RESEARCH ARTICLE

JASIST WILEY

Detecting fake news stories via multimodal analysis

Vivek K. Singh ⓘ | **Isha Ghosh** ⓘ | **Darshan Sonagara**

School of Communication and Information, Rutgers University, New Brunswick, New Jersey

Correspondence
Vivek K. Singh, School of Communication and Information, Rutgers University, 4 Huntington Street, New Brunswick, NJ 08901.
Email: vivek.k.singh@rutgers.edu

Abstract

Filtering, vetting, and verifying digital information is an area of core interest in information science. Online fake news is a specific type of digital misinformation that poses serious threats to democratic institutions, misguides the public, and can lead to radicalization and violence. Hence, fake news detection is an important problem for information science research. While there have been multiple attempts to identify fake news, most of such efforts have focused on a single modality (e.g., only text-based or only visual features). However, news articles are increasingly framed as multimodal news stories, and hence, in this work, we propose a multimodal approach combining text and visual analysis of online news stories to automatically detect fake news. Drawing on key theories of information processing and presentation, we identify multiple text and visual features that are associated with fake or credible news articles. We then perform a predictive analysis to detect features most strongly associated with fake news. Next, we combine these features in predictive models using multiple machine-learning techniques. The experimental results indicate that a multimodal approach outperforms single-modality approaches, allowing for better fake news detection.

1 | INTRODUCTION

The rapid spread of misinformation and fabricated content online has drawn increasing attention in the information science community over the past few years (Conroy, Rubin, & Chen, 2015). As an example, Allcott and Gentzkow (2017) showed the significant role of fake news in the outcome of the 2016 U.S. presidential elections; Gupta, Lamba, Kumaraguru, & Joshi, 2013 analyzed the most viral tweets related to the Boston Marathon blasts in 2013 and found that rumors and misleading content were more likely to be shared than true information. The quick and far-reaching spread of fake news can have serious negative consequences for both individuals and society. First, an increasing presence of fake news can break the authenticity balance of the news ecosystem. Second, fake news persuades readers to accept biased or false beliefs and hence is often used by propagandists to convey political messages or influence

(Khaldarova & Pantti, 2016). To help mitigate these negative effects caused by fake news, to benefit both the public and the news ecosystem, it is critical to build tools that can automatically detect and flag fake news.

The traditional way of verifying online content, that is, via "manual" knowledge-based fact checking, is made difficult—or practically impossible—by the enormous volume of information that is generated online and the rapid rate of diffusion (Conroy et al., 2015). The speed and ease by which information is created and disseminated requires automated approaches that can be used at scale to detect fake news or at least prioritize which articles need closer examination by human experts.

In past research, fake news detection has been studied in two ways: (a) analyzing the *spread* of fake news and (b) analyzing the *content* of fake news. In studying the spread of fake news Tacchini, Ballarin, Della Vedova, Moret, and de Alfaro (2017) show that an analysis of "likes" received by a news article can be used to detect

fake news. Similarly, Bessi et al. (2015) studied the spread of a fake news item across different websites and the attention it received from various users. They find that users who often interact with nontraditional media (e.g., conspiracy sites) are more prone to share misinformation. Along with network analysis techniques, content analysis has also been used to detect fake news. Rubin, Chen, and Conroy et al. (2015) use multiple textual analysis techniques to classify a news item as fake or credible. In an analysis of real versus doctored images, Jin, Cao, Zhang, Zhou, and Tian (2017) identify visual features (clarity, similarity, clustering) to classify whether an image is doctored.

While existing research typically identifies textual or visual features and/or affordances of social media sites to flag an item as fake, it is important to understand the overall composition of a news item as an integration of text and visual content and analyze the credibility of a news story as a *whole* rather than identifying the fakeness of separate components. In theoretical terms, the *Framing Theory* suggests that the way an event or issue is presented defines the issue (De Vreese, 2005). At a broad level, framing deals with the issues of *selection* and *salience*. Simply put, framing is a means of using language and placement as a tool to highlight certain information as being the most important or relevant while leaving out other bits or representing them as unimportant. While framing research has mostly focused on verbal and written framing, Messaris and Abraham (2001) studied visual frames in the portrayal of the African American in media stories and discussed the importance of visual framing. Specific images are often used alongside text to convey an unwritten idea to the audience (Messaris & Abraham, 2001). This implies that textual and visual modalities work together in framing a news item.

According to this theory, a frame is created as a package of key ideas, stock phrases, and images to bolster a particular interpretation of the event. Through strategic repetition and reinforcement, the texts and images provide a dominant interpretation more readily perceivable, acceptable, and memorable than other interpretations. In the case of fake news articles, this ability to implicitly influence readers through the juxtaposition of images and text has the advantage of providing deniability to the author or website, making it an effective strategy for the manipulation of ideas. At the same time, it is harder for automated systems relying only on textual or visual analysis to identify these articles as fake news or misinformation. While the importance of understanding multimodal content for fake news detection has been acknowledged, a majority of the research in this area is still focused on (sophisticated) text processing (e.g., Horne & Adali, 2017; Singhania, Fernandez, & Rao, 2017). Recently, fueled by the

growth in deep learning architectures, there have been multiple attempts at detecting visual fake content (Jin et al., 2017), and lately, some efforts have utilized textual and visual content to detect misinformation in the context of short microblog posts (Wang et al., 2018). However, there is very little theory-driven development on the need for multimodal news detectors and empirical evidence on the value of multimodal approach for full-length news stories, which remain a primary mode for news transmission. This is an important gap in research literature, given that news articles are increasingly becoming multimodal in the last few years (Begley, 2017). Figure 1 shows two such examples of multimodal news stories, each involving a combination of text and image, which is a very common way to present news stories today.

In this paper, we focus on the automatic identification of fake content in online news stories and study the following research questions:

RQ1 *Which visual and textual features are statistically different between fake and credible multimodal news stories?*

RQ2 *How does a multimodal (text + visual features) approach for fake news detection perform compared to unimodal fake news detection approaches?*

This work uses a "Fake News Dataset" from Kaggle (2017) along with a data set of credible news stories and uses textual and visual analysis to derive features. These features are then combined using machine-learning algorithms to build an automated fake news detector.

2 | RELATED WORK

There have been multiple efforts at understanding, detecting, and preventing fake news (Gupta et al., 2013; Horne & Adali, 2017; Rubin, Chen, & Conroy, 2015; Shu, Wang, & Liu, 2019; Starbird et al., 2016). In this section, we present related works that use automated methods for fake news detection.

