

Investigating the emotional appeal of fake news using artificial intelligence and human contributions

Abstract

Purpose:

The creation and dissemination of fake news can have severe consequences for a company's brand. Researchers, policymakers and practitioners are eagerly searching for solutions to get us out of the 'fake news crisis'. Here, one approach is to use automated tools, such as artificial intelligence (AI) algorithms, in conjunction with human inputs to identify fake news. The study in this article demonstrates how AI and machine learning, with its ability to analyze vast amounts of unstructured data, can help us tell apart fake and real news content. Specifically, this study examines if and how the emotional appeal, i.e., sentiment valence and strength of specific emotions, in fake news content differs from that in real news content. This is important to understand, as messages with a strong emotional appeal can influence how content is consumed, processed and shared by consumers.

Design/methodology/approach:

The study analyzes a data set of 150 real and fake news articles using an AI application, to test for differences in the emotional appeal in the titles and the text body between fake news and real news content.

Findings:

The results suggest that titles are a strong differentiator on emotions between fake and real news and that fake news titles are substantially more negative than real news titles. In addition, the results reveal that the text body of fake news is substantially higher in displaying specific negative emotions, such as disgust and anger, and lower in displaying positive emotions, such as joy.

Originality/value:

This is the first empirical study that examines the emotional appeal of fake and real news content with respect to the prevalence and strength of specific emotion dimensions, thus adding to the literature on fake news identification and marketing communications. In addition, this paper provides marketing communications professionals with a practical approach to identify fake news using AI.

Keywords:

Brand communication, Message framing, Machine learning, Emotional appeal, Natural language processing, Emotional branding, Communication model, Real news, Fake news, Artificial intelligence (AI)

Paper type:

Research paper

Fake news and brand communications

“When a headline asks a question, the answer should be no”. Although this may read like a bizarre or even humorous statement, this adage even has a name: Betteridge’s law (Betteridge, 2009). Formulated by the British journalist Ian Betteridge, it is based on the notion that journalists use this style of headline if “they know the story is probably bullshit, and don’t actually have the sources and facts to back it up, but still want to run it” (Betteridge, 2009). This quote certainly rings true today. For example, a 2017 article from YourNewsWire claimed that according to an NPR study, more than 25 million Hillary Clinton votes were fraudulent, suggesting that Clinton had actually lost the popular vote by a huge margin. This claim was false and the study in question was never conducted by NPR; still, it was among the most viral stories shared on Facebook in 2017. Although not a new phenomenon, the generation and impact of fake news and alt- facts have reached new heights, driven mainly by the increasing digitization of information and the explosion of social media (Baccarella *et al.*, 2018; Berthon and Pitt, 2018).

Fake news, defined as news that intentionally present misinformation with the intent to deceive the audience (Bakir and McStay, 2018; Horne and Adali, 2017; Kumar and Shah, 2018), also known as disinformation (Hannah *et al.*, 2015), can have severe consequences for brands, businesses and societies as a whole. A prime example is the 2016 US presidential election and the concern about how false stories on social media may have impacted the election outcome (Allcott and Gentzkow, 2017). This has led a number of commentators from reputable news outlets to suggest that Trump would not have been elected without the influence of fake news (Dewey, 2016; Olson, 2016; Parkinson, 2016; Read, 2016). As a well- informed public is key to any effective democracy, fake news can be especially dangerous in the context of public opinion and political information.

In addition, brands can also be impacted by and impact fake news in a number of ways (Berthon and Pitt, 2018). First, brands can be the direct target of false stories with major consequences: Pepsi’s stock declined 4 per cent when a fake news story about its CEO telling Trump supporters to “take their business elsewhere” went viral. Second, brands can affect fake news stories by consciously or unintentionally associating themselves with dubious content. When this happens, this may be

perceived as brands accepting or even endorsing the false stories, which, on the one hand, may render legitimacy to the false information and, on the other hand, contaminate a brand's image (Berthon and Pitt, 2018). As an illustration, Kellogg's received substantial backlash from consumers for advertising on the alt-fact site Breitbart and eventually removed its ads (Helmore, 2016).

As the above examples illustrate, fake news can threaten the viability of a brand, which is often regarded as a firm's most valuable intangible asset (Aaker, 1991; Keller, 1993). Marketing communication managers are well advised to be vigilant about if and how their brands are associated with fake news. With the rising use of social media, false content can originate from and be shared by any of a brand's stakeholder group – employees, customers, users or third-party news outlets, among others.

Hence, researchers, policymakers and practitioners are eagerly searching for solutions to get us out of the “fake news crisis” (Rothman, 2018). Systemic solutions that involve rethinking the role of marketing communications and branding in creating reality are called for (Berthon and Pitt, 2018); however, these systemic changes will need time.

In the short term, managers require tactical solutions that allow them to identify fake news when designing and executing brand communication strategies. These approaches will need to allow practitioners to flag dubious news at any one of the three core elements of the marketing communication process: the source, the receiver and the message itself. Here, one tactic currently explored is to verify the source of the information, for example, through publishing lists of “most trusted” and “least trusted” news sources (e.g., Strike Social). A problem with this approach is that these lists are static and may be outdated by the time they are published.

Another approach has been to limit the sharing of false content on social media sites like Facebook or Twitter (Baly *et al.*, 2018), which is challenging because fake news generally thrive on social media thanks to the mechanisms of sharing, which amplifies their dissemination. A third strategy is to analyze the news content itself and flag dubious content. For example, fact-checking organizations such as Snopes, FactCheck or PolitiFact manually verify that the claims made in articles are actually factual; however, this is a slow and tedious process that does little to identify false information proactively. On the other

hand, entirely automated fact-checking tools lag behind in terms of availability of authoritative data and are less trusted by humans (Baly *et al.*, 2018).

One approach that has gained more attention recently is to use automated tools, such as artificial intelligence (AI) algorithms, in conjunction with manual fake news identification approaches. At the time of writing this article, social media giant Facebook, for instance, had just announced to expand its team of content specialists, in addition to relying on algorithmic approaches to fake news screening (Hern, 2018; Wakefield, 2019). The premise is that automated tools, such as AI algorithms, can effectively assist humans in the important task of identifying fake news (Baly *et al.*, 2018; Graves, 2018; Hao, 2018; Berthon and Pitt, 2018; Strickland, 2018).

