



# Evaluating the Role of News Content and Social Media Interactions for Fake News Detection

Catherine Sotirakou<sup>(✉)</sup>, Anastasia Karampela,  
and Constantinos Mourlas

Faculty of Communication and Media Studies,  
National and Kapodistrian University of Athens, Athens, Greece  
{katerinasot, anakaram, mourlas}@media.uoa.gr

**Abstract.** Societies across the globe suffer from the effects of disinformation campaigns creating an urgent need for a way of tracking falsehoods before they become widely spread. Although building a detection tool for online disinformation campaigns is a challenging task, this paper attempts to approach this problem by examining content-based features related to language use, emotions, and engagement features through explainable machine learning. We propose a model that, except for the textual attributes, harnesses the predictive power of the users' interactions on the Facebook platform, and forecasts deceptive content in (i) news articles and in (ii) Facebook news-related posts. The findings of the study show that the proposed model is able to predict misleading news stories with a 98% accuracy based on features such as capitals in the main body, headline length, Facebook likes, the total amount of nouns and numbers, lexical diversity, and arousal. In conclusion, the paper provides new insights concerning the false news identifiers crucial for both news publishers and consumers.

**Keywords:** Fake news detection · Disinformation · Fact-checking · Digital journalism · Natural language processing · Machine learning · Explainable AI

## 1 Introduction

The intentional spread of false and concocted information serves many purposes such as financial and political interests, influencing public discourse against marginalized populations, has a negative impact on society and democracy [16, 30], and can expose the public to immediate danger. Examples of false stories that went viral on social media platforms like the “Pizzagate”, a conspiracy theory that threatened the lives of the employees of a pizzeria [29] and coronavirus-related false content that led people to drink toxic chemicals with at least 800 people dead and thousands hospitalized<sup>1</sup>, show that online virality can become

<sup>1</sup> <https://www.bbc.com/news/world-53755067>.

dangerous. More specifically, previous research has found that social bots are crucial in the spread of misinformation [27] since search engines, social media platforms, and news aggregators use algorithms that control the information a user sees. For instance, algorithmic curation on Google can promote a greatly visited news article very high on the search results, thus improving the likelihood of it being shared, read, and emailed. Audience metrics such as page views, likes, shares, and so on, unquestionably influence the number of people who see a given article on their screen. Therefore, experts in disinformation and online radicalization take advantage of these known algorithmic vulnerabilities by creating fabricated accounts which generate fake traffic that results in virality [27]. Virality in turn guarantees that disinformation, trolling rumors, and coordinated campaigns are rapidly propagated across the internet, and as Lotan [12] highlights what we need is “algorithms that optimize for an informed public, rather than page views and traffic”. Nevertheless, after much debate about the need for Facebook to change its algorithm to reduce filter bubbles, and the platform’s avoidance of taking responsibility for the distribution of deceptive content on its News Feed, since mid-December 2016 it started to alter its algorithm to make misleading information to appear lower and Google followed with raising the fact-checked stories higher [3]. However, the Covid-19 pandemic proved those measures were insufficient, while also highlighting the challenges that journalists face as they need to manually check countless requests of potentially deceptive information daily<sup>2</sup>, without sometimes possessing the necessary skills, or having the resources, time, and expert personnel to fight disinformation [3].

The urgent need for disinformation detection led many scientific disciplines in the search for new effective ways to mitigate this problem with promising approaches coming from various fields. In line with this, this paper proposes a computational approach to detect potentially fake information, by identifying textual and nontextual characteristics of both fake and real news articles and then using machine learning algorithms for disinformation prediction. More precisely, we consider two sets of machine-readable features i) content-based, and ii) engagement-based, and we conduct our analysis in two distinct phases. In phase A, only content-based features are explored, while in phase B we add features that correspond to the users’ interactions on Facebook and test them on a subset of the original fake and real news dataset.

