years to tackle this problem and it is also the topic focus of this paper [28].

The remainder of this paper is structured as follows: Section 2 reviews the literature and provides background information. Section 3 presents the dataset used and the experimental approach including feature extraction and selection. Section 4 presents results which are further discussed in section 5. Section 5 presents conclusions and the directions for future work.

2 Background

2.1 Fake news characteristics

In their research, Petty, R. E. & Cacioppo J. T., presented the Elaboration Likelihood model of Persuasion (ELP) theorem, arguing that people are persuaded either by a central route, meaning that all the arguments are examined, or by the peripheral route, which focuses only on the validity of the key concepts of a claim [12]. The peripheral route is frequent in social media, since studies have shown that most of the articles shared are never read [23].

Based on the ELP theorem and Cacioppo's and Petty's findings, Khan, J. Y., et al., argue that fake news targets the peripheral route and therefore their titles contain the most important claims about people and events [6]. The titles' role in fake news mostly serves as the main mechanism of information propagation, where the body just repeats the title's claims [5].

Regarding content characteristics, fake articles are a lot smaller in length using fewer technical words, smaller words, fewer punctuation, fewer quotes, and more lexical redundancy. Also, at the linguistic level, they use simpler language resulting in fewer analytic words, more personal pronouns, fewer nouns, and more adverbs [5]. Finally, a good indicator is the emotional response the article tries to achieve. Strong emotional words and phrases draw more attention and propagate faster [17].

2.2 Fake news categories

Most research split fake news into categories, based on the two basic characteristics: intention and quality of information. On a first level, the author's motive to mislead or not separates fake news to misinformation and disinformation [19]. In the next section we further describe some common fake news categories.

Rumors. A rumor can be defined as "a piece of circulating information whose veracity status is yet to be verified at the time of spreading" [29]. Today rumors flourish in social media and their detection becomes more difficult [9]. Studies have focused on supervised, unsupervised and hybrid methods to separate rumors from real news [1].

Other studies confirmed that the propagation style differs significantly from real news' and is used to classify rumors on the web [9].

Conspiracy theories. This genre of fake news provides explanations for stories of the news referring to entities that exist in the center of attention, but most of the time, these explanations are based in pseudo-scientific results [18]. Conspiracy theories create a way of thinking opposite to the scientific method of explanation, making groups of people with predisposition to them more open to sharing and stand up for misinformation [13].

Click bait. In order for an article to be considered clickbait it needs to have some basic characteristics, including: i) short text, ii) a media attachment, such as image or video and iii) the link to the publisher's article [7]. Most of the publishers in social media use click bait articles, to a greater or lesser extent, to attract more readers. However, journalistic codes of ethics are opposed to these techniques, as they use unethical means to misdirect the readers [13].

Satire. According to B.D Horne, and S. Adali, fake news has more similarities in content with satire than with real news. A basic common characteristic of the two genres is that they use similar persuasion methods based on heuristics and not arguments [5]. Although, most satirical news' primary goal is to entertain rather to mislead the reader the term "satire" has also been used by many webpages that do not have any intent to entertain, but to create fake content without being accused of deception [3, 16].

2.3 Fake news detection methods

According to Potthast, M., et al., the detection methods of fake news can be divided in three categories which are: 1) knowledge-based 2) style-based and 3) content-based [13].

Knowledge – based. Knowledge-based detection method is about identifying the basic claims and statements of the article and comparing them with known facts. This procedure could become either manually or automatic. In manual evaluation a person or a group of people are responsible to judge the validity of the article's main statements [28]. Automatic denotation classifies an article in two stages: Fact extraction and fact checking [13]. For fact extraction the algorithm constructs a knowledge base by mining raw "facts" from the web, and during the fact checking stage, it extracts the basic statements of the article and compares them with the knowledge base facts [28].