Building on this anticipation, in this article, I demonstrate how AI in conjunction with human inputs can help us tell apart fake and real news content. Although there is a large body of research examining the factors that contribute to the dissemination of fake news, there is relatively little work that helps us understand how the content of fake and that of real news differ (Baly *et al.*, 2018; Horne and Adali, 2017). In other words, of the three strategies mentioned above (i.e., to verify the source, analyze the content and limit the sharing), I focus on the content. This study contributes to our understanding of news content by investigating if, and how, fake and real news vary with respect to their emotional appeal. More specifically, this study focuses on understanding whether the overall sentiment valence and strength of specific emotions in fake news articles differ from those in real news articles. This is important to understand, as messages with a strong emotional appeal can impact how content is consumed, processed and shared (Meenaghan, 1995; Taute *et al.*, 2011).

With these aims in mind, this article proceeds as follows: It first provides conceptual foundations of fake news before discussing the current knowledge of fake news identification as it relates to the key elements in the marketing communication process (i.e. the source, the content and the receiver). The next section (Section 4) then describes the empirical study underpinning this paper and presents the study results. Section 5 discusses these results in the context of extant brand communications literature, and the paper concludes with contributions for practice and academia as well as avenues for future research.

Conceptualizing fake news: factuality and intent to deceive

Fake news has become a buzzword, especially in the aftermath of the 2016 US presidential election, so much so that the Collins Dictionary made it the “word of the year” in 2017 (Habgood-Coote, 2018). Early studies on the topic of fake news have defined the phrase in terms of particular types of media content, such as political satire, news parodies or news propaganda (Tandoc *et al.*, 2018). Contemporary discourse defines fake news as viral posts, often on social media, based on fictitious accounts constructed to look like news reports. BuzzFeed’s media editor Craig Silvermann is credited with helping to popularize the term in 2016, when he identified 140 Macedonian websites that fabricated stories to attract audience interest and profitable clicks. One key element of fake news, according to Silvermann’s definition, is an underlying profit motive of its creators (Habgood-Coote, 2018). A second key motivation implied in many definitions of fake news is ideological: Fake news providers create fake stories to promote particular ideas or people that they favor, often by discrediting others (Allcott and Gentzkow, 2017). A similar ideological motif is emphasized by the Ethical Journalism Network (2017), which defines fake news as “information that has been deliberately fabricated and disseminated with the intention to deceive and mislead others into believing falsehoods or doubting verifiable facts” (Ethical Journalism Network, n.d.). These illustrations show that different people have a different understanding of what fake news means. However, for parsimony, it is important to clearly conceptualize fake news in this article. This is the purpose of the following section.

Originally derived as a plural of “new”, news have been defined in various ways, such as the account of recent and interesting events (Kershner, 2005), the presentation of new information (Stephens, 1988) or the account of events that have a significant impact on people (Richardson, 2007). News is often viewed as an output of journalism and is expected to provide independent, reliable, accurate and comprehensive information (Tandoc *et al.*, 2018). Kovach and Rosenstiel (2014) suggest that the key purpose of journalism is to “provide citizens with the information they need to be free and self-governing” (Kovach and Rosenstiel, 2014, p. 12); thus, journalism has

a responsibility to report accurate and true information in the output it produces.

These definitions of fake news also prompt us to think about the connotation of “fake”. The Oxford Dictionary defines fake as “not genuine; imitation or counterfeit” (Oxford English Dictionary, 2018). The word “fake” is often used interchangeably with “not real”, “false”, “fraudulent” or “inauthentic” (Andrea, 2016; Oxford English Dictionary, 2018). Thus, according to this notion, one could judge whether a news item is fake or real based on the degree to which it relies on facts. However, one problem with this unidimensional conceptualization is that it would consider certain news types, such as satire or parody, as fake news; however, these forms play an increasingly important part of the media ecosystem and a critical and informed public (Plangger *et al.*, 2019; Rubin *et al.*, 2016; Tandoc *et al.*, 2018). Satire, for example, relies on actual events but uses humor, sarcasm and exaggeration, rather than objective facts, to portray these events. Satire news outlets promote themselves and their stories as delivering entertainment to their audiences, rather than independent, accurate and reliable information (Rubin *et al.*, 2016; Tandoc *et al.*, 2018).

In a meta-analysis of scholarly articles published on fake news between 2003 and 2017, Tandoc *et al.* (2018) introduce a typology of fake news based on two dimensions, similar to Hannah *et al.*'s (2015) typology of leaked secrets. Each dimension in Tandoc *et al.*'s (2018) typology encompasses a spectrum, ranging from low to high. The first dimension is facticity, i.e. the degree to which fake news relies on facts. High-facticity news items provide accurate and reliable real accounts of events, whereas low-facticity stories often take a broad social context upon which they fabricate fictitious accounts. The second dimension, author intent to deceive, refers to the degree by which the creator of a piece of content intends to misinform people for ideological or profit motivations. Satire news have a low intent to deceive – their main purpose is to entertain their audiences, and they assume an open disclaimer that their stories are not factual (Tandoc *et al.*, 2018). On the other hand, creators of fabricated news content intent to mislead their readers, without any disclaimer.

Recent scholarly work has emphasized the two dimensions proposed by Tandoc *et al.* (2018) to define fake news (Bakir and McStay, 2018; Horne and Adali, 2017; Kumar and Shah, 2018). Following these conceptualizations, the present study defines real news as news articles with a high degree of facticity and low (no) intent to deceive. Fabricated

(or fake) news have a low degree of facticity and a high intention to mislead their audience. Table 1 presents a typology of news definitions; highlighted in italics are the types that this paper focuses on.

Table 1 Typology of news definitions (adapted from Tandoc *et al.* (2018))

Degree of facticity	Author's intention to deceive	
	High	Low
High	Propaganda news	<i>Real news</i>
Low	<i>Fake news</i>	Satire news

Having delineated the key constructs, the following section reviews the existing literature on fake news identification and positions the current study within this literature.

Fake news identification: source, receiver and content

Fake news detection has been extensively examined in the field of communications and computer science. Working from the three key elements of (brand) communication mentioned earlier, i.e. the source, the receiver and the message (i.e. content), the fake news literature can be categorized into three different paradigms: source-based, context-based and content-based studies (Table 2; Potthast *et al.*, 2017).