## 2 Related Work

Fake and manipulated information is circulated in all forms and platforms, unverified videos are shared on Facebook, rumors are being forwarded via messaging apps, while conspiracy theories are being shared by Twitter influencers, and these are only a few of the distribution patterns of disinformation. According to Tandoc and his colleagues [32] the role of social media platforms is crucial to understand the current state of disinformation globally since Facebook and Twitter changed both the news distribution and the trust to traditional media

---

<sup>2</sup> <https://www.poynter.org/coronavirusfactsalliance/>.

outlets. As they vividly note “now, a tweet, which at most is 140 characters long, is considered a piece of news, particularly if it comes from a person in authority” [32]. In this work, we consider real news as defined by Kovach and Rosenstiel [11] to be “independent, reliable, accurate, and comprehensive information”, and “not include unverified facts”, thus disinformation campaigns threaten to curtail the actual purpose of journalism, which is “to provide citizens with the information they need to be free and self-governing” [11]. In addition, to define fake news we use the description by the European Commission [5] “disinformation is understood as verifiably false or misleading information that is created, presented, and disseminated for economic gain or to intentionally deceive the public, and may cause public harm”. Journalists and professional fact-checkers can determine the correctness of potential threats based on their expertise and the use of many digital tools designed to detect a plethora of manipulated elements inside a fake story. Finally, news verification can be a procedure done inside a news outlet that checks all the information before publication or it can be done after the piece is published or shared in social media networks.

The rise of disinformation has attracted strong interest from computer scientists who employ machine learning and other automated methods to help identify disinformation. Fake news detection in computer science is defined as the task of classifying news by its veracity [19] with many studies of this phenomenon aiming to extract useful linguistic and other types of features and then build effective models that can identify and predict fake news from real content. A useful overview of the computational methods used for automated disinformation detection [6] separates two categories, notably machine learning research using linguistic cues, and network analysis using behavioral data. In this section, we will focus only on previous work around the former category, linguistic approaches.

The thought behind linguistic approaches for fake news detection based on content is to find predictive deception elements which can lead to distinguishing the fakeness of news [25]. Rubin et al. [25] built a Support Vector Machine (SVM) model to identify satire and humor articles. Their model performed with 87% accuracy and the results showed that the best predictive features were absurdity, grammar, and punctuation. A similar study from Horne and Adali [9] compared real news against satire articles using also SVM with an accuracy of 91%, and found that headlines, complexity, and style of content are good predictors of satire news. However, when classifying real and fake news the accuracy dropped dramatically. Ahmed et al. [1] experimented with n-grams and examined different feature extraction methods and multiple machine learning models, to find the best algorithm to classify disinformation. The results showed that overall linear-based classifiers are better than nonlinear ones, with the highest accuracy achieved by a Linear SVM. Furthermore, Shu et al. [30] conducted a survey providing a comprehensive review of fake news detection on social media. They discussed existing fake news detection approaches from a data mining perspective, including feature extraction, model construction, and evaluation metrics.

For the fake news corpuses, many researchers use ready-to-use datasets, such as BuzzFeedNews<sup>3</sup>, BuzzFace<sup>4</sup>, BS Detector<sup>5</sup>, CREDBANK<sup>6</sup> and FacebookHoax<sup>7</sup> [29] and others construct their own using potentially false stories from websites marked as fake news by PolitiFact [2, 34]. Wang et al. [34] introduced LIAR, a benchmark dataset for fake news detection about politics created from manually labeled reports from Politifact.com. In this work, the authors used a Convolutional Neural Network and showed that the combination of meta-data with text improves disinformation detection. Asubario and Rubin, [2] downloaded fabricated articles from websites marked as fake news sources by PolitiFact.com and matched them with real news around the same political topics. Their computational content analysis showed that false political news articles tend to have fewer words and paragraphs than the real ones although the fabricated stories have lengthier paragraphs and include more profanity and affectivity. Finally, the titles of the fake stories are bigger and more emotional, including more punctuation marks, demonstratives, and fewer verifiable facts.

