

Credibility-Based Fake News Detection



Niraj Sitaula, Chilukuri K. Mohan, Jennifer Grygiel, Xinyi Zhou,
and Reza Zafarani

Abstract Fake news can significantly misinform people who often rely on online sources and social media for their information. Current research on fake news detection has mostly focused on analyzing fake news content and how it propagates on a network of users. In this paper, we emphasize the detection of fake news by assessing its credibility. By analyzing public fake news data, we show that information on news sources (and authors) can be a strong indicator of credibility. Our findings suggest that an author's history of association with fake news, and the number of authors of a news article, can play a significant role in detecting fake news. Our approach can help improve traditional fake news detection methods, wherein content features are often used to detect fake news.

Keywords Fake news · Misinformation · Credibility assessment · Social media · Data analysis

1 Introduction

In this digital age, news and information are mostly received from various online sources. Surveys have shown that a large number of online users depend on social media for the news: 51% in Australia, 46% in Italy, 40% in the United States, and 36% in the United Kingdom [6]. Hence, fake news can misinform many people who rely on online news and/or social media for the information.

The impact of fake news has been widely discussed after the 2016 U.S. presidential election. A study by Silverman [30] shows that for the top 20 election stories in 2016, the top 20 fake news stories had 8,711,000 shares, reactions, and comments on Facebook. These user engagement numbers were significantly higher

N. Sitaula (✉) · C. K. Mohan · J. Grygiel · X. Zhou · R. Zafarani
Syracuse University, Syracuse, NY, USA
e-mail: nsitaula@syr.edu; mohan@syr.edu; jgrygiel@syr.edu; zhouxinyi@data.syr.edu;
reza@data.syr.edu

© Springer Nature Switzerland AG 2020
K. Shu et al. (eds.), *Disinformation, Misinformation, and Fake News in Social Media*, Lecture Notes in Social Networks,
https://doi.org/10.1007/978-3-030-42699-6_9

than those for the top 20 real stories, with 7,367,000 shares, reactions, and comments on Facebook during the same time period. These concerns motivate assessing news credibility and detecting fake news *before* it spreads online.

Detecting fake news has gained attention from many academic researchers as well as other organizations. Although some fact checking websites exist, such as FactCheck¹ and PolitiFact,² the problem of detecting fake news is far from being solved. Manually verifying each and every fact in the news is extremely difficult with the high volume of data being created and shared every minute. Furthermore, it has become extremely difficult to decide whether a news article is fake or credible, since fake news articles often contain false information as well as some facts. Potthast, et al. [24], observe that fake news articles may contain facts, and credible news articles may contain factual errors. Hence, an automated process to detect fake news based only on content verification may not be effective. If we emphasize past information about the sources (authors or URLs) of news articles, then deception is still possible with new URLs and new fake author names. Our goal is to identify general indicators of news credibility, using both (1) *source* and (2) *content* perspectives to help detect fake news.

Using public data for fake news detection [27], we have analyzed multiple features of information related to the sources and contents of news articles. Our analyses demonstrate that fake news can be distinguished from true news based on features related to source and content. We also observe that while some features exhibit differences between fake and factual news, they do not help better predict fake news.

This paper focuses on finding signals or indicators of news credibility that can help detect fake news. Our findings suggest that the information about authors of news articles can indicate news credibility and help detect fake news. Using only information on the number of authors and the authors' publication history, classifiers were able to obtain >0.75 average F_1 -score. When content related features were added to these features, we observed further improvements when detecting fake news.

In the following, we detail our analysis on various aspects of credibility. We review related work in Sect. 2, followed by a brief description of the dataset used for analysis in Sect. 3. Section 4 provides our analysis of credibility based on the source of the news and Sect. 5 details credibility factors based on the content of news. Based on our analysis on source and content credibility, we build predictive models to detect fake news, which we detail in Sect. 6. We conclude in Sect. 7 and present some directions for future work.

¹<https://www.factcheck.org>

²<https://www.politifact.com/>

2 Related Work

We briefly discuss research on fake news detection and credibility assessment.

Fake News Research Fake news has been an active area of research, where it has been detected often by relying on (1) news content and/or (2) social context information. Often formed from text and images along with news sources (authors or websites), news content has been utilized in various ways to detect fake news. Text has been represented as a set of (subject, predicate, object) features and used to predict fake news by developing link prediction algorithms, i.e., how likely the extracted predicate connects the subject with the specific object [4, 10, 26]. Such textual information can be represented as style features at various language levels as well, e.g., lexicon-level [23, 32, 35, 36], syntax-level [5, 36], semantic-level [23], and discourse-level [15, 25], based on n -grams [23], Term Frequency–Inverse Document Frequency (TF-IDF) [23], Bag-Of-Words (BOWs) [36], neural network word-to-vector embeddings (Word2vec) [18, 36], Part-Of-Speeches (POSS) [5, 11, 36], Context Free Grammars (CFGs) [5, 23, 36], Linguistic Inquiry and Word Count (LIWC) [23], rhetorical relationships among sentences [15, 25], etc.; features can be explicit, i.e., non-latent features such as the frequencies of lexicons, or implicit, i.e., latent features obtained by, for example, WORD2VEC [32, 35]). Recently, news images and source websites have been used in fake news analyses. For example, to investigate news images, Jin et al. [14] defined a set of visual features to predict fake news within a traditional statistical learning framework, and Wang et al. [32] employed a deep neural network (VGG-19) to help extract the latent representation of news images. Baly et al. [1] characterized fake news articles by their source websites, e.g., if they have a Wikipedia page, if their URLs contain digits or domain extensions such as .co, .com, .gov, and their Web traffic information. Nevertheless, few research efforts have focused on the authors who create and write the [true or fake] news, which we investigate in this paper.