Style-based. The most common approach for fake news detection, relying on the research findings from studying the linguistic characteristics of deception [26]. Even though deceptive writers try to mimic the writing style of journalists, there are still some characteristics that could reveal the authenticity of an article, also known as Undeutsch hypothesis [22]. Those characteristics can be split in the following categories:

Lexical features: Describe character and word level signals, such as total words, characters per word, number of unique words etc. [14, 20].

Language features: Syntax in sentence-level describing number of words, syllables per sentence, number of characters per sentence, word types, and number of paragraphs

[14]. Also, they calculate several readability metrics that approximate the appropriate knowledge level that a reader should have to understand the text [11].

Syntactic features: Include frequencies of function words, phrases, and punctuations, and Parts-Of-Speech (POS) tagging [20].

Domain-specific linguistic features: Specifically aligned to news domains, such as quoted words, external links, number of graphs, and average length of graphs [20].

Psycholinguistic features: Category of features based on the linguistics that have to do with the psychological aspect of words. This approach tries to identify the psychological reaction that the article tries to achieve [6].

Content-based. Regarding fake news published in social media, information from the social network can be used, such as user-based information (number of followers of the publisher), post-based information (number of likes, shares etc.) and network-based information (propagation of the news) to effectively tackle the problem [20].

3 Approach

Our experimental approach consists of three parts, shown in Fig. 1: 1) Text Preprocessing, which includes cleaning techniques, 2) Feature Extraction, which includes the extraction and testing of different combinations of feature sets, and 3) Model Testing which tests various ML algorithms. In the following sections we detail these steps.




Figure 1: Flowchart for the experimental approach

3.1 Selected data

The dataset we used was the one provided for the “2nd Int’l TrueFact Workshop: Making a Credible Web for Tomorrow in conjunction with SIGKDD 2020”, created by Kai Shu and contained news related to famous people of the timelines (<https://www.kaggle.com/c/fakenewskdd2020/data>). The dataset consists of 2972 real and 2014 fake news.

3.2 Data engineering/ Feature selection

Lexical features. We extracted 84 lexical features using either manually defined functions or pre-existing ones to be used by the classification task. Regarding text preprocessing at the sentence level, we removed stop words, numbers, and links which are entities that do not contribute significantly to the information of the text. At the word level, we tested three different text representations: i) removing stop words and links only, ii) applying stemming and iii) applying lemmatization. The raw text was used to count Part Of Speech (POS) frequencies and for Named Entities features.

To end up with an optimal feature set we used Decision Tree’s feature importance method based on i) mutual information, ii) gini impurity and iii) information entropy, and kept the set for which a baseline model scored the best results. As baseline model we chose a linear Support Vector Machine (SVM), as it was also used in previous studies [4, 5].

More specifically the steps followed were:

- i. Use the full feature set.
- ii. Sort features based on the selected information-based metric score.
- iii. Measure the performance of the linear Support Vector Classifier.
- iv. Drop the bottom two features.
- v. Repeat 2-4 until one or no features are left.

The results are shown in Table 1. The best lexical feature consists of 23 variables, shown in Table 2, corresponding to 67.02% accuracy.

Table 1: Results of the different information-based criteria

Information metric	Number of best features	Accuracy score
Mutual Information	10	65.7%
Gini Impurity	23	67.02%
Entropy	12	66.4%

Table 2: The best performing features according to gini impurity

%mverbs_freq	%punctuations	%gerund_participle
%first_person_singular	%clauses	noun_diversity
%present_verbs	semantic_redundancy	verb_diversity
%possesive_prn	%modifiers	Subjectivity
%adverbs	%third_person_singular	big_words_ratio
Redundancy	%art	function_words_diversity
%stopwords	%third_person_plural	avg_len_noun_phrase
weighted_sentiment	avg_punct_per_sentence	

Word embeddings. Three different word embedding representations were tested: i) Google’s pre-trained vectors, trained on part of Google News dataset (about 100 billion words), a model that contains 300-dimensional word vectors for 3 million words

and phrases (available at: <https://code.google.com/archive/p/word2vec/>) ii) spacy's word embeddings that includes 1 million different 300-dimensional word vectors designed by using the GloVe algorithm, and iii) our own vectors trained using the word2vec algorithm in the given dataset/corpus.