Several studies related to fake news detection examined social media aiming to extract useful features and build effective models that can differentiate potentially fabricated stories over truthful news. The study of Tacchini et al. [31] focused on whether a hoax post can be identified based on how many people “liked” it on Facebook. Using two different classification techniques, which both provided a performance of 99% accuracy, the research proved that hoax posts have, on average, more likes than non-hoax posts, indicating that the users’ interactions on news posts on social media platforms can be used to predict whether posts are hoaxes. Similarly, the study of Idrees et al. [10] showed that the users’ reactions to Facebook news-related posts are an important factor for determining if they are fake or not. The authors proposed a model based on both users’ comments and expressed emotions (emoji) and suggested that a future Support Vector Machine approach would increase its accuracy. Finally, the work of Reis et al. [24] examined features such as language use and source reliability, while also examining the social network structure. The authors studied the degree of users’ engagement and the temporal patterns and evaluated the discriminative power of the features using several classifiers with the best results obtained by a Random Forest and an XGBoost which both had an F1 score of 81%.

In line with previous work in Communication and Computational Linguistics, this study proposes that the detection of disinformation campaigns can be examined in great detail if it is treated as a classification problem, leveraging explainable machine learning models that can provide new insights on how to identify potentially misleading information. Taking previous findings into consideration, we created a model that uses content-based and engagement features

---

<sup>3</sup> <https://github.com/BuzzFeedNews/2016-10-facebook-fact-check/tree/master/data>.

<sup>4</sup> <https://github.com/gstantia/BuzzFace>.

<sup>5</sup> <https://github.com/bs-detector/bs-detector>.

<sup>6</sup> <http://compsocial.github.io/CREDBANK-data/>.

<sup>7</sup> <https://github.com/gabll/some-like-it-hoax>.

as potential predictors of disinformation. Our goal is twofold, first to examine the effectiveness of the proposed model and second provide conclusions concerning which factors predict fake news stories and especially why particular characteristics of news articles are more important in classifying them as fake. Finding answers to these questions is crucial for journalists, editors, and the audience.

### 3 Model and Feature Extraction

The main purpose of this study is to create an inclusive model to detect disinformation campaigns in (i) news articles and in (ii) Facebook news-related posts. The backbone of the model is structured based on an extensive review of previous studies in both communication and computational linguistics. In the light of the literature, we identify the following types of features:

#### 3.1 Content-Based Features

*Linguistic:* The length of the article and the length of the headline are considered good predictors for potentially false content [2,9], while the use of capitalized words in the body and title of the stories [4] along with certain POS tags such as nouns, demonstratives, personal pronouns, adverbs [2,9] help detect deceptive content. Furthermore, complexity measures like the level of lexical diversity and readability have been used in previous studies with lower levels of complexity to point to fake content [9]. Also, the high number of swear words increases the probability of an article being false [2].

*Emotional:* Emotionality is linked to disinformation in many studies [7,9] with false stories containing more negativity than real news [9] while provocative misleading content on social media has been found to express more anger in an effort to exasperate the audience [7]. In this study, we focus on two different aspects of emotionality to capture i) the actual emotion expressed in the text by measuring the intensity scores for anger, fear, sadness, joy, based on theories of basic emotions [21] and ii) the overall affect that includes the level of valence, arousal, and dominance as described by Russel [26]. The difference between emotion and affect is explained by [28], and defines the emotion as the demonstration of a feeling, whereas the intensity of the non-conscious response of the body to an experience relates to the affect.