Table 2 Elements of communication and corresponding paradigm for fake news detection literature alongside a selection of relevant work

Key element of communication	Source	Receiver	Message	
Paradigm of fake news detection	Source-based	Context-based	Content-based	
			Fact-checking	Style-based
	Baly <i>et al.</i> (2018) Mukherjee and Weikum (2015) Popat <i>et al.</i> (2016)	Kwon <i>et al.</i> (2013) Bessi <i>et al.</i> (2015) Del Vicario <i>et al.</i> (2016) Vosoughi <i>et al.</i> (2018)	Etzioni <i>et al.</i> (2008) Magdy and Wanas (2010)	Chen <i>et al.</i> (2015) Potthast <i>et al.</i> (2017) Home and Adali (2017) Home <i>et al.</i> (2018)

Source-based fake news detection studies attempt to identify the degree of faking by assessing the reliability of the source reporting the information. For example, Baly *et al.* (2018) predict the factuality of reporting and bias of various news sources using a machine learning tool. Mukherjee and Weikum (2015) present a model to predict the trustworthiness of news communities, such as Digg and Reddit, and how trustworthiness interacts with content credibility and user expertise. Popat *et al.* (2016) propose a machine learning tool that leverages the interaction between the language of articles and the reliability of the underlying Web sources to identify and predict false claims.

Context-based fake news detection is mainly concerned with understanding the dissemination of fake news online, thus focusing on the receiver of a message in the communication model. For instance, Kwon *et al.* (2013) propose a model to characterize rumors by examining the temporal, structural and linguistic aspects of their diffusion on social media. Del Vicario *et al.* (2016) conduct extensive quantitative modeling to understand the mechanisms behind the spreading of misinformation on Facebook. In a recent Science publication, Vosoughi *et al.* (2018) investigate the diffusion of 126,000 verified true and verified false news stories distributed on Twitter from 2006 to 2017. They found that false news stories diffused farther and faster, were more novel and inspired replies with different emotions (such as surprise and disgust) compared to true news stories (such as sadness, trust and anticipation).

Content-based fake news studies analyze the textual or visual content to identify dubious content, thus focusing on the third key element in the communications process, i.e. the message. Here, two subfields exist: *knowledge-based* (or “fact checking”) studies use methods borrowed from information retrieval research to extract facts from a given narrative and compare them with known truths. For instance, Etzioni *et al.* (2008) use a tool to extract factual claims made in a given narrative, compare them with facts retrieved from the Web and flag discrepancies. In another study, Magdy and Wanas (2010) demonstrate a statistical model that checks factual claims from a given text and determines how frequently they are supported by facts retrieved from the Web. The underlying assumption of knowledge-based approaches is that Web resources or the frequency by which a fact is mentioned is an indicator of its truth; however, this is problematic and the reliability and credibility of almost any website can be questioned (Potthast *et al.*, 2017). A number of fact-

checking organizations rely on the knowledge-based approach to fake news identification.

The second subfield to content studies, *style-based* fake news detection, focuses on modeling the degree of faking and its manifestation in a given text assuming that deception has its own style. Here, extant studies have examined the textual features, such as semantics and syntax. For example, Potthast *et al.* (2017) use 1,627 articles that have been manually fact-checked by professional journalists from BuzzFeed and apply stylometric analysis using a machine learning algorithm to predict factuality of stories. In another study, Chen *et al.* (2015) examine methods for the automatic detection of clickbait content as a form of deception, by analyzing both textual (i.e., semantics and syntax) and non-textual (i.e., images) elements of content. They conclude that clickbaiting can be identified through “certain linguistic patterns, such as the use of suspenseful language, unresolved pronouns, a reversal narrative style, forward referencing, image placement, reader’s behavior and other important cues” (Chen *et al.*, 2015, p. 18). Featuring prominently among these “other important cues” are emotional features derived from their text analysis, e.g., via sentiment analysis; however, this is not further explored in their study. Horne and Adali (2017) analyze a comprehensive set of features including stylistic, complexity and psychological characteristics in a data set of real, fake and satire news. The authors conclude that fake news use shorter, simpler and more repetitive content. They also found that fake news content displays stronger negative emotions compared to real news; however, the type of emotions in fake news is not further explored. Horne and Adali (2017) also found a strong dissimilarity between the title and text body of news and suggest that the title and the body should be analyzed separately. In a follow-up study, Horne *et al.* (2018) use a data set encompassing over 136,000 news articles from 92 sources, which they characterize based on 130 features from seven categories: structural, sentiment, engagement, topic- dependent, complexity, bias and morality. Their analysis suggests that highly engaged fake articles express more negative emotions, thus confirming previous findings by Bessi *et al.* (2015) and Horne and Adali (2017).

Although a number of studies under the content-based paradigm exist, this field of research is still underdeveloped (Potthast *et al.*, 2017). In particular, the above literature discussion suggests one

interesting finding common across a number of studies under this paradigm: That the content of real and fake news is different with respect to emotions portrayed in it. However, what is currently less understood is the valence of emotions and how different types of emotions differ across different types of narratives. The study presented in this article helps close this gap in our understanding. Specifically, the research uses AI, in addition to human analysis, to determine whether the valence and the strength of specific emotions differ in fabricated and real news articles. The next section will outline the methodology used in this study to achieve this research objective.

A study about the emotions in fake news and real news

The research conducted here was guided by the following research questions:

RQ1. How do the overall sentiment valence and specific emotions displayed in the titles differ with respect to real and fabricated news articles?

RQ2. How do the overall sentiment valence and specific emotions displayed in the article body differ with respect to real and fabricated news articles?

To answer these research questions, I use a data set encompassing political news articles for each of the above two types of news: real and fabricated. This data set was used in a previous study by Horne and Adali (2017) and data collection followed a two-step process. The first step, identifying the appropriate source for each of the two news categories, relied primarily on manual input. Specifically, this step relied on published lists of news outlets that had been previously deemed “trustworthy” or “misleading” by human fact checkers. For instance, the sample source for real news articles encompassed Business Insider’s list of “most trusted and least trusted” news outlets (Engel, 2014) and includes well-established news outlets that all have low (no) intent to

mislead their audience. Fabricated news articles were sourced from Zimdars' (2016) list of misleading news sources which have had at least one news article identified as false on the fact-checking website snopes.com. Table 3 summarizes the sources for each of the three types of news articles in this study.

Table 3 News article sources (adapted from Horne and Adali (2017))

Real news source	Fabricated news sources
ABC	Abcnewsc.com.co
BBC	Before it's news
CBS	DC Gazette
NBC	Ending the Fed
NPR	Infowars
The Economist	Libertywritersnews
The Guardian	Real News Right Now
USA Today	True Pundit
Wall Street Journal	
Washington Post	

The second step involved randomly selecting 75 political news articles from each source group – real news and fabricated news, for a total of 150 articles. Only “hard” news articles were included in the data set, whereas editorials or opinion pieces were excluded. Horne and Adali (2017) made the above- described data set publicly available on the open-source platform github (<https://github.com/rpitrust/fakenewsdata1>), so this platform was accessed to obtain the data set for the current study. Although Horne and Adali (2017) analyzed a comprehensive set of linguistic and psychometric features by which fake and real news content differed, their study did not consider if and how specific types of emotions differ in fake and real news content. This is the focus of my study, for which the results are outlined in the following section.