By testing all three, using 10-fold cross validation to a linear SVM classifier, we ended up with spacy's representations that had the best accuracy results. After that, we chose the best preprocessing method for text which returns the best accuracy results. The base algorithm used for testing was a linear SVM and the Clean text representation provide the best results.

The next step was to use Principal Component Analysis (PCA) to reduce the dimensions of the 300-dimensional embeddings to improve performance. The best results were achieved for 180 dimensions, using the clean text representation. Again, the utilization of linear SVM as benchmark model for comparison was chosen to keep consistency within the feature selection steps.




Figure 2: Accuracy vs. feature set dimensions

Extracting Association Rules (AR) from Named Entities. Another experiment that provided better understanding of the lexical structure at the sentence level was the utilization of Frequent Pattern (FP) Growth algorithm on the set of Named Entities used in the articles. In a first step, spacy library was used to extract the named entities from each article's sentence and then FP Growth was applied to extract rules [2]. Results are presented in section 4.3.

4 Results and evaluation

4.1 Performance results for the ML Algorithms

The first experimental results have been produced by using all the lexical features extracted from the texts. As SVMs use distances to classify their samples, scaling provided a significant performance boost. Tree based algorithms, on the other hand, are not affected by scaling dissimilarities, so we skipped the scaling step for their testing. Results are shown in Table 3.

Table 3: Accuracy scores for different feature sets and models

	SVM	Decision Tree	Random Forest	Gradient Boosting
All Lexical Features	59.8%	61.5%	71.9%	70.8%
All Lexical Features - Scaled	72.1%	-	-	-
Best Lexical Features	58.9%	62.0%	72.3%	70.4%
Best Lexical Features -Scaled	71.7%	-	-	
Word Embeddings	74.7%	63.8%	75.3%	75.3%
Word Embeddings - Reduced dimensions	77.2%	65.5%	73.5%	74.1%
Best lexical features with word embeddings	76.2%	65.8%	75.8%	76.2%
Best lexical features with reduced word embeddings	75.8%	65.2%	75.5%	75.9%

4.2 Performance results for Artificial Neural Networks

The Artificial Neural Networks (ANN) we tested were Convolutional Neural Network (CNN) and Long-Short Term Memory (LSTM), deployed using the Keras library. For the initial embedding layer, we used the pre-trained word embeddings of spacy’s library.

Convolutional Neural Network (CNN). In order to find the best combinations of activation functions we used a simple shallow CNN constructed by four layers. First, there was the embedding layer, which consists of 300 neurons, then the Convolutional layer with 128 neurons, a max pooling layer with 128 neurons and finally, a dense layer with the activation function.

For choosing the best combination of activation functions (Relu, Sigmoid, Softmax, Softsign, Exponential, Tanh) for CNN and Dense layers, and optimizers (Adam, Adadelata, Adamx, SGD) we performed a nested “for” loop testing all the combinations.

Fifteen epochs were used for each of the combinations and the best performance was given with relu as the activation function of the CNN layer, sigmoid activation function for the dense layer, and Adam optimizer, with performance 75.6%. After testing several architectures, the best performance was achieved by the one shown in Table 4, with accuracy performance **79.2%**

Table 4: Best performing architecture for CNN

Layer	Embedding Layer	Convolutional Layer	Convolutional Layer	Max pooling	Dense Layer	Dense Layer
# of neurons	300	300	128	128	28	1

Long-Short Term Memory (LSTM). Similarly, with CNN, for LSTM we tested different combinations of activation functions (Relu, Sigmoid, Softmax, Softsign, Exponential, Tanh) and optimizers (Adam, Adadelata, Adamx, SGD) in a simple network structure using a nested “for” loop.

The architecture of this baseline model consisted of an embedding layer with 300 neurons, a LSTM layer with 300 neurons, a Dropout layer with 300 neurons, a Flatten layer with 300 neurons and finally a Dense layer with one neuron. The best combination was the sigmoid activation function and the adamax optimizer. After testing several different architectures for LSTM, the maximum performance was achieved with the one shown in Table 5, with accuracy performance **75.2%**.