On the other hand, fake news detection models have emerged in recent years by studying how fake news propagates on social media (i.e., using social context information). An example can be seen in the work by Vosoughi et al. [31], which revealed that fake news spreads faster, farther, more widely, and is more popular compared to true news. Currently, methods to predict fake news have investigated the profiles of users [7], their social connections [29, 38], and posts. For instance, Guess et al. [7] found that the age of a user is an important indicator of the frequency that he or she engages in fake news activities; the analysis showed that users over 65 shared fake news approximately seven times more often compared to the younger age group. While remarkable progress has been made to achieve early detection of fake news, very little social context information on news propagation may be available; this motivates the development of approaches that can detect fake news by focusing mainly on news content.

A comprehensive survey of the various approaches for handling fake news problem is given in the work by Zhou and Zafarani [37].

Credibility Assessment The credibility of information (including news) is often evaluated by its quality and believability [3]. Research has specifically focused on assessing the credibility of social context information. TweetCred, a real-time system, scores tweet credibility by using a set of hand-crafted features within a semi-supervised learning framework [9]. Out of all selected features, the most important ones include the number of (unique) characters and words in tweets, indicating the significant correlation between the content of tweets and their credibility. In another relevant study, Gupta, et al. [8] found that a majority of the content generated at the time of crisis are from unknown sources (users) and at the same time, rumors are spread, which emphasized the importance of sources on information credibility. Based on a binary classifier, Castillo et al. [3] evaluated information credibility on Twitter using hand-crafted features from users' posting and re-posting behavior, from the text of the posts, and from citations to external sources, which can achieve a precision and recall value between 0.7 and 0.8. Within a similar framework, O'Donovan, et al. [21], discovered that features such as URLs, mentions, retweets, and tweet length were among the best indicators of the credibility using eight diverse Twitter datasets. Morris, et al. [19], found that the name of a Twitter user and the use of standard grammar, influence the credibility of tweets.

In our work, we have adopted features from earlier findings mostly from the context of microblogging sites such as Twitter, to find how closely they relate to news content credibility. We have not found any credibility study focused on number of authors related to the news, authors' collaboration relationships, and authors' past association with fake news articles. The analysis and findings of this paper provide insights to address these issues and improve fake news identification efforts.

We assess credibility from two broad perspectives: (i) Source and (ii) Content which are discussed in details in Sects. 4 and 5. For each category, we identify

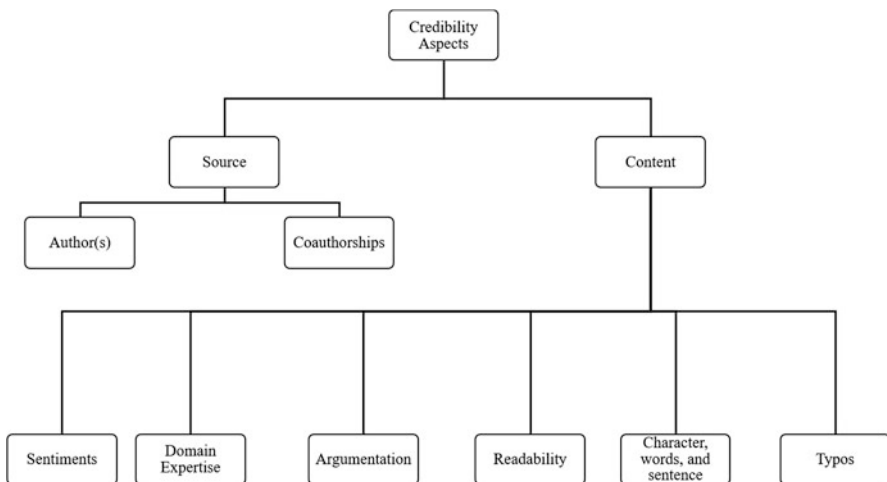


Fig. 1 Hierarchical structure of credibility aspects

Table 1 Data statistics

Data	PolitiFact	BuzzFeed
# Users	23,865	15,257
# News–Users	32,791	22,779
# Users–Users	574,744	634,750
# News stories	240	182
# True news	120	91
# Fake news	120	91

information in fake news that can capture various aspects of credibility, as shown in Fig. 1. Before further elaboration of these categories, we briefly describe the datasets used in our study.

3 Experimental Data

We have used two public datasets for fake news detection, from <https://bit.ly/2mHGnBI> [28, 29, 36, 38]. These datasets are from BuzzFeed news and Politifact. Besides news content and news labels (i.e., *fake* or *true*), the datasets contain information on the social networks of users involved in spreading the news. Statistics of the two datasets are provided in Table 1.

Out of 422 news articles in both datasets, 16 news articles had the same content/text, and were excluded from our analysis. The datasets were processed using *pandas* [17], and *matplotlib* [12] was used for visualization. In the following sections, we will discuss how different aspects of credibility can be captured from such data and demonstrate our findings on these datasets.