First, it is important to note that there have been multiple discussions on the definition of fake news. A widely accepted definition for fake news is "news articles that are intentionally and verifiably false" (Conroy et al., 2015). This definition uses both *intent* and *authenticity* as markers to flag fake news. However, research has also studied satire as a form of fake news even though it is not intentionally written to mislead viewers (Burfoot & Baldwin, 2009). Therefore, for the purpose of this study, we adopt a broader definition that focuses only on the
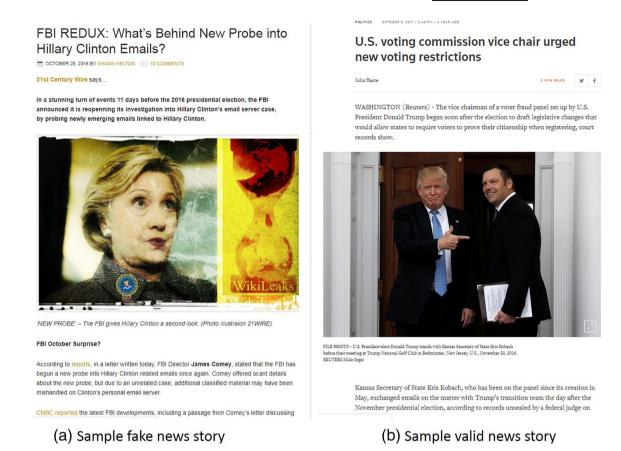
FIGURE 1 Examples of valid and fake news articles considered in this work. (a) Sample fake news story. (b) Sample valid news story. Sources: (a) https://21stcenturywire.com/2016/10/28/fbi-redux-new-probe-into-hillary-clinton-emails/; (b) https://www.reuters.com/article/us-usa-election-kobach/u-s-voting-commission-vice-chair-urged-new-voting-restrictions-idUSKBN1CB2G3 [Color figure can be viewed at wileyonlinelibrary.com]

authenticity of a news item and not its intent, that is, any news story that contains false or inaccurate information is considered "fake news."

Detecting fake news is a complex and multi-dimensional task due to the characteristics of fake news. The detection strategies exploit multiple content-related features (e.g., headline, body text, publisher) and network-related features (e.g., feedback, propagation paths, and spreaders) as cues to detect fake news.

Style-based fake news detection aims to capture and quantify the differences in writing styles between fake and credible news stories. Researchers have examined textual and image properties for this detection. For instance, studies have proposed stylometric analysis (Feng, Banerjee, & Choi, 2012) and linguistic approaches (Pérez-Rosas, Kleinberg, Lefevre, & Mihalcea, 2017) to detect deceptive text on crowdsourced data sets. In a similar vein, authors have shown the value of deep learning approaches for analyzing text at word, sentence, and headline levels for fake news detection (Singhania et al., 2017). Given the impact of social media on the large-scale diffusion of fake news, scholars have also studied the role played by images in the diffusion of fake news. Scholars have used image forensics, content features, and user characteristics to automatically predict whether an image-based tweet is fake or real (Boididou et al., 2015).

A second line of research has used network features or the path through which a news item is spread to assess its credibility. Different groups have designed approaches that utilize contextual information regarding the person speaking/posting to augment the deep learning approaches for fake news detection (Dong et al., 2018; Shu et al., 2019). Images present vivid descriptions of the situation and often attract more attention than a pure text article (Gupta et al., 2013). Given this importance of visual media, images circulated online have been studied to provide clues for the identification of fake news. Fake news articles often exploit individual vulnerabilities and therefore rely on sensational or fake images to provoke anger or an instinctive emotional response in viewers (Jin et al., 2017). Gupta et al. (2013) performed a characterization analysis to understand the temporal, social reputation, and influence patterns for the spreading of fake

images on microblogs. While these approaches are valuable and provide useful insights for fake news detection, they are limited by being restricted to a single modality.

However, it is important to take into account the multimodal nature of today's news stories. Information can be packaged and distributed in many different ways, and it is important to be critical about how and why certain information is being presented in the way that it is. For instance, a news story deliberately crafted to highlight certain key phrases or draw attention toward polarizing images with super-imposed text is likely to evoke a strong emotional response toward the entire news item (Messaris & Abraham, 2001). Therefore, it is important that we design fake news detection models that are able to process different dimensions of the overall news item. Recent attempts at rumor detection and misinformation detection in *microblogs* (e.g., tweets or Weibo posts) have reported the value in using a multimodal approach (Jin et al., 2017; Wang et al., 2018). However, the use of multimodal approaches remains scarce, and its application in the context of full-length news stories (not microblogs) remains empirically untested. Furthermore, this study uses a (information science and media studies) theory-driven perspective for adopting a multimodal framework, which is scarce in the literature.

Figure 1 presents a side-by-side comparison of valid and fake news stories. While the news stories are similar in terms of the modalities used (text and image) and the broad topic covered (politics), the framing of the story and the aspects that attract immediate attention are quite different. While the sample fake news story uses edited images (or "illustrations" in their terminology) and poses questions in the text, the valid news story uses a regular, higher-quality image and more traditional language.

We use multiple theories to ground the analysis of fake news stories in this work. *The Limited Capacity Model of Mediated Motivated Message Processing (LC4MP)* (Lang, 2009) states that individuals are inherently limited in their ability to process the various modalities of information that they encounter. This model is especially useful in understanding how individuals process information-rich, multimodal messages. According to this model, there are a number of conscious and unconscious mechanisms that determine what information an individual will register. Lang's (2009) model has been applied most typically to television-viewing situations to predict what viewers will remember based on various manipulations in structure and content changes in the medium, as well as inducing different motivations in the viewer (Christensen, Bickham, Ross, & Rich, 2015). This implies that people will respond selectively to certain channels of information in each modality, often in proportion to the intensity of the stimulus. For example, intensely negative stimuli such as offensive language or violence or one that challenges the beliefs of the individual will trigger a greater response than normal conversation or a selfie (Fisher, Huskey, Keene, & Weber, 2018; Lang, 2009). Thus, it is important that a machine-learning model is able to capture a wide range of information channels so that it is able to process comprehensively the salient features within each modality to classify a news item as fake. We therefore use this model to identify textual and visual features that could assist the detection of a fake news item.