Analyzing emotions and sentiment using AI

The unit of analysis in this study encompasses the written text in the body and titles of news articles in the data set described above. Following the suggestion by Horne and Adali (2017), the present study will analyze the

text body and titles of real news and fake news separately. The first step in the data analysis included analyzing the data set using the AI application IBM Watson. AI systems, such as IBM Watson, use natural language understanding to assign meaning to text-written everyday human language and to detect language patterns, topicality, sentiment and other linguistic characteristics (Kietzmann *et al.*, 2018; Rubin *et al.*, 2016). In addition, machine learning algorithms enable an AI system to learn by itself, i.e., to improve its performance without the algorithm being changed by a trained coder. Among Watson's suite of applications is a "Natural Language Understanding Service", which uses linguistic analytics to "analyze target phrases in context of the surrounding text for focused sentiment and emotion results" (IBM, 2018).

AI applications, such as IBM Watson, have been used in other academic research, for example, to extract personality profiles from written text (Majumder *et al.*, 2017; Mostafa *et al.*, 2016) or to understand the sentiment portrayed in a body of text (Cambria *et al.*, 2013; Pitt *et al.*, 2018). AI algorithms have also been extensively used to perform linguistic analyses in the context of fake news detection (Chen *et al.*, 2015; Horne and Adali, 2017; Potthast *et al.*, 2017; Rubin *et al.*, 2016). As part of its Natural Language Understanding Service, IBM Watson calculates scores for emotions detected in a narrative based on Plutchik's (2001) work on basic human emotions. The five emotions analyzed by Watson include sadness, joy, fear, disgust and anger, and for each emotion, the score ranges from 0 to 1: A score of 0 means the text does not convey the emotion, whereas a score of 1 indicates that the text expresses the emotion strongly (IBM, 2018). In addition, Watson returns an overall sentiment score indicating the valence (positive/negative/neutral) and strength of sentiment displayed in the focal document. Sentiment scores range from -1 (strong negative sentiment) to 1 (strong positive sentiment); a score of 0 suggests a neutral sentiment. Following the Watson analysis, the second step in the data analysis included further analyzing the emotion and sentiment scores using the content analysis software DICTION (Pitt *et al.*, 2017; Short *et al.*, 2018) for descriptive statistics and the statistical software package SPSS 25.

Table 4 displays the summary statistics for article body and titles. An observation to be made from Table 4 is that the body of real news articles appears to be much longer and uses longer words on average than fabricated news articles. However, the titles of real news appear to be shorter but use longer words on average than fabricated news.

Table 4 Descriptive statistics

	Text body		Title	
	Real news (n=75)	Fabricated news (n=75)	Real news (n=75)	Fabricated news (n=75)
Shortest text (words)	126	100	6	6
Longest text (words)	2316	1700	16	21
Mean total words	699.4	477.1	9.9	11.8
Total words	52458	35781	739	887
Mean word size (characters)	5.06	4.85	5.65	5.45

Sentiment valence by news type

For each news type, the proportions of narratives with positive, negative and neutral sentiment were computed; this was done separately for the article body and the title. Inferential statistics comparing two population proportions (*Z* scores; Hair *et al.*, 2011) were calculated to test whether the proportion of narratives displaying an overall positive (negative, neutral) sentiment differed across each news category. The level of significance was set at $\alpha = 0.05$.

The results in Table 5 suggest that there is no significant difference in the sentiment valence of the article body of fabricated versus real news articles. However, the article titles for each news category differed with respect to sentiment valence (Table 5). Specifically, a larger proportion of fabricated news display a negative sentiment in their titles, and proportionally fewer fabricated articles reveal a neutral sentiment.

Table 5 Sentiment valence in the text body and titles

	Text body			Title		
	Real news (n=75) Proportion	Fabricated news (n=75) Proportion	Z	Real news (n=75) Proportion	Fabricated news (n=75) Proportion	Z
Positive	0.080	0.160	-1.508	0.040	0.053	-0.387
Negative	0.920	0.840	1.003	0.240	0.600	4.467*
Neutral	0.000	0.000	-1.508	0.720	0.347	4.583*
Total	1.000	1.000			1.000	

* $p < 0.05$

Specific emotions by news type

An independent-samples t-test was used to test the hypothesis that the two types of news differ on specific emotion dimensions. The results (Table 6) suggest that emotions in the article body itself differ significantly for three out of the five emotion dimensions analyzed. Fake news articles were significantly less joyful than real news articles ($t_{(144.434)} = 2.073$, $p = 0.040$). In contrast, fabricated news articles expressed significantly higher levels of disgust ($t_{(147)} = -3.720$, $p = 0.000$) and higher levels of anger ($t_{(136.408)} = -3.433$, $p = 0.000$) than real news articles. In other words, the text body of fake news articles was less joyful but angrier and more revulsive than the text body of real news articles. There were no significant differences in the article body of fabricated and real news with respect to the dimensions of sadness ($t_{(147)} = -1.408$, $p = 0.161$) and fear ($t_{(140.389)} = -1.409$, $p = 0.161$).

With respect to the article titles, the results in Table 7 suggest that fabricated news articles displayed significantly higher levels of disgust ($t_{(131.647)} = -3.253$, $p = 0.001$) and higher levels of anger ($t_{(147)} = -3.418$, $p = 0.000$) than real news articles. No significant differences were found in the titles for the emotion dimensions of joy ($t_{(130.789)} = -0.913$, $p = 0.363$), sadness ($t_{(147)} = 0.737$, $p = 0.462$) and fear ($t_{(147)} = 0.653$, $p = 0.515$).

