

# Profiling Fake News Spreaders on Social Media through Psychological and Motivational Factors

Mansoorah Karami  
Arizona State University  
Tempe, AZ, USA  
mkarami@asu.edu

Tahora H. Nazer  
Spotify  
Boston, MA, USA  
tahoranazer@spotify.com

Huan Liu  
Arizona State University  
Tempe, AZ, USA  
huanliu@asu.edu

## ABSTRACT

The rise of fake news in the past decade has brought with it a host of consequences, from swaying opinions on elections to generating uncertainty during a pandemic. A majority of methods developed to combat disinformation either focus on fake news content or malicious actors who generate it. However, the virality of fake news is largely dependent upon the users who propagate it. A deeper understanding of these users can contribute to the development of a framework for identifying users who are likely to spread fake news. In this work, we study the characteristics and motivational factors of fake news spreaders on social media with input from psychological theories and behavioral studies. We then perform a series of experiments to determine if fake news spreaders can be found to exhibit different characteristics than other users. Further, we investigate our findings by testing whether the characteristics we observe amongst fake news spreaders in our experiments can be applied to the detection of fake news spreaders in a real social media environment.

## CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; • **Applied computing** → *Psychology*.

## KEYWORDS

Fake News; Fake News Spreader; Psychological Theories; Social and Behavioral Studies; Social Media

### ACM Reference Format:

Mansoorah Karami, Tahora H. Nazer, and Huan Liu. 2021. Profiling Fake News Spreaders on Social Media through Psychological and Motivational Factors. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media (HT '21)*, August 30–September 2, 2021, Virtual Event, Ireland. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3465336.3475097>

## 1 INTRODUCTION

Due to the increasing amount of our time spent on social media platforms, it is no surprise that people tend to receive their news content through social media more than before. One in five U.S.

adults used social media as their main source of political and election news for the US presidential election in 2020 [29]. The high rate of engagement with online news can be mainly attributed to the nature of the social media platforms themselves. Social Media is typically inexpensive, provides easy access to users, and supports fast dissemination of information that is not possible through traditional media outlets. However, despite these advantages, the quality of news on social media is considered lower than that of traditional news outlets. A factor contributing to this low quality is the widespread nature of fake news articles online. Fake news is a piece of false information published by news outlets to mislead consumers [40, 55].

Fake news has several significant negative effects on civil society. First, people may accept deliberate lies as truths. The likelihood of accepting fake news as true increases after repeated exposure [16], especially when the content aligns with the user's beliefs [17, 53]. Second, fake news may change the way people respond to legitimate news. When people are inundated with fake news, the line between fake news and real news becomes more uncertain. Fake news spreaders make users doubt the nature of real news and create the idea that everything is biased and conflicted, and it is impossible to distinguish fake from real news [25]. Finally, the prevalence of fake news has the potential to break the trustworthiness of the entire news ecosystem. For instance, despite traditional domains such as the New York Times, the Washington Post, and CNN being among the most shared COVID-19-related stories on Twitter, a fake news domain, Gateway Pundit, was ranked 4th in August and 6th in September of 2020 among the most shared domains for URLs about COVID-19 [23]. Therefore, it is critical to develop methods that detect and mitigate fake news, with the purpose of benefiting the general public and the entire news ecosystem.

Detecting fake news is a challenging task because it is designed to be indistinguishable from real news and intentionally misleading. As a result, the features extracted from the content are not enough to build an accurate detection method. Researchers have tackled this problem by using the characteristics of posts, user networks, and user profiles. However, factors that motivate users to spread fake news on social media and their importance in fake news detection have remained understudied. Studies [11, 35] have shown that users are more prone to propagating fake news stories when the situation is uncertain, they are emotionally overwhelmed and anxious, the topic of discussion is of personal importance to them, and they do not have primary control over the situation through their actions. In this work, we provide a methodology to extract factors that are not reliant on currently existing self-reported survey methods [5, 45]. Our analysis aims to shed light on the attitude of spreading fake

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

HT '21, August 30–September 2, 2021, Virtual Event, Ireland.