#### 3.2 Engagement Features

Facebook likes have been identified as significant predictors of hoaxes [24,31], and users' comments and reactions to Facebook news-related posts provide patterns that can point to disinformation [10]. Hereafter, the main features of our model are explained in detail along with the rationale for their selection in Table 1:

**Table 1.** Creation of the features.

Feature	Description
<i>Content-based</i>	
<i>Linguistic</i>	
Body length	The text size in characters. Real news articles are significantly longer than fake news articles [9]
Title length	The title size in characters. The total number of words in fake news titles is higher than in real news titles [9]
Capital letters in the story	In fake news articles are used more capitalized words [9]
Parts of speech	The identification of words as nouns, verbs, adverbs, adjectives, pronouns, prepositions, conjunctions, etc. The study [22] showed that words used to exaggerate, such as superlatives, and modal adverbs are indicative of fake news. However, the survey [14] indicated that trustworthy news writers tend to use more personal pronouns, proper nouns, adverbs, numbers [22] and name entities [25]
Noun/verb	The ratio of nouns to verbs in all words of the text [15]
Lexical Diversity	Refers to the ratio of different unique words in a text [9]
Readability	The Flesch readability score indicates how easy it is for someone to read a particular text, with high readability levels associated with real news [20]
Profanity	The number of swear words is a feature of fake news [9]
Title and body similarity	The relevance of content between the title and the main body, clickbait headlines are often different from the main story [33]
Subjectivity	The quality of news is characterized by the personal author's tone, and personal opinions expressed in a text [23]. Specifically, we measured the degree of weak or strong subjectivity using the MPQA Subjectivity Lexicon [35]
<i>Emotional</i>	
Emotions	For the emotion extraction, the NRC Affect Intensity Lexicon (NRC-AIL) was used that identifies the existence of four basic emotions, anger, fear, joy, and sadness [18]
Affect	We used NRC VAD Lexicon which identifies the sentiments of valence, arousal, and dominance [17]
<i>Engagement</i>	
Likes	The number of likes of the post [10]
Love	Represents more appreciation than liking and expresses more empathy [10]
Wow	Indicates a surprising feeling that the post expresses something unexpected [10]
Haha	Represents a funny reaction, the post causes real laughter or an ironic expression [10]
Sad	Shows sadness about the post's content also is a sign of refusal [10]
Angry	Represents the disliking of the post [10]
Shares	The number of shares may be related to news content truthfulness. [8]
Comments	The total number of comments
Total interactions	The total number of all interactions
Overperforming Score	The overperforming metric is calculated automatically by CrowdTangle <sup>a</sup> based on the performance of similar posts from the same page in similar timeframes

<sup>a</sup> <https://help.crowdtangle.com/en/articles/3213537-crowdtangle-codebook>.

## 4 Method and Dataset

For this study, we collected news articles from both trustworthy and unreliable English-language websites using the Python programming language. The dataset consists of a total of 23.420 articles both real and fake that were published online during the years 2019 and 2020, covering a variety of genres. This paper focuses only on the article level, therefore characteristics such as the overall likes or followers of a Facebook page and other contextual attributes like the genre were not taken into consideration. To construct the dataset, we followed the method of [2] and retrieved 12.420 articles from three widely acknowledged fake news websites, listed in many disinformation indexes such as PolitiFact’s<sup>8</sup> fake news websites dataset and Wikipedia’s list<sup>9</sup>, namely, *dailysurge.com*, *dcgazette.com*, and *newspunch.com*. For real news, we collected a total of 11.000 articles from the following legitimate news sources: *nytimes.com*, *businessinsider.com*, *buzzfeed.com*, *newyorker.com*, *politico.com*, and *washingtonpost.com*. The articles cover various topics and include the article’s full text, title, date, author, and web address (URL). The dependent variable was calculated by setting all stories scraped from fake websites the value of 1 and the truthful articles the value of 0. Furthermore, all articles were processed for stop-words, NaN values, stemming, tokenization, and lemmatization, while articles with less than 1K characters in the main body were deleted since a lot of the fake stories were very small. The total number of articles before the cleaning was 25.020, however, only 19.340 cases were qualified for consideration in the building of the model.

Furthermore, we gathered engagement data from Facebook, through the CrowdTangle platform that belongs to Facebook, and provides access to metrics about public pages and groups. More specifically, we searched for analytics for each article in our dataset published on Facebook using the same headline or URL. However, the query was not always successful because many articles did not appear on Facebook. Thus, we matched only 4822 fake news articles from the original dataset (from *dailysurge.com*, and *newspunch.com*.) with their corresponding Facebook metrics. Finally, to have a balanced dataset we included analytics for the same amount of real articles, resulting in a total of 9.644 articles for inclusion in the model.