According to the *Elaboration Likelihood Model (ELM)*, individuals can be convinced to accept information as true using either the *central* or *peripheral route* (Petty & Cacioppo, 1986). The central processing route requires a careful evaluation of the logic and strength of the arguments presented and an objective assessment of the veracity of a news item. The peripheral (heuristic) route, on the other hand, requires individuals to rely more on heuristic cues such as the perceived agreement of the story with their personal biases or source credibility rather than the actual content of the story (Petty & Cacioppo, 1986). The central route requires individuals to invest significant time and cognitive effort and make a rational and objective assessment to assess the veracity of a news item. The peripheral route, on the other hand, relies on triggering strong emotions and therefore requires very little cognitive effort. For example, if a news item is designed using language and visuals that trigger emotions of pride, fear, or anxiety, individuals are more likely to be swept up in the narrative and suspend rational assessment. Therefore, when attempting to detect fake news, a machine-learning model needs to take into account the emotion conveyed by the news item. When identifying content features, the ELM is helpful in identifying how text or visual properties maybe used to target the peripheral processing route and persuade individuals to accept a misleading news item. Finally, the *Framing Theory*, as described earlier, suggests that the way an event or issue is presented defines the issue (De Vreese, 2005), and this could have implications for the way features are identified for fake news detection.

3 | PROPOSED APPROACH

This work uses the Fake News data set available on the "Kaggle Fake News Dataset" (Kaggle, 2017). The data set contains text and metadata from 244 websites and represents 12,999 news stories gathered using the "BS Detector" chrome extension built by Daniel Sieradski and labeled by users of the plugin as "bs," "biased," "conspiracy" etc. For the purposes of this work, all these articles are considered "fake news" examples. This data set has

been used in multiple recent efforts to detect fake news, including Ahmed, Traore, and Saad (2017) and Pan et al. (2018).

The data set contained some articles with inactive URLs and broken links, which had to be removed from the analysis. The focus of this work lies on multimodal fake news detection; hence, we focus only on news stories that include both text and images. After removing inactive links, only textual articles, and articles with no attributed sources, we were left with 6,022 articles that contained visual and textual data. As an added preprocessing step, we looked at the image content of the article and removed articles where the only image was a logo of the website it had been published on or if it was not a news article (e.g., advertisement, cooking recipe). After the cleaning and preprocessing, we were left with *3,568 articles* as our final data set of fake news articles.

This data set needs to be compared and contrasted with another data set that includes examples of credible news stories. Therefore, we created another data set of credible stories, which includes an equal number of news reports from three news sources: *The New York Times*, Reuters, and Public Broadcasting Service (PBS). The three news sources were selected based on the ratings given to them by the NewsGuard agency while also trying to maintain some variety across the sources. NewsGuard (https://www.newsguardtech.com/) utilizes trained journalists to assess the credibility and transparency of news websites. Their focus is on the use of trained experts rather than algorithms to determine the credibility of sources. They also allow respective news outlets to comment on the assigned ratings before making them public. Their analysis produces nine granular, binary labels for each site, with a points system that is used to derive an overall label for credibility. NewsGuard's methodology is transparent, and it publishes a policy for ethics and conflicts of interest (Nørregaard, Horne, & Adalı, 2019). All the three sources identified—*The New York Times*, Reuters, and PBS—scored the highest rating possible for a combined metric on the credibility of the news source (Nørregaard et al., 2019). While *The New York Times* is a major print and online news house, Reuters is a news agency whose content is utilized by multiple newspapers, and PBS is an American public broadcaster and television program distributor. While all three sources scored very highly on credibility, they were also dissimilar with each other in terms of their focus and hence were selected as candidates for sourcing credible news articles.

We built a data set for credible news stories by collecting 10,000 articles from each of these sources. This data set of 30,000 news articles shrunk to 15,915 credible articles after the data-cleaning process in which articles posted as advertisements or those that did not have images were removed.

4 | FEATURE DESIGN

In this work, we consider three approaches for fake new detection—one based only on textual features, one on visual features, and one combining both.

The features have been identified based on a combination of the theoretical concepts identified in the previous section (Framing Theory, LC4MP, and ELM) and the array of recent empirical results on fake news detection. Based on this review, we identified four broad categories of features: *Content*, *Organization*, *Emotions*, and *Manipulation*.

Content refers to the topics being covered in the text of the news stories or the objects and scene labels assigned to the images in the news stories. This corresponds to the Issue Selection component as per the Framing Theory and Central Persuasion Route as per the ELM. *Organization* refers to how the above content is organized and presented to the audience. This includes features such as the sentence complexity score; words per sentence for the text; and the size of image, height, width, etc. of the visual features. They correspond to the Issue Salience component as per the Framing Theory and the Peripheral Persuasion Route as per the ELM. *Emotions* refer to the emotion-arousing aspects of the news stories. This includes features such as positive/negative emotion scores or the use of explicit content for the text and the facial expressions, nudity, violence, and gore in the visual content. LC4MP and the ELM: Peripheral Route both suggest the use of emotion-arousing content to spread fake news. *Manipulation* refers to the distortion of the content in an effort to convey a certain viewpoint. The corresponding features include the excessive use of personal pronouns in text or the use of image superimposition, edited text added to images, or the sharpness/blur quality of the image. This corresponds to the Peripheral Persuasion Route as per ELM. We therefore find that the LC4MP and ELM theories are both useful in selecting features that might be useful in determining the veracity of a news item.

The features are summarized in Table 1. A total of 124 features were identified with 81 textual features extracted from the Linguistic Inquiry and Word Count (LIWC) tool and 43 visual features extracted from images associated with news stories from three different Application Programming Interfaces (APIs): Google Vision, Amazon Rekognition, and Clarifai. Note that we acknowledge that some features can be interpreted under more than one category, but in this work, we choose to limit each feature to what we believed to be the most relevant category.

TABLE 1 Summary of textual and visual features used for analysis

Type of feature	Theoretical support	Sample text features	Sample visual features	Empirical literature support
Content	Framing: Issue selection ELM: Central route	Topics of personal concern, for example, work, leisure, home, money, religion, death; social words, for example, family, friends, female referents, male referents, perceptual processes, biological processes, drives and needs; words focusing on past, present, future	Topics of image labels (x9), contrast, brightness, colors present (x3), dominant color in the focused area (x3), fraction representation of colors (x3), number of faces (x3), age, gender, celebrity presence (x4)	Shu, Sliva, Wang, Tang, & Liu (2017); Pérez-Rosas et al. (2017); Hong (2013)
Organization	Framing: Issue salience ELM: Peripheral route	Word count, words per sentence, six-letter words, sentence complexity, use of punctuations (e.g., question marks, exclamation points) quotes, use of negations (e.g., no, never, not)	Size, height, width	Conroy et al. (2015); Rubin, Conroy, Chen, & Cornwell (2016); Hong (2013); Horne & Adali (2017)
Emotions	ELM: Peripheral route LC4MP	Affect words, emotional tone, anxiety, anger, sadness, swear words, assent	Emotions portrayed (x8), adult content, medical content—including blood, violent content, smile (x2)	Jin et al. (2017); Gupta et al. (2013)
Manipulation	Framing: Issue salience	Personal pronouns, impersonal pronouns, discrepancy, tentativeness, certainty	Spoofed content, image superimposition, visual text (x2), image sharpness/blur	Pantti & Sirén (2015); Horne & Adali (2017); Lin, He, Tang, & Tang (2009); Farid (2006)

4.1 | Text-based features

Linguistic approaches to identify fake news typically rely on language usage and its analysis (Feng, Banerjee & Choi, 2012). In the current study, we treat each news story as a document and perform text-based analysis.