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8551-0/21/08...\$15.00

<https://doi.org/10.1145/3465336.3475097>

news, rather than detecting potential fake news among the users' tweets. We investigate the following research questions:

**RQ1:** *What motivates users to spread fake news?*

**RQ2:** *What are the differences between users who spread fake news and the users who spread real news in terms of motivational factors determined in RQ1?*

**RQ3:** *Can we use the differences between fake news spreaders and other users to profile fake news spreaders?*

We make the following contributions in this work:

- We study a novel problem by investigating the relationship between motivational factors and spreading fake news on social media in Sections 2 and 3. This study lays the foundation of exploiting the effectiveness of these factors for preventing the dissemination of fake news;
- We perform a statistical comparative analysis between users who spread fake news and other users. We hypothesize that users who are keen to spread disinformation may have different motivating characteristics compared to other users. Our study shows that they are significantly different in terms of motivational factors in Section 4.1;
- We conduct experiments on real-world datasets to show whether the extracted features add additional value to the representation of the users' tweets in Section 4.2.

## 2 MOTIVATIONAL FACTORS IN SPREADING FAKE NEWS

To investigate the factors that motivate users to spread fake news on social media (RQ1), we begin with enumerating a list of motivational factors in spreading fake news: *uncertainty, anxiety, lack of control, relationship enhancement, and rank*.

### 2.1 Uncertainty

Spreading fake news can be a sense-making activity in ambiguous situations. The frequency of fake news sharing increases in uncertain situations such as global pandemics, when the details of the spread are unknown, or political events like an election when people are unsure of the results [11, 35]. When a crisis happens, people first seek information from credible sources [9]. If there is no such information, people tend to form unofficial social networks to make predictions based on their own judgment and fill the information gap [37]. This might result in the generation of fake news such as injecting one's self with bleach or consumption of highly-concentrated alcohol to kill the SARS-CoV-2 novel coronavirus [34]. As the uncertainty increases, the reliance on firm beliefs and the unity among the users with the same ideology or in the same group reduces. Hence, users are more prone to accept new information, even false, as a compromise to resolve the uncertain uncertainty. Uncertainty can cause emotions like anxiety and anger which can affect the spread of fake news in other ways [28, 51].

### 2.2 Emotions

Emotional pressure can play an important role in spreading fake news and can be triggered by emotions like anxiety. Anxiety can make people prone to spreading unproven and less accurate claims when transmitting information [11]. In high anxiety situations, fake news can serve to justify the user's feelings and relieve emotional

tension [1]. Fake news might be used as a method of expressing emotions in anxious situations that allows people to talk about their concerns and receive feedback informally. This results in sense-making and problem-solving [48]. For example, during the devastating Hurricane Harvey, 2017, a fake news story accusing Black Lives Matter supporters of blocking first responders from reaching the affected area was spread by more than one million Facebook users [14]. Believing and spreading such fake news stories may help the people in disaster areas cope with the anxiety caused by delays in relief efforts [12]. Anxiety can reduce the effects of motivated reasoning [51]. Users with higher traits of anxiety pay closer attention to the contemporary information even when it is against their ideology or bias [27].

### 2.3 Lack of control

Fake news represents a method for coping with uncertain and uncontrollable situations. When people are uncertain about their situation, they tend to experience the feeling of lack of control [4], which leads to anxiety. When people do not have primary control over their situation, they resort to secondary control strategies consisting of emotional responses such as predicting the worst to avoid disappointment or attributing events to chance. Some additional secondary control themes are interpreting the meaning of the event and predicting future events [49]. Studies show that fake news can be used as a way of regaining control over uncertain and uncontrollable circumstances. An example of a situation that temporarily induces the feeling of losing control and the subsequent attempt to regain control is when individuals enter a new environment. In such a setting, individuals engage in activities to manifest their control, boost their performance, and increase satisfaction in order to engage more with their environment [3]. The desire to actively engage with the environment can make social media newcomers more vulnerable to fake news because they are less familiar with the environment and more eager to gain the attention of their network.