For feature engineering, many Python libraries were used such as the *py-readability-metrics*<sup>10</sup> package and the Natural Language Toolkit<sup>11</sup> (NLTK) to perform basic text analysis and filtering. After the features of every category (content-based, engagement-based) were created, redundant features were identified by using a correlation matrix, and the ones with a correlation higher than 0.7

<sup>8</sup> <https://www.politifact.com/article/2017/apr/20/politifacts-guide-fake-news-websites-and-what-they/>.

<sup>9</sup> [https://en.wikipedia.org/wiki/List\\_of\\_fake\\_news\\_websites](https://en.wikipedia.org/wiki/List_of_fake_news_websites).

<sup>10</sup> <https://pypi.org/project/py-readability-metrics/>.

<sup>11</sup> <https://www.nltk.org/>.

were removed from the data. Furthermore, several similar features were removed using clustering techniques. For the model, the Decision Tree and the Random Forest classifier from Scikit-learn Python library<sup>12</sup> were used, and we compared their results to find the one with the highest prediction accuracy. Afterward, the importance of each feature in this fake news classification problem was determined.

## 5 Data Analysis and Findings – In Two Distinctive Phases

For the data analysis, we separated the experiment into two phases based on the two different datasets. In Phase A, the original dataset was used for the evaluation of the importance of only the content-based features, notably the linguistic and emotional features. Then, in Phase B, a subset of the dataset that included Facebook activity (engagement features) was used twice. First, using only the engagement features as predictor variables, and then with all the features. The aim at this stage was to add the predictive power of the engagement features and check their effects on the accuracy scores. Furthermore, the overall goal of the analysis is to explore the different sets of features to be able to understand what elements of a story increase the probability of it being fake, thus we opted for models that are not complete black boxes but provide in-depth explanations of the classifier’s predictions, such as tree-based models [13]. For all the experiments, 70% of the stories were used for training and the remaining 30% for testing, and three classification methods were used for the evaluation of the model, namely F-measure (F1), precision, and recall.

### 5.1 Phase A - Evaluating the Importance of the Proposed Model Content-Based Features


For phase A of the experiment, the original dataset (fake and real articles) was used to discover the most significant content-based features that can classify an article before publication, meaning that engagement features were not being considered at this stage. The two different classification methods were applied, and the algorithm with the highest accuracy was the Random Forest classifier with an F1-Score of 91%. Our main interest lies in the feature importances of the classifier that will enable us to interpret what matters most as the model constructs its decision trees, therefore except for calculating the contribution of every feature on the prediction (see Fig. 1), we also used the ELI5<sup>13</sup> Python package for “Inspecting Black-Box Estimators” to measure the permutation importance (Table 2).

---

<sup>12</sup> <https://scikit-learn.org/stable/>.

<sup>13</sup> <https://eli5.readthedocs.io/en/latest/overview.html>.





**Fig. 1.** Feature importance score for the content-based features

Figure 1 shows the importance of the content-based features. The category of linguistic features is the most significant with capital letters in the body of the article, POS tags (nouns, adpositions, particles), lexical diversity, headline length, article length, and weak subjectivity to be amongst the top-ten important predictors. From the emotional features, arousal is the only significant attribute for detecting false content.

## 5.2 Phase B - Combining the Content-Based Features with the Engagement Features

The objective of this phase is to examine if the combination of the textual characteristics of an article (content-based features), together with audience metrics (engagement features), provides better accuracy in distinguishing the fake from real news. In this stage, we used the smaller dataset that includes the engagement features, and ran the models twice; first, we examined the performance results based only on the engagement features, and then we combined all the features. The results of the two phases are presented in Table 3. When we ran the model the first time using only the engagement features, the random forest correctly classified 95.8% of news-related posts into either fake or real class, showing that even without any textual features such as headline length or lexical diversity the model performs well based on users' interactions with the Facebook platform. Furthermore, the total number of Facebook users who "liked" the post was the most important feature, followed by the overperforming score, calculated by CrowdTangle based on the performance of similar posts from the same page in similar timeframes.