Content: This refers to the topics covered in the news stories. Specifically, following Horne and Adali (2017), we focus on LIWC-based categories of textual content. An extensive list of LIWC-based features is available in Tausczik and Pennebaker (2010), including features such as the number of words related to the past or present and work or relaxation.

Organization: These features are based on natural language processing to understand the syntax, text style, and grammatical elements of each article content and title. We therefore use punctuations (e.g., question marks, exclamation points), quotes, use of negations (e.g., no, never, not), and grammar to analyze news stories. We further analyze the overall complexity of the news story. The importance of the impression and relatability of the overall article is stressed by both the framing theory and previous research investigating fake news (Hong, 2013; Horne & Adali, 2017).

Emotions: The presence of emotion also forms an important part of the narrative and provides insights complementary to the purely factual information (Messaris & Abraham, 2001). The ELM considers sentiment polarity a major determinant of informational quality (Osatuyi & Hughes, 2018). We therefore identify features associated with the emotion or sentiment of the text. We use LIWC dictionaries to measure affect words designed to evoke an emotional response from the audience.

Manipulation: Prior literature suggests that excessive use of second-person pronouns (e.g., "You") can be used to persuade the reader in specific ways and is less likely to be used in objective news articles (Horne & Adali, 2017). Furthermore, words associated with discrepancy, tentativeness, and certainty could be associated with fake news. Hence, we consider these features part of our model.

4.2 | Visual features

Images are also considered powerful framing tools because they are less intrusive than words and require less cognitive load. Therefore, the peripheral, rather than central, processing route may be activated, and audiences will be more likely to accept the visual frame without question (Messaris & Abraham, 2001).

Content: We examine the content of images as a clue to identify a news item as fake. Some of the features identified are the objects shown in the image, the presence and number of faces, and the scene labels, which were clustered into topics or themes based on the Latent Dirichlet Allocation (LDA) method (Blei, Ng, & Jordan, 2003). We also use an analysis of color components and properties of images to serve as an indicator of fake news.

Organization: We use features such as size, width, etc. of the image as cues identifying fake news. Past research with images on microblogs has shown these to be associated with fake news (Papadopoulou, Zampoglou, Papadopoulos, & Kompatsiaris, 2017).

Emotions: These features include expressions or emotions portrayed by faces in the image and the presence of violence. Previous literature has shown the importance of visual imagery when attempting to convey or underline the overall news story (Messaris & Abraham, 2001). A comparison of credible and fake news in literature shows fake news to be eye-catching and visually striking (Pantti & Sirén, 2015). Most fake images also depict disturbing events, such as an accident, abuse, injury, and conflict (Jin et al., 2017). Based on this literature, we analyze visual content for the presence of adult or pornographic content, explicit images depicting blood and gore, and violence.

Image Manipulation: As shown in existing literature, the presence of tampering or manipulation is a strong indicator of fake news (Lin et al., 2009). Here, a tampered image is one where a part of the original image has been manipulated. The presence of text superimposed on an image could also serve as an indicator of fake news.

5 | EXPLORATORY ANALYSIS

One of the goals of this work is to explore how various textual and visual features relate to the fakeness or credibility of news stories. As a first step, we applied the standard normalization technique (*value − minimum value/[maximum value − minimum value]*) to map the values of these features on a scale of 0–1. Next, in order to find positively and negatively associated features with fake and credible classes, we calculated the relative difference (where the negative value tends toward fake class and the positive value tends toward credible class) for each feature for both the classes. We present the top 20 most significant features in separating fake news from credible news in Table 2.

5.1 | Textual features

As shown in Table 2, we found that, in terms of the "content," a higher number of words pertaining to the present, higher use of pronouns and the higher number

TABLE 2 Top 20 features with the most differences in values observed in fake news and credible news categories

Modality	Category	Feature	Credible average	Fake average	Percentage increase or decrease = 100 * (credible-fake)/fake
Textual features	Content	**Words related to present**	**0.2098**	**0.2998**	**−30.04**
	Content	**Total pronouns (e.g., he, they, her, your)**	**0.2569**	**0.3242**	**−20.76**
	Content	**Auxiliary verbs (e.g., be, do, have)**	**0.3048**	**0.3768**	**−19.11**
	Content	*Words related to work*	*0.2386*	*0.169*	*41.2*
	Content	*Words related to past*	*0.2984*	*0.2215*	*34.7*
	Organization	**Filler words**	**0.0018**	**0.0049**	**−64.18**
	Organization	*Word per sentence*	*0.1312*	*0.1001*	*31.1*
	Emotions	**Swear words**	**0.0168**	**0.0384**	**−56.27**
	Emotions	**Anger words**	**0.0622**	**0.0961**	**−35.26**
	Emotions	**Sexual words**	**0.0169**	**0.0237**	**−28.85**
	Manipulation	**Personal pronoun—You**	**0.0208**	**0.0533**	**−60.95**
Visual features	Content	*Most common color: [red, green, blue] components.*	*[0.34, 0.33, 0.30]*	*[0.19, 0.15, 0.33]*	*[76.8,118.9,12.4]*
	Content	*Dominant color in the focused area: [red, green, blue] components.*	*[0.55, 0.55, 0.60]*	*[0.48, 0.43, 0.41]*	*[14.4, 26.4, 45.9]}*
	Content	Topic {#2,#4,#8,#9} from LDA modeling of labels	{0.07,0.08, 0.10, 0.17}	{0.15, 0.03, 0.17, 0.09}	{−52.5, 125.3, −42.0, 86.7}
	Organization	*Image size*	*0.0744*	*0.0379*	*96.13*
	Organization	*Image height*	*0.1811*	*0.106*	*70.83*
	Organization	*Image width*	*0.216*	*0.1414*	*52.72*
	Emotions	**Violence present**	**0.0257**	**0.0357**	**−27.9**
	Manipulation	**Artificial text**	**0.1179**	**0.3245**	**−63.68**
	Manipulation	**Likeliness of memes, text, or face editing**	**0.0612**	**0.1371**	**−55.34**

Note: Bold text indicates that the feature is associated with higher odds of fake news. Italic text is associated with higher odds of credible news. Some features are complex (e.g., three channels for red, green, blue colors), and all subfeatures are included in such cases.

of auxiliary verbs used were common in fake news articles. On the other hand, the number of words related to work and temporal past were more commonly used in credible news articles. In terms of organization, we found that credible news articles used fewer filler words and had more words per sentence. In terms of emotions, credible news articles made lesser use of words associated with swearing, anger, and sexuality. Finally, the personal pronoun "you" was much more frequently used in fake news articles.