### 2.4 Relationship Enhancement

Enhancing social relations is one of the goals of spreading unverified information [11] and conspiracies [20]. In the formation of short term social relations with the purpose of generating a positive effect on others, people may spread information to attract or hold attention without verifying its authenticity. People prefer to follow, engage with, and endorse those that they like. This liking is caused by three factors: (1) people that are similar to us, (2) people that pay us compliments, and (3) people that cooperate with us towards mutual goals [8]. All of these factors can be observed in social cyber communities and cyber tribes. In other words, users on social media tend to share and like news more when it comes from the sources or the communities they follow [33]. Also, in friendly relations, negative fake news might be transmitted as a warning mechanism to provide information that is relevant, useful, and prevents harmful consequences [52]. Essentially, group members want to show that they are looking out for one another.

### 2.5 Rank

Researchers observed that employees with the lowest rank have a high contribution to the circulation of information that is confirmed

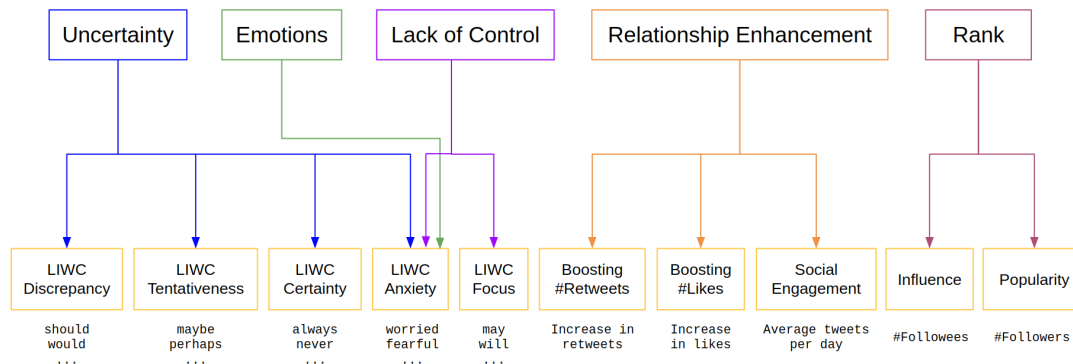


Figure 1: Summary of the metrics used to measure features of users who spread fake news with sample words used.

to be false [30]. Moreover, spreading fake news can be considered as a self-enhancing mechanism. Self-enhancement motivates individuals to either consciously spread the content for personal gain or unconsciously spread the content that is consistent with their previous beliefs and biases [5, 11]. Therefore, the motivation to improve social rank can encourage users to spread fake news on social media. There is not a specific way of defining a rank for users on social media. Measures such as popularity or influence are attempts of indicating ranks on social networks.

### 3 DATA PREPARATION FOR STUDYING FAKE NEWS SPREADERS

We use two datasets from FakeNewsNet repository [41] to explain our method and experiments in Section 4: PolitiFact and GossipCop<sup>1</sup>. The data consists of fake and real stories identified by PolitiFact and GossipCop websites. To identify the users, we utilized a crawler to collect all the tweets referring to those stories. Recent posts, profile, and list of followers and friends of Twitter users who posted about fake and real stories were also collected. We treat each user as a document and the users’ recent posts are considered as the content of that document. We based this method on the ex-

Table 1: Statistics of the datasets.

Dataset	News Spreader	# of Users	# of Tweets
PolitiFact	Real	12,511	88,930
	Fake	6,309	45,184
GossipCop	Real	2,794	134,055
	Fake	11,560	267,803

periments done by Yang et al. [54], using the same parameter of about 150 words per document. Since each tweet consists of at most 280 characters (which on average is about 55 words), if there is a mention or retweet of three or more fake news, that user is tagged as a user who is keen to spread fake news, otherwise, that user is labeled as real news spreader.

All the features introduced in Section 2 can be extracted from the content (tweets), activity (tweeting behavior and likes), or network (followees and followers) of Twitter users in PolitiFact

and GossipCop datasets. For measuring the features that are extracted from content, we exploit Linguistic Inquiry and Word Count (LIWC) [46]. LIWC includes a dictionary that lists a set of words for some psychologically-relevant categories including positive and negative emotions, social relationships, honesty and deception.