Table 6 Emotion dimensions in the text body

	Real news text body (n=75)		Fabricated news text body (n=74)		Independent-samples t-test			
	Mean	SD	Mean	SD	t	df	p (2-tailed)	Cohen's d
Sadness	0.417	0.154	0.451	0.146	-1.408	147	0.16	n/s
Joy _a	0.43	0.178	0.365	0.201	2.073	144.434	0.040*	0.345
Fear _a	0.166	0.151	0.205	0.186	-1.409	140.389	0.16	n/s
Disgust	0.302	0.176	0.411	0.183	-3.720	147	0.000***	-0.614
Anger _a	0.195	0.131	0.281	0.172	-3.433	136.408	0.001***	-0.588

Table 7 Emotion dimensions in the article title

	Real news text body (n=74)		Fabricated news text body (n=74)		Independent-samples t-test			
	Mean	SD	Mean	SD	t	df	p (2-tailed)	Cohen's d
Sadness	0.256	0.155	0.238	0.140	0.737	147	0.462	n/s
Joy _a	0.159	0.144	0.186	0.211	-0.913	130.789	0.363	n/s
Fear	0.147	0.170	0.131	0.127	0.653	147	0.515	n/s
Disgust _a	0.216	0.149	0.315	0.216	-3.253	131.647	0.001*	-0.567
Anger _a	0.178	0.135	0.255	0.1391	-3.418	147	<0.001***	-0.563

^aEqual variances not assumed as Levene's test is significant at the p<0.05 level; *p < 0.05; ***p < 0.0001

Discussion

If, as alluded to earlier, fake news is about stories with low degree of facticity and high intent to deceive for ideological or profitable motivations, then one would expect the content of fake and real news to substantially differ. Previous studies have confirmed that substantial differences exist and have suggested that the emotional appeal is one important feature by which to tell fake and real news apart (Bessi *et al.*, 2015; Chen *et al.*, 2015; Horne and Adali, 2017). However, what was lacking in the extant literature was an understanding of if and how specific emotions differ between fake and real news. This is the research objective of the present study. Specifically, this work was guided by the following research questions: How do overall sentiment valence and specific emotions displayed in titles differ with respect to real and fabricated news articles? And, how do the overall sentiment valence and specific emotions displayed in the article body differ with respect to real and fabricated news articles? The following section discusses the findings related to these research questions.

Titles are a strong differentiator on emotions between fake and real news

When looking at only the titles of fake and real news, the results suggest substantial dissimilarities between the two types of news. First, fake news titles are significantly more “emotional” and less neutral in sentiment than titles of real news articles, thus confirming previous findings by Horne and Adali (2017) and Horne *et al.* (2018). Table 8 provides an illustrative example of a fabricated news title displaying a high (negative) emotional tone in comparison to a real news title with neutral sentiment.

Table 8 Example of fabricated news title and strong sentiment valence

Fabricated news title – Negative sentiment	Real news title – Neutral sentiment
BREAKING! Dems Caught Counting Hillary Votes 6 TIMES In THIS City, TIME FOR PRISON!	Report: White House preparing response to Russian election interference

To better explain these findings, I turn to the existing literature on message framing in marketing and brand communication. Marketing communication strategies can be roughly divided into rational versus emotional appeals (Dens and Pelsmacker, 2010). A rational message framing aims to inform consumers using factual claims, whereas an emotional framing aims to evoke emotions in consumers (Johar and Sirgy, 1991). An emotional appeal in a message facilitates message processing in a number of ways, for example, through containing less information to be processed by consumers, improved levels of attention (Heath, 2007) and higher levels of recall (Ambler and Burne, 1999; Friestad and Thorson, 1986). Given consumers have limited capacity to absorb all the information contained in a message (Lang, 2000), an emotional appeal may be more effective at persuading consumers than a rational message frame. Furthermore, studies suggest that a majority of links shared on social media are never clicked, only the titles are read by users (Wang *et al.*, 2016). This suggests that titles may serve as an important mechanism for fake news to audience persuasion through strong emotional appeals.

Fake news content is substantially more “disgustful” and “angry” and less “joyful”

The results of my study also revealed that negative emotions are much more prominent in fabricated news than real news, in both article title and text body. Specifically, fabricated news articles display significantly higher levels of anger and disgust and substantially lower levels of “joy” in their article body than real news stories. Again, to make sense of these findings, one can turn to the communications literature. Although research into the impact of specific emotions, such as disgust, is rare, a number of studies have considered negative emotions and their impact on message processing (Dens *et al.*, 2008). Extant studies suggest that negative message framing, such as offensive or shocking content, can be effective at cutting through the “information clutter”, improve attention and facilitate message retrieval and recall (Cotte and Ritchie, 2005; Dahl *et al.*, 2003). This can be explained using Mandler’s (1982) schema incongruity framework according to which consumers process messages on the basis of existing schemata, i.e. cognitive structures representing consumers’ knowledge and expectations about brands, products and

events (Dens *et al.*, 2008). According to Mandler, a message that is incongruent with existing structures will attract more attention, more extensive processing and better recall than information consistent with existing schemata (Mandler, 1982). This suggests that negative message framing may be a means for fabricated news to draw an audience's attention and facilitate recall.

Contributions to literature, managerial implications and future research

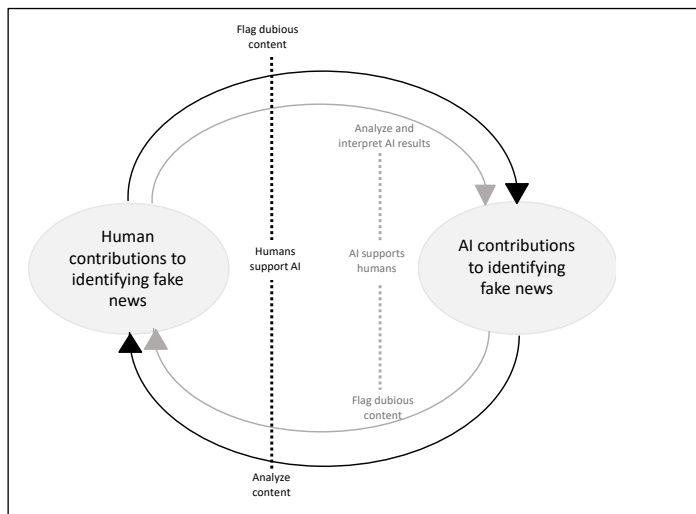
This study makes a number of contributions to academia and practice. First, it contributes to the body of knowledge on the distinguishing characteristics of fake versus real news content. Previous studies have suggested that fake news and real news are different with respect to linguistic style and particularly the emotional tone. The present study confirms previous findings that fake news are substantially more negative than real news. In addition, and this is a new contribution, the study in this article has demonstrated which specific emotions differ in fake versus real news narratives. Specifically, the study found that fake news is substantially higher in anger and disgust and lower in joy than real news articles. This opens up a number of avenues for future research (discussed in the section below) and contributes to our understanding of message framing as a key element of any communication process.