5.2 | Visual features

In terms of content, fake news articles tended to have lower scores for all the three (red, green, blue) color channels, that is, they used darker images. This was true for the overall image, as well as the parts that were in focus. The topics or themes of the visual scenes captured in the images were also quite different between fake and credible news articles. In terms of organization, fake news used images that were smaller in height, width, and overall pixel size. Together with the color observation above, this suggests that they used less professionally captured photos. Images in fake news articles were also more likely to contain violence. Finally, fake news articles were more likely to use "manipulated" images in terms of the presence of artificial text and the likelihood of the image involving memes, text, or face editing. A more detailed discussion on these findings is presented in the Discussion section.

6 | CLASSIFIER: AUTOMATIC FAKE NEWS DETECTION

Next, we built an automatic fake news classifier using machine learning with the discussed features. We tried three modality sets: *text*, *visual*, and *text + visual*. As the data set included 3,568 fake news articles, to create a balanced data set, 3,568 credible news articles (out of 15,915) were randomly selected in each iteration to train and evaluate the classification algorithm. One hundred such iterations were used to reduce variance in the results, and a 70–30 train/test split was used to evaluate the classifier's performance.

Given the human effort required to identify fake news samples and the consequent modest sample size, we chose to adopt relatively simple classification models. We use Python Scikit-learn's implementation of multiple well-known algorithms such as Logistic Regression (Alpaydin, 2014), Linear Discrimination Analysis (Balakrishnama & Ganapathiraju, 1998), Quadratic Discriminant Analysis (Tharwat, 2016), K-Nearest Neighbors (Peterson, 2009), Naïve Bayes (Rish, 2001), Support Vector Machine (Noble, 2006), Classification and Regression

Tree (Alpaydin, 2014), and Random Forest (Pal, 2005; Pedregosa et al., 2011) analyses. For each classifier, we select the hyperparameters based on 10-fold cross-validation within the training set.

As the total number of features is relatively large (124 = 81 textual +43 Visual), we apply the Recursive Feature Elimination (RFE) technique to select a subset of features for all three modality sets in order to reduce training time, prevent overfitting, and improve the generalization of the models. RFE's goal is to find the threshold where the model does not lose in performance while keeping the number of features at its lowest. We used the SkLearn package called RFECV—Recursive Feature Elimination with Cross Validation—to identify the right number of features for the models. We used 10-fold cross-validation that yielded 20 features as the appropriate feature size.

The results for the accuracy obtained on the testing set (averaged over 100 iterations) using different algorithms and considering the three different models—text, visual, and both—are shown in Table 3. We notice a consistent trend of the multimodal, that is, text + visual approach, outperforming both the unimodal approaches.

In order to further analyze the performance of the classifiers, we consider multiple well-known metrics such as precision, recall (true positive rate), false negative rate, F1 score, and area under the receiver operating characteristic curve (AUROC) (Alpaydin, 2014; Zuva & Zuva, 2012). Precision can be defined as the probability given that an item is retrieved (here, identified as "fake news"), it will be relevant (i.e. indeed be "fake news"). Recall is the probability given that an item is relevant, it will retrieved (Zuva & Zuva, 2012). F1 score is the weighted harmonic mean of precision and recall. The Receiver Operating Characteristic (ROC) curve is a plot of true positive rate (TPR) versus false positive rate (FPR) across different thresholds. AUROC is a metric that defines how an algorithm performs over the ROC space. A score of 1.0 corresponds to a perfect classifier, and 0.5 corresponds to a completely random classifier (Alpaydin, 2014).

To optimize hyperparameters, we choose to optimize the F1 score as it is a function of both precision and recall. As Random Forest yielded the best performance in

TABLE 3 Results of each machine-learning model (accuracy) with different modalities

Testing accuracy—100 iterations			
Algorithm	Text	Visual	Text + visual
Logistic regression	81.93	74.69	84.43
Linear discriminant analysis	81.90	76.33	85.36
Quadratic discriminant analysis	77.68	70.29	81.43
K nearest neighbors	82.07	69.88	83.31
Naïve Bayes	77.71	64.86	78.46
Support vector machines	81.11	70.32	82.48
CART (classification and regression trees)	77.92	90.78	91.35
Random forest	86.24	93.84	95.18
AVERAGE (across classifiers)	80.82	76.37	85.25

TABLE 4 Results of best-performing machine-learning model with different modalities

Features	Confusion matrix [TP, TN, FP, FN]	Accuracy	Precision	Recall (true positive rate)	Miss rate (false negative rate)	F1 score	AUROC
Textual	[950, 896, 174, 120]	86.24	84.56	86.64	11.21	85.10	93.88
Visual	[1,022, 969, 101, 48]	93.84	91.84	92.54	4.49	92.19	98.14
Textual + visual	[1,042, 992, 78, 28]	95.18	93.18	94.88	2.62	94.02	98.86

the previous table, we zoom in and consider multiple metrics for Random Forest in Table 4. As can be seen from Table 4, the multimodal (textual + visual) approach again performs better than only textual and only visual models in terms of all the metrics considered. Note that a similar trend of multimodal approach outperforming only textual and only visual approaches was observed across the different algorithms (e.g., Logistic Regression, Linear Discriminant Analysis etc.) considered. Interested readers are pointed to

TABLE 5 Results of Wilcoxon signed rank test to compare the performance of different multimodal and unimodal approaches in terms of accuracy and area under the receiver operating characteristic curve (AUROC)

Comparison groups		p-value	Z- stat	Z- critical
Accuracy				
Text	Text + visual	8.88E−16	−8.68223	2.58
Visual	Text + visual	4.44E−16	−8.68239	2.58
AUROC				
Text	Text + visual	4.44E−16	−8.68177	2.58
Visual	Text + visual	1.11E−16	−8.68178	2.58

the detailed results in the supplementary information (Also available online at: https://docs.google.com/spreadsheets/d/1Z7BdeL57Gj2WOUHEvFWRqkyweiMo1WNHUWpu7EKL3do/edit?usp=sharing).

The results in Table 4 imply that textual and visual features alone are reasonably accurate in detecting fake news (accuracies of 86.24% and 93.84%, respectively). This would explain the reasonably large body of literature that has focused on a single modality. However, these modalities, if combined, can help the classifier make correct decisions in borderline cases. This is shown by the increase in accuracy from textual to textual + visual features. As we can see from Table 4, in relative terms, this results in a 10.37% increase in accuracy, a similar 10.48% increase in F1 score, and an increase from 93.88 to 98.86 in the AUROC score. Another way to think of this is that more than 60% of the errors remaining after textual analysis (or over 20% of the errors remaining after visual analysis) are removed with the use of multimodal analysis.