**Table 2.** Permutation Importance for the top 10 combined features

Feature	Weight
Capitals in article	0.0758 $\pm$ 0.0095
Likes	0.0732 $\pm$ 0.0092
Title’s Length	0.0595 $\pm$ 0.0038
Numbers	0.0190 $\pm$ 0.0045
Overperforming	0.0128 $\pm$ 0.0053
Arousal	0.0088 $\pm$ 0.0023
Nouns	0.0063 $\pm$ 0.0027
Comments	0.0049 $\pm$ 0.0033
Readability Score	0.0040 $\pm$ 0.0006
Strong Subjectivity	0.0040 $\pm$ 0.0016

**Table 3.** Accuracy of machine learning classifiers

Features	Measures in %	Machine Learning Classifiers	
		Decision Tree	Random Forest
News Content Features	Accuracy	84.1	91.0
Engagement Features	Accuracy	94.4	95.8
News Content Features + Engagement Features	Accuracy	94.9	98.0

As we can observe the combination of the content-based and engagement features proved to have greater predictive power compared to any single group of features. First on the top 3 of the permutation importance table (see Table 2) is the number of capital letters in the body of the article with the significance of this feature remaining stable in both datasets, while the second is the number of likes, followed by the length of the headline, which was very important also in phase A. Moreover, POS tags like numbers and nouns are significant predictors, while the overperforming score is the fifth most significant characteristic. Similar to phase A, arousal is the only emotional feature that contributes to the prediction, while the total number of comments a news post received, the readability score, and the expressed subjectivity are of lower importance.

## 6 Discussion of the Results

In general, as depicted in Table 2, the content-based features and especially the linguistic ones are the most informative for distinguishing real from fake news articles. The results are in line with previous studies [9, 15, 20, 22] which found that textual attributes can forecast the probability of a news item being deceptive. The second most important category of features is the engagement features,

with the number of Likes being the best amongst them. Interestingly, the emotional category of features is third, (that belongs to content-based) with only the Arousal being significant for the prediction. The weights of the features show what matters most for the classifications and seem to relate well with the proposed categories of features.

Overall, the findings support several features recognized in studies methodologically close to this one. More specifically, features related to words in Capital letters were highlighted in the study of Horne and Adali [9], along with the Headline Length, and the Article Length that was significant also in the work of Marquardt [15]. Facebook Likes are essential for the model's predictions and have also been identified to distinguish hoax posts [24,31], and the use of the audience reactions on the platform considered to provide patterns that can point to disinformation [10]. Furthermore, our results show that the syntax of the fake news articles is very significant, and this is one of the features recognized by many researchers in the past, specifically, that false stories include more Adverbs [2,9,22], fewer Nouns [9,15], more personal Pronouns [2,20,22], fewer Numbers [22], and more demonstratives, [2]. Additionally, Lexical Diversity and Subjectivity proved to be significant in phase A in line with the findings of [9] that false stories have less lexical complexity and more self-referential words. On the contrary, characteristics often related to disinformation like profanity, negative sentiment [15] and anger [7] were not identified by the model as significant indicators of falsity.

## 7 Conclusion and Future Work

Many studies related to disinformation in news articles and social media treat fake news detection as a text classification problem, therefore extract features and build effective models that can predict false stories [6]. Accordingly, our study employed content-based and engagement features drawn from previous theoretical constructs in an attempt to model online disinformation campaigns and cast light on its significant identifiers. To this end, we created two datasets, one that included real and fake news and a subset of the original that contained the audience's interactions to the same articles posted on Facebook. Then we performed a number of experiments, comparing the different sets of features and two tree-based classifiers. Our findings revealed that the content-based features such as Capitals in the article, Headline Length, POS tags, and the engagement feature of Facebook Likes were the most important predictors of deceptive online stories. The results provided us with insights of fake news attributes useful in the light of combating disinformation, in terms of proposing a machine learning approach to automatically detect false stories and of pointing to certain telling characteristics of these falsehoods that could be incorporated in media literacy education programs to bolster resilience against this devastating phenomenon.