Table 1: Best performing architecture for LSTM.

Layer	Embedding Layer	LSTM Layer	Dropout	Flatten	Dense Layer	Dense Layer
# of neurons	300	300	300	300	30	1

4.3 AR extracted from Named Entities of the articles

We used FP-Growth to extract AR from the Named Entities of the articles. The results, show in Table 6, indicate that, fake and real news do not significantly differ with regards to the entities they use, confirming that fake news tend to use similar terms with real ones, limiting the scope of linguistic approaches. The only different entity in real news was the “work of art” item referring to titles of books, songs, etc.

Table 62: 4.3 Items in rules extracted from Named Entities

Real news AR	Fake news AR
'PERSON'	'PERSON'
'DATE'	'DATE'
'ORG'	'ORG'
'CARDINAL'	'GPE'
'GPE'	'CARDINAL'
' WORK_OF_ART '	'DATE', 'PERSON'
'DATE', 'PERSON'	'ORG', 'PERSON'
'ORG', 'PERSON'	'DATE', 'ORG'
'DATE', 'ORG'	'DATE', 'ORG', 'PERSON'
'DATE', 'ORG', 'PERSON'	'GPE', 'PERSON'

5 Discussion

We extracted 84 style-based features. By using the Gini impurity metric, we ended up with an optimal subset of 23 linguistic features, which improved accuracy, providing at the same time explainable results. We further discuss some of the best performing features, information captured from the text and whether results were expected.

Third person singular / Third person plural / First Person Singular: The use of first-person singular is more frequent in real articles. This is because writers who try to deceive readers tend to separate themselves from the information they propagate [26].

Our findings, shown in Table 7, agree with the expected results, showing higher usage of the first person in real news and higher usage of the third person, plural and singular, in fake news.

Table 73: Feature statistics for two classes

Statistic	3 rd person singular		3 rd person plural		1 st Person Singular	
	Real news	Fake news	Real news	Fake news	Real news	Fake news
Mean	1.96%	2.32%	0.32%	0.47%	0.88%	0.67%
Std. deviation	1.43%	1.52%	0.52%	0.61%	1.86%	1.47%

Percentage of modal verbs: Modal verbs indicate uncertainty (would, could, might etc.) and are most often used by deceivers. This happens because deceivers are not sure about the information they propagate, so they tend to hypothesize and imply correlations about events that do not have clear connection between them [26]. Again, our findings, shown in Table 8, were expected, showing higher usage of modal verbs in fake news.

Table 8: Statistics of “% modal verbs” for two classes

Statistic	Real news	Fake news
Mean	3.93 %	4.54 %
Standard deviation	1.85 %	1.85 %

By observing the probability distributions of the two most important lexical features (% modal verbs frequency – Fig. 3, and % first person singular – Fig. 4) we observe that there is a high amount of overlap making them difficult to be separated.

Although, style-based features provide explainable results, their performance was overpassed by the vector representations of texts. More specifically, SVM achieved 72.1% accuracy using only the lexical features, but 77.2% accuracy using the embeddings with 180 dimensions, and 76.2% by combining the full 300-dimensional embeddings with the 23 best lexical features (Table 3).

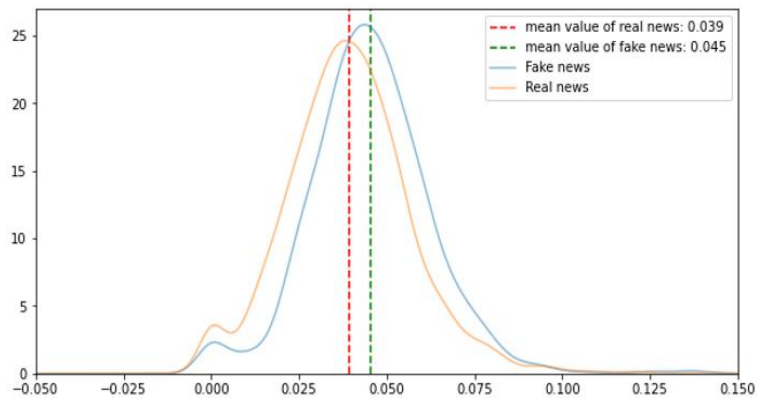


Figure 3: Probability distribution of “% modal verbs” for two classes.