To test the significance of the improvements observed, we undertook a Wilcoxon Signed Rank Test (which does not need the normality assumption required for Paired Sample *t*-tests) between the multimodal approach and the

Maryland Trump Supporter: They Switched My Vote to Hillary Yet another report of vote flipping

A woman in Hollywood, Maryland is the latest in a number of early voters to claim that her ballot was switched to Hillary Clinton after she had tried to vote for Donald Trump. Maryland Trump Supporter: They Switched My Vote to Hillary "https://t.co/▓▓▓▓▓ pic.twitter.com/▓▓▓▓▓▓"

▓▓▓▓▓▓▓▓▓ (@▓▓▓▓▓▓) October 28, 2016 Noting that she had seen reports on the news of votes being flipped, the woman said, "I went in and voted a straight Republican ticket and thank God I went back and checked and they had switched my vote from Trump to (Hillary)." She said that she had to get the vote changed back by alerting election officials, who simply told her to vote for a second time.

...

FIGURE 2 A sample case that was detected correctly as fake news by the multimodal classifier but misclassified by the only textual and the only visual models [Color figure can be viewed at wileyonlinelibrary.com]

FIGURE 3 Use of artificial text in images used in fake news articles. Sources: https://barenakedislam.com/; http://img.youtube.com/vi/F5FYHE8VQ58/0.jpg [Color figure can be viewed at wileyonlinelibrary.com]

unimodal (only text and only visual) approaches. As shown in Table 5, the gains were found to be significant, with *p*-values consistently being below 0.001. This was true for comparisons based on accuracy and AUROC.

The results shown in Tables 3–5 demonstrate the proficiency of the multimodal approach.

To better understand the working of the multimodal approach, we analyzed a news story that was correctly detected as fake news by the multimodal classifier but misclassified by the unimodal classifiers (Figure 2).

In this case, the visual model found the story to be credible. Looking under the hood, we noticed there were eight features (e.g., size, height, width, darker colors) that were closer to average fake news scores and another 12 (e.g., sharpness of the facial image and the topics generated from image labels) that were closer to average credible news scores. In fact, we have no evidence to suggest that the image itself is untrue or manipulated. The text model also found the news story to be credible. Upon examination, 11 of the features (e.g., smaller word count, use of auxiliary verbs and informal language) were closer to fake news, but there were 9 features (e.g., higher words per sentence, use of temporal words, focus on present) that were closer to credible news. A ratio of 11:9 suggests that the algorithm was unable to make a clear choice, and ultimately, the algorithm selected the "credible news" label, which in this case was incorrect.

The multimodal model had the opportunity to utilize the most useful predictive features across both text and visual features. We found that 14 features in this case were closer to fake news (e.g., smaller word count, use of auxiliary verbs, informal language, and image height, width, size), and only 6 (e.g., higher words per sentence, focus on the present, words describing hearing/witnessing) were closer to credible news. Hence, what was a toss-up in the case of only visual features (8:12) and only textual features (11:9) became a clear case (14:6) in favor of fake news classification based on the multimodal analysis. Hence, multiple weak pieces of evidence across

modalities when combined together allowed for enough confidence for the classifier to identify the story as "fake news." We must note that this is only one of the many possible scenarios. However, this analysis, taken together with the empirical results presented in Tables 3–5, motivates the use of multimodal approaches for detecting fake news.

7 | DISCUSSION

RQ1 asks which visual and textual features are indicative of a news item being fake. The results obtained in Table 2 show multiple such features associated with fake news.

Several of these results, especially those pertaining to textual features, support results from previous work investigating fake news. For instance, the use of more personal pronouns and swear words by fake news stories have been reported earlier (Horne & Adali, 2017; Rubin et al., 2016). Fake news articles also differ from credible articles in language complexity and the way the news story is presented. For instance, fake news articles engaged in much higher use of auxiliary and filler words, which could indicate the presence of less original content with a high amount of redundancy. According to ELM, the central route, which requires assessment of information quality and source credibility, is the more effortful form of information processing. Hence, fake news articles might be targeting the peripheral route of information processing or the emotional pathway, as also suggested by LC4MP, to present simpler, emotion-laden content to the readers.

These findings are also in line with some of the recent theoretical and empirical work that connects LC4MP and ELM with misinformation. Work by Fisher et al. (2018) and Kirkwood and Minas (2020) suggests that an individual's a priori beliefs play an important role in how the information is processed. News stories that match with

one's prior beliefs tend to be analyzed via the peripheral route, while stories that do not match one's prior beliefs are analyzed via the central route. Hence, there is motivation for fake news creators to create news stories that are both emotionally arousing and consistent with their prior beliefs. While we do not know the a priori beliefs of potential readers in this work, it would make sense for fake news articles to target the peripheral route of information processing using the emotional pathway or simple redundant messaging.

In visual analysis, while previous research has presented an in-depth analysis of detecting evidence of manipulation and doctoring in images (Farid, 2006; Papadopoulou et al., 2017), we analyze features of images used in fake news articles. The images themselves may not be tampered with, but we did find some patterns in the images that fake news articles tend to use. For instance, darker colors have been known to elicit more negative connotations than lighter colors (Meier, Robinson, Crawford, & Ahlvers, 2007). Research has also shown this mental association to occur automatically rather than as a result of cognitive analysis; it has also been shown to be true across cultures (Meier et al., 2007). This implies that fake news stories consciously use images that trigger emotions of fear, anger, or suspicion in its users, further triggering an emotional response.

Apart from the use of color, the content of images in a fake news article was also more likely to have violent images, with a greater likelihood of containing blood and gore. Research studying violent imagery has found that people are often more attracted to violent imagery and experience an instinctive response (Goldstein, 1998). Hence, violent images may be used by articles trying to push their readers toward making associations that are unrelated to the logic and quality of information presented. Another interesting result is the high likelihood of the presence of artificial text in images used in fake news stories (see Figure 3 for some examples from the data set). This result again points to the tendency of fake news to target peripheral information processing.

Overall, the fake news strategy appears to be to present information that is often redundant but emotionally arousing. The focus is on catchy phrases, clichés, and slogans rather than on capturing specific factual details that support the truthfulness of the news being presented. This is consistent with the LC4MP model for real-time processing of mediated messages (Lang, 2009).