4 Source Credibility

In this section, we present our analysis to derive credibility from the news source, i.e. news URL, number of authors of the news, coauthorship relation to credibility, author(s) affiliations, and history of credibility of authors.

Earlier research has focused on assessing source credibility by looking at the URL associated with any news article [1], including features such as whether a website contains the `https` prefix, numbers, or `.gov`, `.co`, `.com` domain extensions. In our data, 354 news articles used the `http` prefix, 15 used the `https` prefix, and 37 had no URL. Out of the 354 URLs with `http` prefix, only 154 were fake news. Surprisingly, 14 out of 15 news articles with the `https` prefix were fake. These observations contradict past studies: having `https` in a URL does not imply credibility or help differentiate fake news from true news. We also explored whether site names can inform credibility of the news. In the dataset among those which had site names field, we found 87 unique site names with 7 site names common between

true and fake news. The total of 112 news of dataset were from these site names with 57 fake and 55 true news. It suggests that distinguishing fake news only based on site names can be extremely difficult as site names of around 27.58 % news could be associated with both fake and true news.

Other studies [3, 8, 9, 19] have shown that specific users information can be good indicators of credibility on Twitter. Hence, we seek generic user information that can capture credibility and help detect fake news. We group such *source* information into information on:

1. *Author(s)* of the news articles, and
2. *Coauthorships*, i.e., author collaborations.

We now discuss each of these subcategories in detail.

4.1 *Credibility Signals in Author(s)*

If a news article does not provide any information on its authors, its credibility can be questioned. An earlier study found that rumors mostly spread during the times of crises on Twitter and the majority of such rumors are posted by unknown sources/users [8]. However, having the name(s) of the author(s) is insufficient, because fake names or fake profiles can be easily created. Previous work has also looked at whether the news source is Wikipedia or a verified social media account, or contains other attributes to verify its credibility [1].

Thus, credibility assessment methods require multiple steps. First, we simplify credibility assessment to focus on the number of authors in the two types of news, i.e., whether it contains no authors, one author, two authors, or more. We found that the average number of authors is 0.66 for fake news and 1.97 for true news. The median number of authors is 0 for fake news and 2 for true news. Figure 2 provides the distribution of the number of authors for true and fake news.

From Fig. 2, we observe that the number of authors of a news article does have some correlation to its credibility. If an article has more than one author, it is more likely to be credible, and news with no author name is more likely to be fake news. The Pearson correlation coefficient between labels (true/fake) and number of authors is 0.406. It is difficult to draw similar inferences when news articles have only one author. From the figure, we observe that there are almost equal proportions of fake and true news for articles with a single author. A p -value of <0.05 was obtained after running Shapiro-Wilk normality test for the number of authors, which indicates that the distribution is not normal. As the distributions are not normal, we cannot compute the significance in differences in mean values between fake news and true news. Hence, we used Mann-Whitney U test on number of authors on two types of news, i.e., fake and true news. The p -value of <0.001 shows that the median number of authors in these two types of news can capture credibility. In our later analyses, we discuss ways to add past association of the authors with fake news to tackle the case when there is only one author.




Fig. 2 Frequency of number of authors in fake and true news. If an article has more than one author, it is more likely to be true news

4.2 Credibility Signals in Coauthorships

Earlier studies on coauthorship networks (e.g., Newman [20]) found that (1) a small number of influential individuals exist in such networks and (2) disconnecting such individuals from the network can result in a set of small disjoint networks. Such observations motivate us to explore whether such influential authors exist in the network of news authors, and if they do, can they help assess the credibility of the news, and their coauthors. Hence, we extended our analysis by looking at the network of news authors and classifying them into three groups: authors who are (i) only associated with fake news, (ii) only associated with true news, and (iii) associated with both fake and true news. The objective is to analyze news credibility, given the *position* of the author in this network as well as its *neighbors* (other authors). This approach allows us to understand whether fake news authors collaborate only with other fake news authors, or if they also collaborate with true news authors. We raise similar questions for true news authors and those who publish both.

Among the 237 unique authors in our data, 87 authors were authors of at least two or more news articles. To have sufficient historical data, authors whose names occurred only once (in our data) were excluded from the analysis. For simplicity, we only considered news articles whose authors were in the set of 87 authors. To provide clear insights on coauthorships, we assign these 87 authors to one of three groups:




Fig. 3 Authors collaboration network, showing authors that publish only fake news (red), only true news (green), or both (yellow)

1. *True-news authors*, only associated with two or more true news stories;
2. *Fake-news authors*, only associated with two or more fake news stories;
3. *Fake+True authors*, who have published both fake news and true news.

For these groups, the coauthorship network among authors is shown in Fig. 3, where green nodes represent *True-news authors*, red nodes represent *Fake-news authors*, and yellow nodes represent *Fake+True authors*. Dashed lines connect authors that have collaborated only once, whereas solid lines connect authors who have collaborated more than once. We notice that only 12.7% of 87 authors are involved in both fake and true news, whereas the majority were either exclusively involved in fake news or true news. We also observe that fake news authors are often either the only author (of the fake news article), or are more likely to collaborate with other *Fake-news authors* (rather than with *True-news authors* or *Fake+True authors*). We had similar observations for *True-news authors* (green nodes in Fig. 3) and *Fake+ True authors* (yellow nodes in Fig. 3).