In addition, the study offers a methodological contribution by demonstrating how computer algorithms, specifically AI tools with their ability to process natural language and “learn by themselves” (Kietzmann *et al.*, 2018), can assist researchers in their data analysis to detect patterns in large sets of unstructured qualitative data. Without an automated tool, such as IBM's Watson, the sheer task of manually coding a data set consisting of more than 85,000 words seems daunting, if not unachievable. Here, the computer algorithms made data analysis substantially less time-consuming. However, this does not suggest that human contributions are eliminated from the research process and the data analysis. On the contrary, the numerical output from IBM Watson has lent itself well to further quantitative analysis by researchers, for

example for hypothesis testing, as demonstrated in this article. These quantitative analyses could be replicated by other scholars.

Furthermore, a number of managerial implications flow from the findings of this research. First, this study provides managers with a practical approach by which to identify fake news as part of their brand communication plans and implementations. Specifically, in this study, I have demonstrated a tactic that practitioners can use to identify dubious content originating from third-party news outlets. This is a complex undertaking given the volume and velocity of communications that are produced and consumed every day. Here, AI tools can provide valuable support for managers by doing the “heavy lifting” on the data analysis. Trained observers can then further analyze, verify and interpret the results produced by AI tools. Alternatively, or in addition to as shown in this study, humans can be enlisted to flag spurious content or their sources to be used as an input in the AI analysis. This is consistent with Berthon and Pitt’s (2018) call to “introduce a human layer” (Berthon and Pitt, 2018, p. 222) into solutions to deal with the fake news crisis. This suggests that both human and machines offer valuable contributions to identifying fake news. Figure 1 summarizes the relationship between human and machine contributions to fake news identification.

Figure 1 Relationship between human and artificial intelligence (AI) contributions to fake news identification



Second, this article identifies specific emotional dimensions by which to tell fake and real news apart. This is important for brand managers because a brand's association, willingly or unwillingly, with fake news can significantly contaminate a brand's image with substantial consequences, as illustrated at the beginning of this article. Thus, practitioners can use the insights from this study to minimize placing their brand ads next to fake news and thus better "manage their brands in this post-rational world" (Berthon and Pitt, 2018, p. 222).

Third, the results of this study may also serve as a note of caution for communication professionals or writers, be it the creators of fake content or non-fake ad copy. In an attempt to cut through the clutter and draw consumers' attention, well-known brands and organizations have used highly controversial and negative message appeals relying on anger or disgust (Dens *et al.*, 2008). The European clothing manufacturer Benetton, for example, created a brand campaign including advertisements displaying a photograph of the blood-stained clothes of a dead Croatian soldier or a man dying of AIDS (Crawshaw, 1995). Although this shock strategy succeeded in drawing attention, brand managers risk that these types of content may be flagged as fake because, as illustrated in this study, fake news is substantially more likely to exhibit specific negative emotions. Thus, managers can use the insights gained from this study to "check" their message appeal for being too negative and adjust accordingly.

This study is not without limitations. First, the data set contained political news articles only. As fake news in politics can have profound consequences for societies, as alluded to at the beginning of this article, it is critical to be able to detect and flag dubious political news reporting. However, this article did not investigate whether and to what extent the differences in emotional valence and specific emotions manifest in other, non-political fake narratives, which could be explored in future studies. Second, this study focused on communications originating from one important stakeholder group, i.e. third-party news organizations. Future studies could investigate fake news originating from other stakeholder groups, such as consumers (e.g. user-generated content on Twitter or other social media sites) or employees (e.g. employee reviews on Glassdoor or other employee review sites). Third, although AI tools and specifically IBM Watson have been used in previous research studies, Watson's analysis of emotions constrains the researcher to the five human emotions reported above, whereas other basic human emotions

(e.g. surprise, anticipation and trust) are excluded from the analysis. Thus, this study cannot offer any insights if and how fake narratives differ with respect to these omitted emotions; however, this could be an avenue for further research. Third, when using a large data set that contained more than 85,000 words for the text body itself, there is the possibility of a Type I error (false-positive), i.e., the test statistics show a significant result that may actually be trivial in practice. I attempted to overcome this limitation by reporting the effect size, Cohen's *d*, but further studies could confirm the findings.

Concluding remarks

The identification of fake news is a complex issue. The study presented here describes a method on how to tell fake news content apart from real news content. Specifically, it analyzes the differences in the emotional framing of the message content in fake and real news using a data set of 150 political news articles. Using an AI tool for analyzing these narratives, in addition to skilled human input, it finds that there is a significant difference between the overall emotional sentiment portrayed in the titles, and that fake news are substantially more negative with regard to the emotion dimensions disgust and anger than real news articles. In addition, fake news stories are markedly less "joyful" than real news narratives. This study adds to the literature on fake news, in addition to providing a number of implications for marketing communication professionals. In summary, it suggests that human approaches to fake news identification can be augmented by automated tools, like AI, and in turn, human intelligence can augment "machine intelligence". Armed with powerful tools, like AI, managers can be more effective and efficient at telling fake and real news apart.