#### 3.1 Feature Extraction

We pre-processed all tweets to remove punctuation, URLs, hashtags, and mentions. To measure a feature using an LIWC category, we find the percentage of words in a tweet that belongs to that category. Next, we explain how each of the features are measured:

- Selected LIWC features: there are three LIWC categories that are related to *uncertainty*: Discrepancy (e.g. should, would, and could), Tentativeness (e.g. maybe, perhaps, and guess), and Certainty (e.g. always and never). These categories are abbreviated as *discrep*, *tentat*, and *certain*, respectively. *Anxiety* can be measured using the LIWC Anxiety category (*anx*) which includes words such as nervous, afraid, and tense. We use LIWC Future Focus (*futurefocus*) to measure *lack of control*, this category includes words such as may, will, and soon.
- Social Engagement: we measure social engagement on Twitter using the average number of tweets per day.
- Position in the Network: this feature can be quantified using a variety of metrics when the network structure is known. However, in the case of social networks between Twitter users, we do not have complete structure and even collecting local information is time consuming due to the rate limitation of Twitter APIs. Hence, we use the information available in our datasets and extract *influence* using the number of followees and *popularity* using the number of followers of each user.
- Relationship Enhancement and Liking: improving the relation to other social media users and gaining more attention from the community is one of the motivations for spreading fake news. If the number of retweets and likes of a user’s fake news tweet is higher than the average number of the user’s retweets and likes, it can indicate that this user has enhanced their social relation and initiated conversation.

<sup>1</sup>The source code for this work is available at <https://github.com/mansourehk/profiling>.

Consequently, we use the difference between these two values as indicators of relationship enhancement.

The summary of features and the metrics used to measure them is presented in Figure 1. As explained in Section 2, some metrics overlap in interpreting the motivational features.

## 4 PROFILING FAKE NEWS SPREADERS

We study the differences between users who spread fake news and the ones who spread real news in terms of the feature categories introduced in Section 2 (RQ2). We perform two experiments towards this goal. First, we set our null hypothesis to be that these two groups have the same mean in five categories of features. If the results of t-test shows that there is a significant difference (p-value less than 0.05) between two groups, we can reject the null hypothesis. Second, we use the proposed features to see if a supervised model (a classifier) can use these features to discriminate the two groups of users. The byproduct of this experiment also validates if these features can add more information to the representation of the user’s data (RQ3).

Feature Name	t-statistic Politifact	t-statistic GossipCop
Tentativeness	-22.91**	17.80**
Discrepancy	-16.13**	16.43**
Certainty	-8.67**	16.36**
Anxiety	-1.06*	1.02
Lack of Control	-14.66**	19.18**
Social Engagement	-13.12**	-0.41
Influence	-0.26	2.74**
Popularity	-2.69*	2.10*
Boosting #tweets	-3.78**	9.70**
Boosting #likes	-1.54	-0.70

**Table 2: Comparison between fake news and real news spreaders in terms of psychological features. The features that are significantly different between users who spread fake news and the users who spread real news are marked with \*\* (p-value <0.005) or \* (p-value <0.05)**

### 4.1 Significance of Psychological Features

As shown in Table 2, for the features, we observe a significant difference (p-value < 0.05) between users who share fake news and the ones who share real news in terms of most of the features. Notably, the T-test for Anxiety in the GossipCop dataset is not significant. We reason that this is because GossipCop articles are mostly concerned with celebrity news and regardless of the true or false nature of the news being propagated, the news has a low potential of causing distress or being propagated to address the anxiety of the spreader. On the other hand, Politifact contains news articles about a variety of topics including natural disasters and politics. Because these topics are serious, there is more potential for these articles to be circulated due to the anxiety of the spreader.

In social engagement, Politifact users who share fake news have significantly lower number of tweets per day but we do not observe a significant difference among users in the GossipCop users. We used the Influence (number of followers) and popularity (number of

followers) as indicators of the status in the network. Except for the Influence in the Politifact users, we observe a significant difference among users who spread real news and the ones spreading fake news. We also expected to see different number of likes and retweets for the fake news tweets which is only observed in reshares feature; indicating that Boosting #likes can be a motivation for spreading fake news on social media.