Figure 4: Probability distribution of "% first person singular" for two classes.

The combination of spacy’s Named Entity Recognition functionality with FP-Growth for the extraction of AR between name entities, in real and fake news, produced results which, once again, support the claim that fake and real news have quite similar structure, limiting linguistic approaches. The only entity that differentiates the two classes is the “work of art” item, referring to titles of books, songs, movies etc., for which we assume it is a characteristic of the specific dataset and not a general characteristic.

Finally, the best performance was achieved by CNN, agreeing with the state-of-the-art trend to utilize ANNs in text classification tasks [25].

6 Conclusions and future work

The scope of this paper was to review the state of the art on fake news detection methods and find an efficient way to perform text classification by utilizing the linguistic information of the content. We used a variety of ML algorithms with the ANNs outperforming. Since ANNs brought promising results, for future work, we could focus our research on text classification methods using ANNs. Also, the ANNs that we utilized are not considered deep, leaving space for additional searching and experimentations with deep ANNs.

On the part of style-based features, most of the psychological studies in which those features are based are conducted in real time communication, where fake articles belong to a different category. Although they bring quite satisfying explainable results, there are a lot of differences in the process of writing a fake text and telling lies in real time and first person.

Also, additional statistical analysis could be applied, like t-statistic test to normalized data, to check if the means of two sample/classes distributions for the most important features are significantly different from each other. Finally, it would be very interesting to test the style-based method in some of the most common fake news datasets and compare our results and findings.

References

1. Alzanin, S. M., & Azmi, A. M. (2018). Detecting rumors in social media: A survey. *Procedia computer science*, 142, 294-300.
2. Ghafari, S.M., Tjortjis, C. (2019). A Survey on Association Rules Mining Using Heuristics. *WIRES Data Mining and Knowledge Discovery*, 9(4).
3. Golbeck, J., Mauriello, M., Auxier, B., Bhanushali, K. H., Bonk, C., Bouzaghrane, M. A., ... & Visnansky, G. (2018). Fake news vs satire: A dataset and analysis. In *Proc, 10th ACM Conf. Web Science* (pp. 17-21).
4. Gravanis, G., Vakali, A., Diamantaras, K., & Karadais, P. (2019). Behind the cues: A benchmarking study for fake news detection. *Expert Systems with Applications*, 128, 201-213.
5. Horne, B. D., Adali, S., & Sikdar, S. (2017). Identifying the social signals that drive online discussions: A case study of reddit communities. In *26th IEEE Int'l Conf. Computer Communication and Networks (ICCCN)* (pp. 1-9).
6. Khan, J. Y., Khondaker, M., Islam, T., Iqbal, A., & Afroz, S. (2019). A benchmark study on machine learning methods for fake news detection. *arXiv preprint arXiv:1905.04749*.
7. Kiesel, J., Mestre, M., Shukla, R., Vincent, E., Adineh, P., Corney, D., Stein, B. and Potthast, M., (2019). Semeval-2019 task 4: Hyperpartisan news detection. In *Proc. 13th Int'l Workshop on Semantic Evaluation* (pp. 829-839).
8. Koukaras P., Tjortjis C., Rousidis D. (2020). Social Media Types: Introducing a Data Driven Taxonomy. *Computing*, 102(1), 295-340.
9. Liu, Y., & Xu, S. (2016). Detecting rumors through modeling information propagation networks in a social media environment. *IEEE Transactions on computational social systems*, 3(2), 46-62.
10. Orso, D., Federici, N., Copetti, R., Vetrugno, L., & Bove, T. (2020). Infodemic and the spread of fake news in the COVID-19-era. *European Journal of Emergency Medicine*.
11. Pérez-Rosas, V., Kleinberg, B., Lefevre, A., & Mihalcea, R. (2017). Automatic detection of fake news. *arXiv preprint arXiv:1708.07104*.
12. Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. In *Communication and persuasion* (pp. 1-24). Springer, New York, NY.
13. Potthast, M., Kiesel, J., Reinartz, K., Bevendorff, J., & Stein, B. (2017). A stylometric inquiry into hyperpartisan and fake news. *arXiv preprint arXiv:1702.05638*.
14. Reis, J. C., Correia, A., Murai, F., Veloso, A., & Benevenuto, F. (2019). Supervised learning for fake news detection. *IEEE Intelligent Systems*, 34(2), 76-81.
15. Rousidis D., Koukaras P., C. Tjortjis (2020). Social Media Prediction A Literature Review. *Multimedia Tools and Applications*, 79(9-10), 6279-6311.
16. Rubin, V. L., Conroy, N., Chen, Y., & Cornwell, S. (2016). Fake news or truth? using satirical cues to detect potentially misleading news. In *Proc, 2nd workshop computational approaches to deception detection* (pp. 7-17).
17. Ruchansky, N., Seo, S., & Liu, Y. (2017). Csi: A hybrid deep model for fake news detection. In *Proc, 2017 ACM Conf. Information and Knowledge Management* (pp. 797-806).
18. Shahsavari, S., Holur, P., Tangherlini, T. R., & Roychowdhury, V. (2020). Conspiracy in the time of corona: Automatic detection of covid-19 conspiracy theories in social media and the news. *arXiv preprint arXiv:2004.13783*.
19. Sharma, K., Qian, F., Jiang, H., Ruchansky, N., Zhang, M., & Liu, Y. (2019). Combating fake news: A survey on identification and mitigation techniques. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(3), 1-42.