Note: www.uccs.edu/lbecker/

References

- Aaker, D.A. (1991), *Managing Brand Equity*, The Free Press, Simon and Schuster, New York, NY.
- Allcott, H. and Gentzkow, M. (2017), "Social media and fake news in the 2016 election", *Journal of Economic Perspectives*, Vol. 31 No. 2, pp. 211-236.
- Ambler, T. and Burne, T. (1999), "The impact of affect on memory of advertising", *Journal of Advertising Research*, Vol. 39, pp. 25-34.
- Andrea, M. (2016), "Aesthetics of fake: an overview", *Aesthetics*, Firenze University Press, Via Cittadella, p. 59.
- Baccarella, C., Wagner, T., Kietzmann, J.H. and McCarthy, I. P. (2018), "Social media? It's serious! Understanding the dark side of social media", *European Management Journal*, Vol. 36 No. 4, pp. 431-438.
- Bakir, V. and McStay, A. (2018), "Fake news and the economy of emotions", *Digital Journalism*, Vol. 6 No. 2, pp. 154-175. Baly, R. Karadzhev, G. Glass, J. and Nakov, P. (2018), "Predicting factuality of reporting and bias of news media sources", *ArXiv Preprint ArXiv:1810.01765*.
- Berthon, P.R. and Pitt, L.F. (2018), "Brands, truthiness and post-fact: managing brands in a post-rational world", *Journal of Macromarketing*, Vol. 38 No. 2, pp. 218-227.
- Bessi, A., Coletto, M., Davidescu, G.A., Scala, A., Caldarelli, G. and Quattrociocchi, W. (2015), "Science vs conspiracy: collective narratives in the age of misinformation", *PLoS One*, Vol. 10 No. 2, p. e0118093.
- Betteridge, I. (2009), "TechCrunch: irresponsible journalism", *TechCrunch*, 23 February, available at: <https://web.archive.org/web/20090226202006/http://www.technovia.co.uk/2009/02/techcrunch-irresponsible-journalism.html> (accessed 22 November 2018).
- Cambria, E., Schuller, B., Xia, Y. and Havasi, C. (2013), "New avenues in opinion mining and sentiment analysis", *IEEE Intelligent Systems*, Vol. 28 No. 2, pp. 15-21.
- Chen, Y., Conroy, N.J. and Rubin, V.L. (2015), "Misleading online content: recognizing clickbait as 'false news'", *Proceedings of the 2015 ACM on Workshop on Multimodal Deception Detection*, pp. 15-19.
- Cotte, J. and Ritchie, R. (2005), "Advertisers' theories of consumers: why use negative emotions to sell?", *ACR North American Advances*.
- Crawshaw, S. (1995), "Benetton sued over shock ads", *The Independent*, 22 January.
- Dahl, D.W., Frankenberger, K.D. and Manchanda, R.V. (2003), "Does it pay to shock? Reactions to shocking and nonshocking advertising content among university students", *Journal of Advertising Research*, Vol. 43 No. 3, pp. 268-280.
- Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G. and Quattrociocchi, W. (2016), "The spreading of misinformation online", *Proceedings of the National Academy of Sciences*, Vol. 113 No. 3, pp. 554-559.
- Dens, N. and Pelsmacker, P.D. (2010), "How advertising strategy affects brand and USP recall for new brands and extensions", *International Journal of Advertising*, Vol. 29 No. 2, pp. 165-194.
- Dens, N., Pelsmacker, P.D. and Janssens, W. (2008), "Exploring consumer reactions to incongruent mild disgust appeals", *Journal of Marketing Communications*, Vol. 14 No. 4, pp. 249-269.
- Dewey, C. (2016), "Facebook fake-news writer: 'I think Donald Trump is in the White House because of me'", *The Washington Post*, 17 November, available at: www.washingtonpost.com/news/the-intersect/wp/2016/11/17/facebook-fake-news-writer-i-think-

donald-trump-is-in-the-white-house-because-of-me/?noredirect=on&utm_term=.ff04779f63b8 (accessed 22 November 2018).

Engel, P. (2014), "Here are the most- and least-trusted news outlets in America", Business Insider, 21 October, available at: <http://uk.businessinsider.com/here-are-the-most-and-least-trusted-news-outlets-in-america-2014-10> (accessed 11 November 2018).

Ethical Journalism Network (n.d.), "The EJM definition of fake news", available at: <https://ethicaljournalismnetwork.org/tag/fake-news/page/2>

Etzioni, O., Banko, M., Soderland, S. and Weld, D.S. (2008), "Open information extraction from the Web", Communications of the ACM, Vol. 51 No. 12, pp. 68-74.

Friestad, M. and Thorson, E. (1986), "Emotion-eliciting advertising: effects on long term memory and judgment", Advances in Consumer Research, Vol. 13 No. 1, pp. 111-116.

Graves, L. (2018), Understanding the Promise and Limits of Automated Fact-Checking, Reuters Institute and University of Oxford, Factsheet, available at: <https://agency.reuters.com/content/dam/openweb/documents/pdf/news-agency/report/reuters-institute-graves-factsheet-180228.pdf> (accessed 30 November 2018).

Habgood-Coote, J. (2018), "The term 'fake news' is doing great harm", The Conversation, 27 July, available at: <https://theconversation.com/the-term-fake-news-is-doing-great-harm-100406> (accessed 3 December 2018).

Hair, J.F.J., Wolfinbarger, M., Money, A.H., Samouel, P. and Page, M.J. (2011), Essentials of Business Research Methods, 2 edition, Routledge, Armonk, N.Y.

Hannah, D.R., McCarthy, I.P. and Kietzmann, J. (2015), "We're leaking, and everything's fine: how and why companies deliberately leak secrets", Business Horizons, Vol. 58 No. 6, pp. 659-667.

Hao, K. (2018), "Even the best AI for spotting fake news is still terrible", MIT Technology Review, 3 October, available at: www.technologyreview.com/s/612236/even-the-best-ai-for-spotting-fake-news-is-still-terrible/ (accessed 27 December 2018).

Heath, R. (2007), "Emotional persuasion in advertising: a hierarchy-of-processing model", available at: <https://researchportal.bath.ac.uk/en/publications/emotional-persuasion-in-advertising-a-hierarchy-of-processing-mod> (accessed 19 December 2018).

Helmore, E. (2016), "Kellogg pulls ads from Breitbart News: site isn't 'aligned with our values'", The Guardian, 29 November, available at: www.theguardian.com/media/2016/nov/29/kellogg-pulls-ads-from-breitbart-news-alt-right (accessed 17 December 2018).

Hern, A. (2018), "Facebook advertises for 'fake news' fact checkers", The Guardian, 8 June, available at: www.theguardian.com/technology/2018/jun/08/facebook-advertises-for-fake-news-fact-checkers

Horne, B.D. and 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", presented at the International AAAI Conference on Web and Social Media, available at: <https://arxiv.org/pdf/1703.09398.pdf>

Horne, B.D. Dron, W. Khedr, S. and Adali, S. (2018), "Sampling the news producers: a large news and feature data set for the study of the complex media landscape", ArXiv

Preprint ArXiv:1803.10124

IBM (2018), "IBM clouds/natural language understanding", 8 November, available at: <https://console.bluemix.net/docs/services/natural-language-understanding/getting-started.html#analyze-phrase>

Johar, J.S. and Sirgy, M.J. (1991), "Value-expressive versus utilitarian advertising appeals: when and why to use which appeal", *Journal of Advertising Research*, Vol. 20 No. 3, pp. 23-33.

Keller, K.L. (1993), "Conceptualizing, measuring, and managing customer-based brand equity", *Journal of Marketing*, Vol. 57 No. 1, pp. 1-22.

Kershner, J.W. (2005), *The Elements of News Writing*, Pearson Allyn and Bacon, Boston.

Kietzmann, J., Paschen, J. and Treen, E. (2018), "Artificial intelligence in advertising how marketers can leverage artificial intelligence along the consumer journey", *Journal of Advertising Research*, Vol. 58 No. 3, pp. 263-267.