However, the results of this study are based on a set of assumptions producing the following limitations. First of all, the dataset was built based on the fundamental assumption that all the articles from the sources listed as fake

news websites by Politifact are 100% fake. Undoubtedly, there are better ways of constructing a fake news corpus such as asking fact-checkers to verify the potentially deceptive stories before incorporating them into the dataset or opting for a human-in-the-loop approach where the model would not rely so heavily on Artificial intelligence but include more sophisticated human judgment. Except for the dependent variable of our model not being the optimal one, there is the limitation of the English language thus it is uncertain how the model would behave with datasets in other languages. Based on the current study, future work could use a more diverse dataset and design a study in which human fact-checkers define false stories based on certain features and their respective significance and then correlate their judgment with the feature importances of the model, or focus on rule extraction and investigate more closely the effect of each feature on disinformation detection.

**Acknowledgements.** The research work was supported by the Hellenic Foundation for Research and Innovation (HFRI) and the General Secretariat for Research and Technology (GSRT), under the HFRI PhD Fellowship grant (GA. 14540).

## References

1. Ahmed, H., Traore, I., Saad, S.: Detection of online fake news using N-Gram analysis and machine learning techniques. In: Traore, I., Woungang, I., Awad, A. (eds.) ISDDC 2017. LNCS, vol. 10618, pp. 127–138. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-69155-8\\_9](https://doi.org/10.1007/978-3-319-69155-8_9)
2. Asubiaro, T.V., Rubin, V.L.: Comparing features of fabricated and legitimate political news in digital environments (2016–2017). *Proc. Assoc. Inf. Sci. Technol.* **55**(1), 747–750 (2018)
3. Bakir, V., McStay, A.: Fake news and the economy of emotions: problems, causes, solutions. *Digit. J.* **6**(2), 154–175 (2018)
4. Bradshaw, S., Howard, P.N., Kollanyi, B., Neudert, L.M.: Sourcing and automation of political news and information over social media in the united states, 2016–2018. *Polit. Commun.* **37**(2), 173–193 (2020)
5. Commission, E.: Joint communication to the European parliament, the European council, the European economic and social committee and the committee of the regions: action plan against disinformation (2018)
6. Conroy, N.K., Rubin, V.L., Chen, Y.: Automatic deception detection: methods for finding fake news. *Proc. Assoc. Inf. Sci. Technol.* **52**(1), 1–4 (2015)
7. Freelon, D., Lokot, T.: Russian twitter disinformation campaigns reach across the american political spectrum. *Misinformation Review* (2020)
8. Granik, M., Mesyura, V.: Fake news detection using Naive Bayes classifier. In: 2017 IEEE First Ukraine Conference on Electrical and Computer Engineering (UKRCON), pp. 900–903. IEEE (2017)
9. Horne, B., Adali, S.: 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: Proceedings of the International AAAI Conference on Web and Social Media, vol. 11 (2017)
10. Idrees, A.M., Alsheref, F.K., ElSeddawy, A.I.: A proposed model for detecting Facebook news' credibility. *Int. J. Adv. Comput. Sci. Appl. (IJACSA)* **10**(7), 311–316 (2019)