To further investigate these observations, for each author in the coauthorship graph, we compute the number of coauthors (i.e., graph neighbors) who only post true news, only post fake news, and those who post both. Using these three numbers, we can represent any author as a 3D point and plot all authors in 3D space, as shown in Fig. 4. We observe that the credibility of authors who have had multiple coauthorships are easily distinguishable, as they often collaborate with the same type of authors. Hence, knowing the author's credibility, we can infer the credibility of coauthors. For authors with no neighbors (i.e., coauthors), they are indistinguishable.

In sum, *homophily* exist in authorship [34], where authors who write only true news are less likely to collaborate with authors who write fake news. These observations also indicate that if there are groups of authors associated with some news, by knowing credibility of any author, we may be able to infer the credibility of the news and its other authors. But, *how can we determine an author's credibility?* Two observations help us address this question:




Fig. 4 3D plot of authors. An author is represented in terms of three values: the number of coauthors of him or her who only post true news, fake news, or both

I. Affiliations provide information on credibility. Some author names are associated with known organizations such as ABC news, Associated Press, Politico, and CNN.

Hence, affiliations of authors with well-recognized organizations may indicate that a news article is not fake. While we did not explicitly consider the author relationships with the organizations, these findings support the hypothesis that working for a credible organization allows one to infer author's credibility. Similarly, we found some unrealistic author names such as Fed up, Ny evening, About the Potatriot, and About Stryker associated with fake news. This finding corroborates the earlier observation by Gupta, et al. [8], where rumors on Twitter were shown to have been spread by unknown sources/users.

II. Historical record provides information on credibility. The 87 authors that were selected were related to 172 news articles. In Table 2, we looked at how different types of authors relate to the two types of news. We observe that around 28% of the news articles have authors who post both fake and true news. However it is unclear how this information can help infer credibility of the authors or the news. To tackle this issue, we looked into the history of authors' credibility, i.e., their past associations with true/fake news articles, in order to explore whether these can capture the credibility of other articles authored later by the same authors.

Table 2 Distribution of count of news articles based on coauthors

News type	Number of authors (author type(s))	Number of articles
True news	One author (<i>True-news authors</i>)	37
True news	One author (<i>Fake+True news authors</i>)	18
True news	Multiple authors (<i>True-news authors</i>)	59
True news	Multiple authors (<i>Fake+True news authors</i>)	6
Fake news	One author (<i>Fake news authors</i>)	23
Fake news	One author (<i>Both news authors</i>)	18
Fake news	Multiple authors (<i>Fake news authors</i>)	5
Fake news	Multiple authors (<i>Both news authors</i>)	6

Not all news articles in our data had information on their publication date. As articles with no publication date could not help with historical analysis, we filtered them, reducing our data to 289 news articles. Among these 289 articles, we focused on authors of at least two news articles, which resulted in 69 authors of 163 news articles. For each author, their published news articles were sorted chronologically and we analyzed whether they contradicted their past behavior anytime in the future. We only found 11 authors that contradicted, i.e., either they had fake news in the past and were associated with true news in future, or they posted true news in the past but were associated with fake news in the future. However, the majority of authors (84%) showed consistent behavior. Thus, past information on authors' credibility provides insights on the credibility of other articles authored by them.

5 Content Credibility

Next, we discuss credibility based on the content of the news. Our goal is to explore how various characteristics of a news article content (which includes the text from the title and the body of an article) can help assess its credibility. These characteristics are compared for fake and true news articles. Previous research on credibility on Twitter [3, 9, 19, 21] has shown that there exist various indicators of credibility within content. Here, we search for such indications of credibility in the following: (1) sentiments expressed, (2) domain expertise in the article, (3) arguments used, (4) text readability, (4) characters, words, and sentences used, and (5) typos.

5.1 Credibility Signals in Sentiments

Castillo et al. [3] identified connections between sentiments expressed and credibility, whereas O'Donovan, et al. [21], found that positive sentiments may not indicate credibility in tweets. Such studies encouraged us to study the relationship between credibility and sentiments in news, answering questions such as:

- Are sentiments expressed in fake news different from those in true news?
- Is fake news written with negative, neutral, or positive sentiments?
- Can sentiments help infer credibility?

For sentiment analysis, we compute sentiment intensity (a numeric value) for each sentence in the news articles using VADER [13] sentiment analyzer available in NLTK (natural language toolkit) [2]. Using the standard threshold [13], three labels are assigned to each sentence: positive, negative, or neutral. We represent each news article with the three fractions of negative, positive, and neutral sentences in the article, e.g., $\frac{\text{number of positive sentences}}{\text{total number of sentences}}$. Some statistics on such proportions of each type of sentiment in fake and true news are in Table 3.

The mean and median values in Table 3 show that (1) proportion of neutral sentiments is slightly higher in true news compared to fake news and (2) negative sentiments proportions are higher in fake news compared to true news. However, it is still difficult to infer whether sentiments are good indicators of credibility. For further analysis, we explored whether sequence of expressed sentiments in articles differ, i.e. is fake news more likely to have sequences of sentences with positive sentiments followed by other sentences with positive sentiments? There are 9 (3×3) possible types of sequences that one can get with positive, negative, and neutral sentences. We label each sentence pair as one of these types. The mean and median proportion for each one of these nine types in both fake and true news are provided in Table 4.