20. Shu, K., Mahudeswaran, D., Wang, S., Lee, D., & Liu, H. (2018). Fakenewsnet: A data repository with news content, social context and dynamic information for studying fake news on social media arXiv:1809.01286, 8.
21. Tsiara E., C. Tjortjis (2020). Using Twitter to Predict Chart Position for Songs. In Proc. 16th IFIP Int'l Conf. Artificial Intelligence Applications & Innovations, pp. 62-72, 2020.
22. Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *science*, 185(4157), 1124-1131.
23. Wang, L. X., Ramachandran, A., & Chaintreau, A. (2016). Measuring click and share dynamics on social media: a reproducible and validated approach. In 10th Int'l AAAI Conf. Web and Social Media.
24. Wu, L., Li, J., Hu, X., & Liu, H. (2017). Gleaning wisdom from the past: Early detection of emerging rumors in social media. In Proc, 2017 SIAM Int'l Conf. data mining (pp. 99-107).
25. Zhang, Y., & Wallace, B. (2015). A sensitivity analysis of (and practitioners' guide to) convolutional neural networks for sentence classification. arXiv preprint arXiv:1510.03820.
26. Zhou, L., Twitchell, D. P., Qin, T., Burgoon, J. K., & Nunamaker, J. F. (2003). An exploratory study into deception detection in text-based computer-mediated communication. In Proc. 36th IEEE Int'l Conf. System Sciences, 2003 (pp. 10-pp).
27. Zhou, X., Jain, A., Phoha, V. V., & Zafarani, R. (2020). Fake news early detection: A theory-driven model. *Digital Threats: Research and Practice*, 1(2), 1-25.
28. Zhou, X., Zafarani, R., Shu, K., & Liu, H. (2019). Fake news: Fundamental theories, detection strategies and challenges. In Proc, 12th ACM Int'l Conf. web search and data mining (pp. 836-837).
29. Zubiaga, A., Aker, A., Bontcheva, K., Liakata, M., & Procter, R. (2018). Detection and resolution of rumours in social media: A survey. *ACM Computing Surveys*, 51(2), 1-36.