11. Kovach, B., Rosenstiel, T.: *The elements of journalism: what newspeople should know and the public should expect*. Three Rivers Press (CA) (2014)
12. Lotan, G.: Networked audiences: attention and data-informed. *The New Ethics of Journalism: Principles for the 21st Century*, pp. 105–122 (2014)
13. Lundberg, S.M., et al.: From local explanations to global understanding with explainable AI for trees. *Nat. Mach. Intell.* **2**(1), 56–67 (2020)
14. Mahyoob, M., Al-Garaady, J., Alrahaili, M.: Linguistic-based detection of fake news in social media. *Forthcom. Int. J. Engl. Linguist.* **11**(1) (2020)
15. Marquardt, D.: Linguistic indicators in the identification of fake news. *Mediatization Stud.* **3**, 95–114 (2019)
16. Marwick A., Kuo R., C.S.J., Weigel, M.: *Critical disinformation studies: a syllabus* (2021)
17. Mohammad, S.: Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 English words. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (volume 1: Long Papers)*, pp. 174–184 (2018)
18. Mohammad, S.M.: Word affect intensities. arXiv preprint [arXiv:1704.08798](https://arxiv.org/abs/1704.08798) (2017)
19. Olivieri, A., Shabani, S., Sokhn, M., Cudré-Mauroux, P.: Creating task-generic features for fake news detection. In: *Proceedings of the 52nd Hawaii International Conference on System Sciences* (2019)
20. Pérez-Rosas, V., Kleinberg, B., Lefevre, A., Mihalcea, R.: Automatic detection of fake news. arXiv preprint [arXiv:1708.07104](https://arxiv.org/abs/1708.07104) (2017)
21. Plutchik, R.: A general psychoevolutionary theory of emotion. In: *Theories of Emotion*, pp. 3–33. Elsevier (1980)
22. Rashkin, H., Choi, E., Jang, J.Y., Volkova, S., Choi, Y.: 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*, pp. 2931–2937 (2017)
23. Reinemann, C., Stanyer, J., Scherr, S., Legnante, G.: Hard and soft news: A review of concepts, operationalizations and key findings. *Journalism* **13**(2), 221–239 (2012)
24. Reis, J.C., Correia, A., Murai, F., Veloso, A., Benevenuto, F.: Supervised learning for fake news detection. *IEEE Intell. Syst.* **34**(2), 76–81 (2019)
25. Rubin, V.L., Conroy, N., Chen, Y., Cornwell, S.: Fake news or truth? Using satirical cues to detect potentially misleading news. In: *Proceedings of the Second Workshop on Computational Approaches to Deception Detection*, pp. 7–17 (2016)
26. Russell, J.A.: Core affect and the psychological construction of emotion. *Psychol. Rev.* **110**(1), 145 (2003)
27. Shao, C., Ciampaglia, G.L., Varol, O., Flammini, A., Menczer, F.: The spread of fake news by social bots, vol. 96, p. 104. arXiv preprint [arXiv:1707.07592](https://arxiv.org/abs/1707.07592) (2017)
28. Shouse, E.: Feeling, emotion, affect. *M/c J.* **8**(6) (2005)
29. Shu, K., Mahudeswaran, D., Wang, S., Lee, D., Liu, H.: Fakenewsnet: a data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big Data* **8**(3), 171–188 (2020)
30. Shu, K., Sliva, A., Wang, S., Tang, J., Liu, H.: Fake news detection on social media: a data mining perspective. *ACM SIGKDD Explor. Newsl.* **19**(1), 22–36 (2017)
31. Tacchini, E., Ballarin, G., Della Vedova, M.L., Moret, S., de Alfaro, L.: Some like it hoax: automated fake news detection in social networks. arXiv preprint [arXiv:1704.07506](https://arxiv.org/abs/1704.07506) (2017)
32. Tandoc, E.C., Jr., Lim, Z.W., Ling, R.: Defining “fake news” a typology of scholarly definitions. *Digit. Journal.* **6**(2), 137–153 (2018)

33. Tromble, R.: The (MIS) informed citizen: indicators for examining the quality of online news. Available at SSRN 3374237 (2019)
34. Wang, W.Y.: “liar, liar pants on fire”: a new benchmark dataset for fake news detection. arXiv preprint [arXiv:1705.00648](https://arxiv.org/abs/1705.00648) (2017)
35. Wilson, T., Wiebe, J., Hoffmann, P.: Recognizing contextual polarity in phrase-level sentiment analysis. In: Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing, pp. 347–354 (2005)