Table 4 shows that positive sentences that are immediately followed by neutral sentences are more in true news compared to fake news. In fake news, negative sentences followed by other sentiments occur more often than in true news. Overall,

Table 3 Sentiment proportions in news

News type	Positive		Neutral		Negative	
	Mean	Median	Mean	Median	Mean	Median
Fake news	0.34	0.32	0.28	0.27	0.38	0.38
True news	0.33	0.33	0.34	0.31	0.34	0.33

Table 4 Statistics on sequences of sentences with different sentiments

Sentence sequence	Mean		Median	
	Fake news	True news	Fake news	True news
Positive sentence followed by a positive sentence	0.14	0.13	0.1	0.11
Positive sentence followed by negative sentence	0.10	0.09	0.10	0.09
Positive sentence followed by neutral sentence	0.09	0.11	0.08	0.10
Negative sentence followed by positive sentence	0.11	0.09	0.10	0.09
Negative sentence followed by negative sentence	0.16	0.14	0.13	0.11
Negative sentence followed by neutral sentence	0.11	0.12	0.10	0.11
Neutral sentence followed by positive sentence	0.08	0.09	0.07	0.10
Neutral sentence followed by negative sentence	0.09	0.08	0.09	0.07
Neutral sentence followed by neutral sentence	0.09	0.13	0.06	0.10

the uniformity of values in Table 4 on sentiment sequences in news articles suggests that sentiment sequences may be weak indicators of credibility. In contrast to earlier findings [3, 21], our results show that relying on sentiments alone may provide only a weak indication of credibility.

5.2 Credibility Signals in Domain Expertise

Research has shown that the presence of signal words and expert sources enhance credibility [33]. As our data was collected during the U.S. 2016 presidential election, we looked into the use of words from NCSL (National Conference of State Legislatures), which included 150 words from <https://bit.ly/1iMTzXa>. We studied whether there exist differences between fake news and true news in terms of usage frequencies of these words. We found that the average number of words from the NCSL word list were 4.37 in fake news and 7.46 in true news. The medians for the number of words were 3 for fake news, and 4 for true news. The Shapiro-Wilk test on the number of NCSL words on both fake and true news had a p -value of < 0.05 , showing the sample is not normally distributed on number of words for both types of news. With a small difference of one word, it is difficult to argue the importance of domain words/phrases, so, we further looked into distinct words that are present in one type of news and not in the other, shown in Table 5.

Later in our experiments, we will show how the occurrence of words shown in Table 5, in addition to other information, allows one to detect fake news.

5.3 Credibility Signals in Argumentation

To build strong arguments in a news article, one can rely on providing data and references.

Frequent occurrences of numbers or digits may indicate that a news article is well-researched, containing verifiable data; similarly, the occurrence of hyperlinks and URLs may indicate citations suggesting that an article is supported by external sources. The connections between URLs and credibility have been studied earlier on tweets [3, 9]. Similarly, the findings from Koetsenruijter [16] suggests that the

Table 5 List of NCSL words in fake and true news

Words in fake news, but not in true news	petition, legislator, impeachment, adhere
Words in true news, but not in fake news	fiscal, calendar, precedent, bipartisan, convene, interim, caucus, nonpartisan, statute, decorum, veto, repeal, constituent, chamber

Table 6 Distribution of digits in fake and true news

Mean		Median	
Fake news	True news	Fake news	True news
490.82	739.33	424	461

presence of numbers in an article conveys credibility. Hence, we studied whether there is a difference between fake news and true news in terms of numbers of URLs.

We found that only 18 news articles contained URLs, of which 7 were fake news and 11 were true news. With presence of URL in small proportion of dataset, it is difficult to assess credibility strength based on this feature. Table 6 shows the distribution of number of digits used in fake and true news. Our findings suggest that there are differences between fake news and true news based on the use of numbers in the news content, and that it is likely that true news is supported with facts that include numbers.

The Shapiro-Wilk test gave a p -value of <0.05 , showing that the sample is not normally distributed on number of words for both type of news. The Mann-Whitney U test shows that there is a difference in medians with a p -value of 0.011, i.e., the greater occurrence of digits in news articles indicates credibility.

5.4 Credibility Signals in Readability

The study by Horne et al. [11] suggested readability as an important feature to distinguish fake news from true news. To compare readability differences between fake and true news, we used the Flesch-Kincaid reading-ease test on the text of the news. The mean readability scores were found to be 67.32 for fake news and 65.30 for true news. Similarly, the median scores are 68.33 and 65.38 for fake news and true news, respectively. Contrary to our expectation, fake news readability was higher than that of true news. This raises a series of other interesting questions such as:

- Is fake news more readable?
- Is ease of reading why users engage more with fake news than true news?

Further analysis of the news content may reveal insights on such questions. The Shapiro-Wilk test, with p -value <0.05 indicates that the sample is not normally distributed. The Mann-Whitney U test with a p -value of 0.02 shows differences in medians and that poorer readability may indicate credibility.