RQ2 investigates if a combination of image and text-based features improves fake news detection beyond that detected by unimodal (text only or image only) analysis. The results shown in Tables 3–5 demonstrate a significant jump in the performance of the classification algorithms in terms of multiple metrics such as accuracy, ROC area, and F-score when using both visual and textual features compared to only textual or only visual features. The results indicate that the textual and visual components work with each other to inform the readers' opinions, perceptions, knowledge, and beliefs about a particular topic. While textual and visual analysis certainly provide clues to identify fake news, the framing of images and text in connection with each other influences the audiences' response (Messaris & Abraham, 2001). Framing is found to activate specific thoughts and ideas for news audiences (Entman, 1993; Messaris & Abraham, 2001). In the context of fake news detection, this provides clear justification for analyzing visual and textual features as complementary signals to detect fake news.

It was interesting to note that, among the single-modality classifiers, the highest performance observed was for a visual classifier (not textual classifier). One way to interpret this is that, in a typical news article, an impactful image alongside text is known to evoke emotions and frame citizens' perceptions of politics (Grabe & Bucy, 2009). Images are also considered powerful framing tools because they are less intrusive than words and, as such, require less cognitive load. Hence, a classifier that can understand these aspects of visual communication can perform reasonably well at fake news detection (although not as well as a multimodal approach).

Finally, we note that the results obtained here are not intended to close the door on improving results further using more sophisticated text or visual analysis. Rather, these results are intended to motivate newer theory-driven research that uses *multimodal* analysis for fake news detection. The present study provides a starting point to identify other potentially relevant features for future research.

This study has some limitations. First, we would like to expand the data set. The current data set has a limited number of news articles, and the fake news articles include those that have been identified with variants such as "bs," "biased," "conspiracy" etc. based on crowdsourcing. In the future, we would like to replicate the key ideas of this work (i.e., multimodal analysis) using a bigger, more robustly labeled data set. Next, we would like to perform a user study to understand the cognitive mechanisms behind the acceptance and continued belief in fake news. For instance, recent work suggests that a priori beliefs of an individual significantly affect how that individual engages with a news report (Kirkwood & Minas, 2020). It is therefore

important to understand the cognitive processes guiding people toward clicking and sharing this news. Finally, this work focuses on full-length news articles. The dynamics of short-form news transmission via social media are important and different from what is considered here. Interested readers are pointed to other relevant efforts (e.g., Jin et al., 2017; Wang et al., 2018) that focus on those aspects.

Despite these limitations, this work has multiple implications for information science research. Building technologies that can automatically detect the veracity of a news item is an important research priority. Furthermore, beyond the creation of automated tools, the work undertaken here gives insights on the phenomenon of fake news in terms of its composition and the psychological underpinnings. In future, the multimodal text and visual features identified here could be used to create interactive dashboards that help users understand their own (fake) news consumption. For instance, users could interactively engage with news items, observe different textual and visual features, compare differences with typical fake and credible news articles, and then create an understanding of the fake news phenomenon.

8 | CONCLUSION

Considering the substantial effects of fake news on recent political events, the automatic detection of fake news has important practical consequences. This exploratory study investigates the efficiency of multimodal (text + visual) analysis in detecting fake news. The features identified based on multiple information processing and media theories are evaluated using several machine-learning algorithms and combined to create multimodal fake news detectors. The results identify textual and visual features that are more likely to be associated with fake news. They also suggest that multimodal analysis can help improve the performance of purely textual or purely visual fake news detectors. These results pave the way for better understanding of the fake news phenomenon, including its psychological underpinnings, and aid further research on multimodal fake news detection.

ORCID

Vivek K. Singh https://orcid.org/0000-0002-8194-2336
Isha Ghosh https://orcid.org/0000-0002-0879-3514

REFERENCES

Ahmed, H., Traore, I., & Saad, S. (2017). Detection of online fake news using N-gram analysis and machine learning techniques. In *International conference on intelligent, secure, and dependable systems in distributed and cloud environments* (pp. 127–138). Cham, Switzerland: Springer.

Allcott, H., & Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, *31*(2), 211–236.

Alpaydin, E. (2014). *Introduction to machine learning*. Cambridge, MA: MIT Press.

Balakrishnama, S., & Ganapathiraju, A. (1998). Linear discriminant analysis—A brief tutorial. *Institute for Signal and Information Processing*, *18*, 1–8.

Begley, J. (2017). The rise of the image: Every NY times front page since 1852 in under a minute. *Colossal*. Retrieved from https://www.thisiscolossal.com/2017/02/the-rise-of-the-image-every-ny-times-front-page-since-1852-in-under-a-minute/

Bessi, A., Coletto, M., Davidescu, G. A., Scala, A., Caldarelli, G., & Quattrociocchi, W. (2015). Science vs conspiracy: Collective narratives in the age of misinformation. *PLoS One*, *10*(2), e0118093.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, *3* (January), 993–1022.

Boididou, C., Andreadou, K., Papadopoulos, S., Dang-Nguyen, D. T., Boato, G., Riegler, M., & Kompatsiaris, Y. (2015, September). Verifying multimedia use at MediaEval 2015. *MediaEval*.

Burfoot, C., & Baldwin, T. (2009). Automatic satire detection: Are you having a laugh? In *Proceedings of the ACL-IJCNLP 2009 Conference Short Papers* (pp. 161–164). Association for Computational Linguistics.

Christensen, C. G., Bickham, D., Ross, C. S., & Rich, M. (2015). Multitasking with television among adolescents. *Journal of Broadcasting & Electronic Media*, *59*(1), 130–148.

Conroy, N. J., Rubin, V. L., & Chen, Y. (2015). Automatic deception detection: Methods for finding fake news. In *Proceedings of the 78th ASIS&T Annual Meeting: Information Science with Impact: Research in and for the Community* (p. 82). American Society for Information Science.

De Vreese, C. H. (2005). News framing: Theory and typology. *Information Design Journal & Document Design*, *13* (1), 51–62.

Dong, M., Yao, L., Wang, X., Benatallah, B., Sheng, Q. Z., & Huang, H. (2018). DUAL: A deep unified attention model with latent relation representations for fake news detection. In *International Conference on Web Information Systems Engineering* (pp. 199–209). Cham, Switzerland: Springer.

Entman, R. M. (1993). Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, *43*(4), 51–58.

Farid, H. (2006). Digital doctoring: How to tell the real from the fake. *Significance*, *3*(4), 162–166.

Feng, S., Banerjee, R., & Choi, Y. (2012). Syntactic stylometry for deception detection. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers*-Volume 2 (pp. 171–175). Association for Computational Linguistics.

Fisher, J. T., Huskey, R., Keene, J. R., & Weber, R. (2018). The limited capacity model of motivated mediated message processing: Looking to the future. *Annals of the International Communication Association*, *42*(4), 291–315.