Kovach, B. and Rosenstiel, T. (2014), *The Elements of Journalism*, 3rd ed., Three Rivers Press, New York, NY.

Kumar, S. and Shah, N. (2018), "False information on web and social media: a survey", presented at the *Social Media Analytics: Advances and Applications*.

Kwon, S., Cha, M., Jung, K., Chen, W. and Wang, Y. (2013), "Prominent features of rumor propagation in online social media", *IEEE 13th International Conference on Data Mining*, presented at the 13th International Conference on Data Mining, pp. 1103-1108.

Lang, A. (2000), "The limited capacity model of mediated message processing", *Journal of Communication*, Vol. 50 No. 1, pp. 46-70.

Magdy, A. and Wanas, N. (2010), "Web-based statistical fact checking of textual documents", *Proceedings of the 2nd International Workshop on Search and Mining User-Generated Contents, SMUC '10*, ACM, New York, NY, pp. 103-110.

Majumder, N., Poria, S., Gelbukh, A. and Cambria, E. (2017), "Deep learning-based document modeling for personality detection from text", *IEEE Intelligent Systems*, Vol. 32 No. 2, pp. 74-79.

Mandler, G. (1982), "The structure of value: accounting for taste", *Affect and Cognition: The Seventeenth Annual Carnegie Symposium on Cognition*, Erlbaum, Hillsdale, pp. 3-36.

Meenaghan, T. (1995), "The role of advertising in brand image development", *Journal of Product & Brand Management*, Vol. 4 No. 4, pp. 23-34.

Mostafa, M., Crick, T., Calderon, A.C. and Oatley, G. (2016), "Incorporating emotion and personality-based analysis in user-centered modelling", in Bramer M. and Petridis M. (Eds), *Research and Development in Intelligent Systems XXXIII*, SGA1, presented at the *International Conference on Innovative Techniques and Applications of Artificial Intelligence*, Springer, Cham, pp. 383-389.

Mukherjee, S. and Weikum, G. (2015), "Leveraging joint interactions for credibility analysis in news communities", *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, ACM, pp. 353-362.

Olson, P. (2016), "How Facebook helped Donald Trump become President", *Forbes*, 9 November, available at: www.forbes.com/sites/parmyolson/2016/11/09/how-facebook-helped-donald-trump-become-president/#11d968cb59c5 (accessed 22 November 2018).

Oxford English Dictionary (2018), "Main definitions of 'fake'", Oxford University Press, available at: <https://en.oxforddictionaries.com/definition/fake> (accessed 3 December 2018).

Parkinson, H.J. (2016), "Click and elect: how fake news helped Donald Trump win a real election", *The Guardian*, 14 November, available at: www.theguardian.com/commentisfree/2016/nov/14/fake-news-donald-trump-election-alt-right-social-media-tech-companies (accessed 22 November 2018).

Pitt, C., Kietzmann, J., Botha, E. and Wallström, Å. (2018), "Emotions and sentiment: an exploration of artist websites", *Journal of Public Affairs*, Vol. 18 No. 2, pp. e1653.

Pitt, C., Plangger, K.A., Botha, E., Kietzmann, J. and Pitt, L.F. (2017), "How employees engage with B2B brands on social media: word choice and verbal tone", *Industrial Marketing Management*, 74.

Plangger, K.A., Matheson, K. and Kietzmann, J. (2019), "The serious side to funny cartoons: understanding public perception through cartoon content analysis", *Journal of Wine Research*.

Plutchik, R. (2001), "The nature of emotions: human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice", *American Scientist*, Vol. 89 No. 4, pp. 344-350.

Popat, K., Mukherjee, S., Strötgen, J. and Weikum, G. (2016), "Credibility assessment of textual claims on the web", *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, ACM, pp. 2173-2178.

Potthast, M., Kiesel, J., Reinartz, K., Bevendorff, J. and Stein, B. (2017), "A stylometric inquiry into hyperpartisan and fake news", No. arXiv preprint arXiv:1702.05638.

Read, M. (2016), "Donald Trump won because of Facebook", *NY Magazine*, 9 November, available at: <http://nymag.com/intelligencer/2016/11/donald-trump-won-because-of-facebook.html> (accessed 22 November 2018).

Richardson, B. (2007), *The Process of Writing News: From Information to Story*, Pearson, Boston:

Rothman, J. (2018), "In the age of A.I., is seeing still believing?", *The New Yorker*, 12 November, available at: www.newyorker.com/magazine/2018/11/12/in-the-age-of-ai-seeing-still-believing (accessed 29 November 2018).

Rubin, V., Conroy, N. and Cornwell, S. (2016), "Fake news or truth? Using satirical cues to detect potentially misleading news", *Proceedings of the Second Workshop on Computational Approaches to Deception Detection*, pp. 7-17.

Short, J.C., McKenny, A.F. and Reid, S.W. (2018), "More than words? Computer-aided text analysis in organizational behavior and psychology research", *Annual Review of Organizational Psychology and Organizational Behavior*, Vol. 5 No. 1, pp. 415-435.

Stephens, M.A. (1988), *History of News*, Penguin, New York, NY.

Strickland, E. (2018), "AI-human partnerships tackle 'fake news'", *IEEE Spectrum*, Vol. 55 No. 9, pp. 12-13.

Tandoc Jr, E.C., Lim, Z.W. and Ling, R. (2018), "Defining 'fake news': a typology of scholarly definitions", *Digital Journalism*, Vol. 6 No. 2, pp. 137-153.

Taute, H.A., McQuitty, S. and Sautter, E.P. (2011), "Emotional information management and responses to emotional appeals", *Journal of Advertising Research*, Vol. 40 No. 3, pp. 31-44.

Vosoughi, S., Roy, D. and Aral, S. (2018), "The spread of true and false news online", *Science* (New York, N.Y.), Vol. 359 No. 6380, pp. 1146-1151.

Wakefield, J. (2019), "Facebook employs UK fact-checkers to combat fake news", *BBC News*, 11 January, available at: www.bbc.com/news/technology-46836897

Wang, L.X., Ramachandran, A. and Chaintreau, A. (2016), "Measuring click and share dynamics on social media: a reproducible and validated approach", *Proceedings of Tenth International AAAI Conference on Web and Social Media*, presented at the Tenth International AAAI Conference on Web and Social Media, AAAI Publications.

Zimdars, M. (2016), "False, misleading, clickbait-y, and/or satirical 'news' sources", 13 December, available at: www.justicemitchell.com/justice-mitchell/2016/12/13/false-misleading-clickbait-y-andor-satirical-news-sources (accessed 11 November 2018).