5.5 Credibility Signals in Characters, Words, and Sentences

In TweetCred system [9], the number of characters and number of words were among the important features to evaluate credibility of tweets. Earlier work has also shown that tweet length is one of the indicators for credibility [9, 21]. Hence,

Table 7 Distribution of words in title and text

News type	# Words in title		# Words in text		# Sentences		# Words # Sentences	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Fake news	14.16	14.0	490.82	424.0	19.44	16.0	26.80	25.32
True news	12.09	11.0	739.33	461.0	26.59	17.0	27.69	27.50

Table 8 Distribution of characters and words

News type	# Characters		# Characters # Words		# Special characters	
	Mean	Median	Mean	Median	Mean	Median
Fake news	2052.76	1803.0	4.23	4.20	6.79	4.0
True news	3144.28	1899.0	4.27	4.27	11.04	6.0

the news content length may also be an indicator for the credibility. The mean and median of words in title and text of the news are shown in Table 7 along with the number of words per sentences. Table 7 shows that fake news text is shorter in terms of number of words compared to true news. Similarly, the length of true news articles and their number of sentences were found to be higher on average compared to fake news.

In similar research, it has been shown that presence of special characters, e.g., colon [9], exclamation mark, and question mark [3], can help assess credibility. To check if fake news contains more special characters compared to true news, we select !, #, \$, %, *, +, -, ?, @, | as special characters and count their occurrences in each type of news. Table 8 provides the mean and median of number of characters and special characters in news text, indicating that special characters are more often observed in true news.

5.6 Credibility Signals in Typos

As suggested by Morris et al. [19], the use of standard grammar and spellings enhances credibility. So, we checked if there are significant differences in typographical errors between fake and true news. To find if there are more typos in fake news compared to true news, we used the words from the NLTK corpus, containing 235,892 English words. For each news content, we counted the number of words with typos and normalized it by the total number of words in the content. Words from both content and NLTK corpus were lower-cased before checking for typos. Contrary to our expectation, we found on average 0.19 and 0.22 typos in fake news and true news, respectively. The median number of typos were 0.20 for fake news and 0.21 for true news. The Shapiro-Wilk test indicated that the sample is

not normally distributed, and using Mann-Whitney U test, we obtained a p -value <0.001 , showing that typos may indicate credibility.

6 Results

Based on our discussion on credibility aspects in Sects. 4 and 5, we built different fake news prediction models to predict fake news. From the attributes discussed, we obtained 26 features in the following categories:

- Number of authors in the news;
- Sentiments (counts of positive, neutral, and negative sentiments in text, and sequence of sentiments);
- Number of NCSL words that are only present in fake news;
- Number of NCSL words that are only present in true news;
- Flesch-Kincaid reading-ease score;
- Number of words in the title;
- Number of characters, special characters, words, sentences, digits, and typos;
- Words per sentences;
- Characters per words; and
- Past history of the author.

Using the above features, we trained a fake news prediction model with seven different classifiers (to account for learning bias) using scikit-learn package [22] and ten-fold cross validation. The F_1 -scores for these classifiers are in Table 9, where we found that the logistic regression and linear Support Vector Machines performed well among different classifiers. The best classification was achieved by logistic regression, with an 0.80 average F_1 -macro score.

Comparing Source-Credibility and Content-Credibility We studied the importance of assessing each type of credibility (source and content) by predicting fake news independently using each category of features. For source-credibility, we only considered three features: number of authors, as well as the numbers of past fake and true news stories authored by them in the past. Surprisingly, with these three

Table 9 Average F_1 scores for all features (highest values are shown in bold font)

Classifier	F_1 -micro	F_1 -macro	F_1 -weighted
SVM (RBF kernel)	0.74	0.74	0.74
Linear SVM	0.79	0.79	0.79
Logistic regression	0.80	0.80	0.80
Random forest	0.76	0.76	0.76
AdaBoost	0.74	0.74	0.74
Naive bayes	0.69	0.69	0.69
Gradient boosting decision tree	0.77	0.77	0.77

Table 10 Average F_1 scores obtained by source-credibility features (highest values are shown in bold font)

Classifier	F_1 -micro	F_1 -macro	F_1 -weighted
SVM (RBF kernel)	0.75	0.75	0.75
Linear SVM	0.75	0.75	0.75
Logistic regression	0.75	0.74	0.74
Random forest	0.76	0.76	0.76
AdaBoost	0.77	0.77	0.77
Naive bayes	0.75	0.75	0.75
Gradient boosting decision tree	0.77	0.76	0.76

Table 11 Average F_1 scores obtained by content-credibility features (highest values are shown in bold font)

Classifier	F_1 -micro	F_1 -macro	F_1 -weighted
SVM (RBF kernel)	0.64	0.63	0.63
Linear SVM	0.68	0.68	0.68
Logistic regression	0.67	0.67	0.67
Random forest	0.63	0.63	0.63
AdaBoost	0.60	0.60	0.60
Naive bayes	0.58	0.57	0.57
Gradient boosting decision tree	0.65	0.65	0.65

features, we find that the classification performance does not degrade much, as shown in Table 10. However, the best classifier was then AdaBoost, which indicates that the classifier performance is feature-dependent. Similarly, when using the 23 content-credibility features, the best F_1 -score achieved was 0.68, less than when using only source-credibility features, which achieved 0.77. The results are shown in Table 11.

By comparing the performance of source-credibility and content-credibility features, we find that assessing source credibility plays a stronger role in detecting fake news. Adding content-credibility features with source-credibility features can further improve fake news detection.

Feature Importance Analysis We also identified the most important features that can capture credibility in news.