Goldstein, J. (Ed.). (1998). *Why we watch: The attractions of violent entertainment*. Oxford, England: Oxford University Press.

Grabe, M. E., & Bucy, E. P. (2009). *Image bite politics: News and the visual framing of elections*. Oxford, England: Oxford University Press.

Gupta, A., Lamba, H., Kumaraguru, P., & Joshi, A. (2013). Faking sandy: Characterizing and identifying fake images on twitter during hurricane sandy. In *Proceedings of the 22nd International Conference on World Wide Web* (pp. 729–736). ACM.

Hong, S. C. (2013). Scare sells? A framing analysis of news coverage of recalled Chinese products. *Asian Journal of Communication*, *23*(1), 86–106.

Horne, B. D., & Adali, S. (2017). *This just in: Fake news packs a lot in title, uses simpler, repetitive content in text body, more similar to satire than real news*. Association for the Advancement of Artificial Intelligence.

Jin, Z., Cao, J., Zhang, Y., Zhou, J., & Tian, Q. (2017). Novel visual and statistical image features for microblogs news verification. *IEEE Transactions on Multimedia*, *19*(3), 598–608.

Kaggle. (2017). *Getting Real about Fake News: Text & metadata from fake & biased news sources around the web*. [Dataset]. Retrieved from https://www.kaggle.com/mrisdal/fake-news

Khaldarova, I., & Pantti, M. (2016). Fake news: The narrative battle over the Ukrainian conflict. *Journalism Practice*, *10*(7), 891–901.

Kirkwood, L., & Minas, R. (2020). Approaching fake news at the expense of truth: A psychophysiological study of news on social media. In *Proceedings of the 53rd Hawaii International Conference on System Sciences*.

Lang, A. (2009). The limited capacity model of motivated mediated message processing. In *The SAGE handbook of media processes and effects* (pp. 193–204). Newbury Park, CA: Sage.

Lin, Z., He, J., Tang, X., & Tang, C. K. (2009). Fast, automatic and fine-grained tampered JPEG image detection via DCT coefficient analysis. *Pattern Recognition*, *42*(11), 2492–2501.

Meier, B. P., Robinson, M. D., Crawford, L. E., & Ahlvers, W. J. (2007). When" light" and" dark" thoughts become light and dark responses: Affect biases brightness judgments. *Emotion*, *7*(2), 366.

Messaris, P., & Abraham, L. (2001). The role of images in framing news stories. In *Framing public life* (pp. 231–242). Abingdon-on-Thames, England, UK: Routledge.

Noble, W. S. (2006). What is a support vector machine? *Nature Biotechnology*, *24*(12), 1565.

Nørregaard, J., Horne, B. D., & Adalı, S. (2019). NELA-GT-2018: A large multi-Labelled news dataset for the study of misinformation in news articles. In *Proceedings of the International AAAI Conference on Web and Social Media* (Vol. 13, no. 01, pp. 630–638).

Osatuyi, B., & Hughes, J. (2018). A tale of two internet news platforms-real vs. fake: An elaboration likelihood model perspective. In *Proceedings of the 51st Hawaii International Conference on System Sciences*.

Pal, M. (2005). Random forest classifier for remote sensing classification. *International Journal of Remote Sensing*, *26*(1), 217–222.

Pan, J. Z., Pavlova, S., Li, C., Li, N., Li, Y., & Liu, J. (2018). Content based fake news detection using knowledge graphs. In *International semantic web conference* (pp. 669–683). Cham, Switzerland: Springer.

Pantti, M., & Sirén, S. (2015). The fragility of photo-truth: Verification of amateur images in Finnish newsrooms. *Digital Journalism*, *3*(4), 495–512.

Papadopoulou, O., Zampoglou, M., Papadopoulos, S., & Kompatsiaris, Y. (2017, June). Web video verification using con-textual cues. In *Proceedings of the 2nd International Workshop on Multimedia Forensics and Security*, ACM (pp. 6–10).

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Vanderplas, J. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12* (October), 2825–2830.

Pérez-Rosas, V., Kleinberg, B., Lefevre, A., & Mihalcea, R. (2017). Automatic detection of fake news. *arXiv preprint arXiv: 1708.07104*.

Peterson, L. E. (2009). K-nearest neighbor. *Scholarpedia*, *4*(2), 1883.

Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. In *Communication and persuasion* (pp. 1–24). New York, NY: Springer.

Rish, I. (2001, August). An empirical study of the naive Bayes classifier. In *IJCAI 2001 Workshop on Empirical Methods in Artificial Intelligence* (Vol. 3, No. 22, pp. 41–46).

Rubin, V., Conroy, N., Chen, Y., & Cornwell, S. (2016). 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).

Rubin, V. L., Chen, Y., & Conroy, N. J. (2015). Deception detection for news: Three types of fakes. In *Proceedings of the 78th ASIS&T Annual Meeting: Information Science with Impact: Research in and for the Community* (p. 83). American Society for Information Science.

Shu, K., Sliva, A., Wang, S., Tang, J., & Liu, H. (2017). Fake news detection on social media: A data mining perspective. *ACM SIGKDD Explorations Newsletter*, *19*(1), 22–36.

Shu, K., Wang, S., & Liu, H. (2019). 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.

Singhania, S., Fernandez, N., & Rao, S. (2017). 3han: A deep neural network for fake news detection. In *International Conference on Neural Information Processing* (pp. 572–581). Cham, Switzerland: Springer.

Starbird, K., Spiro, E., Edwards, I., Zhou, K., Maddock, J., & Narasimhan, S. (2016, May). Could this be true? I think so! Expressed uncertainty in online rumoring. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 360–371). ACM.

Tacchini, E., Ballarin, G., Della Vedova, M. L., Moret, S., & de Alfaro, L. (2017). Some like it hoax: Automated fake news detection in social networks. *arXiv preprint arXiv: 1704.07506*.

Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computer-ized text analysis methods. *Journal of Language and Social Psychology*, *29*(1), 24–54.

Tharwat, A. (2016). Linear vs. quadratic discriminant analysis classifier: A tutorial. *International Journal of Applied Pattern Recognition*, *3*(2), 145–180.

Wang, Y., Ma, F., Jin, Z., Yuan, Y., Xun, G., Jha, K., ... & Gao, J. (2018). Eann: Event adversarial neural networks for multimodal fake news detection. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 849–857). ACM.

Zuva, K., & Zuva, T. (2012). Evaluation of information retrieval systems. *International Journal of Computer Science & Information Technology*, *4*(3), 35.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

> **How to cite this article:** Singh VK, Ghosh I, Sonagara D. Detecting fake news stories via multimodal analysis. *J Assoc Inf Sci Technol*. 2021; 72:3–17. https://doi.org/10.1002/asi.24359