While there can be various combinations of features to search for the optimal features, we combined both the feature selection and a hand-tailored approach (testing with trial and error), which led to 13 features with the best F_1 score: number of authors in the news and past history of authors, presence of domain words, readability, number of words, characters, special characters, and typographical errors. Table 12 shows that all classifiers performed best with these selected features, even better than using the original 26 features. Also, features that were found to be of least importance were sentiments and count of digits in the text.

Table 12 Average F_1 score with only 13 features (highest values are shown in bold font)

Classifier	F_1 -micro	F_1 -macro	F_1 -weighted
SVM (RBF kernel)	0.77	0.77	0.77
Linear SVM	0.80	0.80	0.80
Logistic regression	0.80	0.80	0.80
Random forest	0.77	0.77	0.77
AdaBoost	0.77	0.77	0.77
Naive bayes	0.75	0.75	0.75
Gradient boosting decision tree	0.77	0.77	0.77

Table 13 Average F_1 score using 13 features on separated datasets (highest values are shown in bold font)

Classifier	F_1 -macro (Politifact)	F_1 -macro (Buzzfeed)
SVM (RBF kernel)	0.79	0.75
Linear SVM	0.82	0.77
Logistic regression	0.82	0.76
Random forest	0.77	0.72
AdaBoost	0.78	0.67
Naive bayes	0.79	0.69
Gradient boosting decision tree	0.80	0.74

Comparing Tables 9 and 12, we can observe that using these 13 selected features, all the classifiers perform better than using the original 26 features. While our results did not outperform all other models as discussed in [36], where best model had a 0.892 F_1 score, our model used comparatively fewer and new features compared to the models discussed in the work.

With only 3 source-credibility features on Politifact data, the best classifier achieved an average F_1 -macro score of 0.83 and with only content-credibility features, best score was 0.66 (see Table 13). This observation shows that content-credibility has very little to add to the performance of fake news prediction in the data. Similarly, for Buzzfeed news data, the best classifier was able to obtain an average F_1 -macro score of 0.76 with only source-credibility features, whereas with content-credibility features it obtained 0.66. Thus, adding content-credibility features only slightly improved the performance. Our content-credibility features are comparatively fewer than earlier studies, so we emphasize our findings with source-credibility features, which we did not find in earlier research.

7 Conclusion

We have analyzed credibility of news, emphasizing features related to source and content of the articles. Our results based on source of the news (Sect. 4) show that number of authors of the news is a strong indicator of credibility. We found

that when the news article has no authors, it is more likely to be fake news. Our findings on collaboration of authors suggests that authors who are engaged in true credible news are less likely to collaborate with authors who are associated with fake news. This indicates that for a news article with multiple authors, by knowing the credibility of one author, we can infer the credibility of the news as well as other coauthors. Furthermore, we found that authors' affiliations with well-recognized organizations can be a signal for credibility. The results also suggest that credibility history of authors can provide insights on credibility of other articles from the same author.

Similarly, we investigated credibility based on various content-related aspects of the news (Sect. 5). The results show that sentiments expressed in news articles are weak indicators of credibility. We observed that the use of numbers in true news articles occurred more often than in fake news, perhaps because true news is supported with facts that include numbers. Comparing the number of words and sentences in true news and fake news showed that on average, true news had more words and sentences than fake news. Surprisingly, we observed more typos in true news than in fake news. Our analyses also showed that domain expertise on topics discussed in news can enhance fake news detection.

After analysis of individual features, we used our findings to build predictive models to detect fake news. The F_1 -score of 0.80 obtained by predictive models built with source-credibility features show that with a small number of features, one can still detect fake news reasonably well. Using fewer features can lead to less complex models. Hence, our simple approach provides a straightforward fake news detection framework with a few features that can quickly detect fake news.

Stronger conclusions require further research on additional machine learning features, other predictive models, and datasets. We have not yet explored word-based sentiments in our analysis, where one can consider negated positive words, or number of negative and positive words in sentences. Another avenue to explore is to study sentiments in paragraphs, which may show less variation compared to our results. Furthermore, the news content can include images (and other media), as well as the number of user interactions, which may provide more insights on the differences between fake and true news.

References

1. Baly, R., Karadzhov, G., Alexandrov, D., Glass, J., Nakov, P.: Predicting factuality of reporting and bias of news media sources. arXiv preprint arXiv:1810.01765 (2018)
2. Bird, S., Klein, E., Loper, E.: Natural Language Processing with Python: Analyzing Text With the Natural Language Toolkit. O'Reilly Media, Inc, Beijing (2009)
3. Castillo, C., Mendoza, M., Poblete, B.: Information credibility on twitter. In: Proceedings of the 20th International Conference on World Wide Web, pp. 675–684. ACM (2011)
4. Ciampaglia, G.L., Shiralkar, P., Rocha, L.M., Bollen, J., Menczer, F., Flammini, A.: Computational fact checking from knowledge networks. PLoS One **10**(6), e0128193 (2015)

5. Feng, S., Banerjee, R., Choi, Y.: Syntactic stylometry for deception detection. In: Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers, vol. 2, pp. 171–175. Association for Computational Linguistics (2012)
6. Fletcher, R., Nielsen, R.K.: Are people incidentally exposed to news on social media? a comparative analysis. *New Media Soc.* **20**(7), 2450–2468 (2018)
7. Guess, A., Nagler, J., Tucker, J.: Less than you think: prevalence and predictors of fake news dissemination on facebook. *Sci. Adv.* **5**(1), eaau4586 (2019)
8. Gupta, A., Kumaraguru, P.: Twitter explodes with activity in Mumbai blasts! a lifeline or an unmonitored daemon in the lurking? Technical report (2012)
9. Gupta, A., Kumaraguru, P., Castillo, C., Meier, P.: Tweetcred: real-time credibility assessment of content on twitter. In: International Conference on Social Informatics, pp. 228–243. Springer (2014)
10. Hassan, N., Arslan, F., Li, C., Tremayne, M.: Toward automated fact-checking: detecting check-worthy factual claims by claimbuster. In: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1803–1812. ACM (2017)
11. Horne, B.D., 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: Eleventh International AAAI Conference on Web and Social Media (2017)
12. Hunter, J.D.: Matplotlib: A 2d graphics environment. *Comput. Sci. Eng.* **9**(3), 90 (2007)
13. Hutto, C.J., Gilbert, E.: Vader: a parsimonious rule-based model for sentiment analysis of social media text. In: Eighth International AAAI Conference on Weblogs and Social Media (2014)
14. Jin, Z., Cao, J., Zhang, Y., Zhou, J., Tian, Q.: Novel visual and statistical image features for microblogs news verification. *IEEE Trans. Multimedia* **19**(3), 598–608 (2017)
15. Karimi, H., Tang, J.: Learning hierarchical discourse-level structure for fake news detection. arXiv preprint arXiv:1903.07389 (2019)
16. Koetsenruijter, A.W.M.: Using numbers in news increases story credibility. *Newspaper Res. J.* **32**(2), 74–82 (2011)
17. McKinney W.: Data Structures for Statistical Computing in Python, Proceedings of the 9th Python in Science Conference, 51–56 (2010)
18. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 (2013)
19. Morris, M.R., Counts, S., Roseway, A., Hoff, A., Schwarz, J.: Tweeting is believing? understanding microblog credibility perceptions. In: Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work, pp. 441–450. ACM (2012)
20. Newman, M.E.: Coauthorship networks and patterns of scientific collaboration. *Proc. Natl. Acad. Sci.* **101**(suppl 1), 5200–5205 (2004)
21. O’Donovan, J., Kang, B., Meyer, G., Höllerer, T., Adalii, S.: Credibility in context: an analysis of feature distributions in twitter. In: 2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing, pp. 293–301. IEEE (2012)
22. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011)
23. Pérez-Rosas, V., Kleinberg, B., Lefevre, A., Mihalcea, R.: Automatic detection of fake news. arXiv preprint arXiv:1708.07104 (2017)
24. Potthast, M., Kiesel, J., Reinartz, K., Bevendorff, J., Stein, B.: A stylometric inquiry into hyperpartisan and fake news. arXiv preprint arXiv:1702.05638 (2017)
25. Rubin, V.L., Lukoianova, T.: Truth and deception at the rhetorical structure level. *J. Assoc. Inf. Sci. Technol.* **66**(5), 905–917 (2015)
26. Shi, B., Weninger, T.: Discriminative predicate path mining for fact checking in knowledge graphs. *Knowl. Based Syst.* **104**, 123–133 (2016)

27. Shu, K., Mahudeswaran, D., Wang, S., Lee, D., Liu, H.: Fakenewsnet: a data repository with news content, social context and dynamic information for studying fake news on social media. arXiv preprint arXiv:1809.01286 (2018)
28. Shu, K., Sliva, A., Wang, S., Tang, J., Liu, H.: Fake news detection on social media: a data mining perspective. *ACM SIGKDD Explor. Newsletter* **19**(1), 22–36 (2017)
29. Shu, K., Wang, S., Liu, H.: Beyond news contents: the role of social context for fake news detection. In: *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pp. 312–320. ACM (2019)
30. Silverman, C.: This analysis shows how viral fake election news stories outperformed real news on facebook (2016). <https://www.buzzfeednews.com/article/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook#.vtQpz9DKd>
31. Vosoughi, S., Roy, D., Aral, S.: The spread of true and false news online. *Science* **359**(6380), 1146–1151 (2018)
32. Wang, Y., Ma, F., Jin, Z., Yuan, Y., Xun, G., Jha, K., Su, L., Gao, J.: Eann: event adversarial neural networks for multi-modal fake news detection. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 849–857. ACM (2018)
33. Wogalter, M.S., Kalsher, M.J., Rashid, R.: Effect of signal word and source attribution on judgments of warning credibility and compliance likelihood. *Int. J. Ind. Ergon.* **24**(2), 185–192 (1999)
34. Zafarani, R., Abbasi, M.A., Liu, H.: *Social media mining: an introduction*. Cambridge University Press, New York (2014)
35. Zhang, J., Cui, L., Fu, Y., Gouza, F.B.: Fake news detection with deep diffusive network model. arXiv preprint arXiv:1805.08751 (2018)
36. Zhou, X., Jain, A., Phoha, V.V., Zafarani, R.: Fake news early detection: a theory-driven model. arXiv preprint arXiv:1904.11679 (2019)
37. Zhou, X., Zafarani, R.: Fake news: a survey of research, detection methods, and opportunities. arXiv preprint arXiv:1812.00315 (2018)
38. Zhou, X., Zafarani, R.: Network-based fake news detection: a pattern-driven approach. arXiv preprint arXiv:1906.04210 (2019)