Politifact Dataset		
Model	Accuracy	F1-Score
Bert	78.61	58.02
Bert+Features	90.0	83.11
GossipCop Dataset		
Model	Accuracy	F1-Score
Bert	90.36	94.61
Bert+Features	92.69	96.02


**Table 3: Performance and F1-score of downstream task using sentence representation of BERT with and without proposed features.**

### 4.2 Fake News Spreader Classifier

In the next step we examine if adding the proposed features to the representation of a text will give us a better fake news spreader classifier as our downstream task. For this section, we use Bidirectional Encoder Representations from Transformers (BERT) [10] as our model for sentence embedding. Each Twitter user’s timeline data were represented by a vector  $H_B$  by BERT and their motivational factors were represented by a vector  $H_f$ . We examined the performance of BERT as our baseline classification model by only using  $H_B$ . For the Bert+Features model the vectors are concatenated together,  $H = \text{Concat}(H_B, H_f)$ , and were passed through a feed-forward neural network to balance the effect of the two vectors [19]. The performance of the two models is presented in Table 3. This experiment shows that by adding the proposed features to context representation, we can add more value to the representation of the data for classifying fake news versus real news spreaders. We give an illustrative visualization of the embeddings of 1000 random samples of Politifact and GossipCop Datasets by both BERT and BERT+Features with the t-SNE tool [47] (Figure 2). We observe that BERT+Features can learn more discriminative embeddings.

## 5 RELATED WORK

Fake news detection has gained increased research attention in recent years. Fake news detection methods generally focus on using *news content* or *social context* features [42]. In news content-based approaches, features are extracted linguistically or visually. Linguistic-based features capture textual information that appears in fake news content [36], such as lexical and syntactic features. Visual-based features are used to identify fake images [15] that are intentionally created. For social context-based approaches, the features include user-based, post-based, and network-based. User-based features are extracted from user profiles [43] to measure their



**Figure 2: The t-sne visualization of features for the BERT ((a)- and (b)-left) and BERT+Features ((a)- and (b)-right). Adding the motivational features will create a better discriminative embeddings.**

characteristics and credibility [7]. Post-based features represent users’ social responses in terms of stance [18], topics [26], or credibility [7]. Network-based features are extracted by constructing specific networks, such as diffusion networks [21].

Recent advancements utilize deep learning approaches to learn representations from the text, image, and network data for fake news detection. For example, convolution neural networks have been utilized to learn textual information from news content [50], and recurrent neural networks have been adopted to learn temporal representations from the propagation of fake news [24, 38]. Shu *et al.* utilize a network embedding approach to model the relations of the publisher, news, and users to detect fake news [44]. Deep geographic models such as graph convolution neural networks have been exploited to model the propagation networks as well as to detect fake news [31].

On the other hand, the task of profiling users has long been studied in both formal and informal languages [2, 39]. Several explicit (from metadata) and implicit cues can be used for user profiling such as user identity linkage, age, and gender [43]. Moreover, in recent research, factors reported from disciplines such as psychology and behavioral studies helped in profiling users as well. For instance, Giachanou *et al.* [13] combined emotional signals with a Long Short Term Memory (LSTM) neural network to detect credible versus non-credible claims. Furthermore, Cardaioli *et al.* [6] utilized behavioral-based features such as Big Five personality and stylometric features for discriminating a fake news spreader.

## 6 CONCLUSION AND FUTURE WORK

In this work, we investigated whether the psychological features that are observed in users who spread fake news in behavioral studies on human subjects are also visible in social media users who spread fake news on social media. Towards this goal, we introduced five categories of features based on psychological theories that can be quantified for social media users. Based on our observations on two real-world datasets, we observed that (i) social media users who spread fake news are significantly different in terms of the majority of these features and (ii) these features have predictive power in the detection new and unobserved fake news spreaders.

This study is a first step towards understanding users who are exposed to fake news. We wish to continue this work by looking into what other psychological features associated with individual

differences such as personality traits may be related to fake news susceptibility [22] as well as providing stronger proxy for each feature. Moreover, we plan to study the application of the psychological features in improving the performance of state-of-the-art fake news detection algorithms along with providing explanations for such algorithms [32].

## ACKNOWLEDGEMENTS

The authors would like to thank Dr. H. Russell Bernard (ASU) and Tyler Black (ASU) for their comments on the manuscript. This material is, in part, based upon works supported by Parallax Advanced Research Corporation (11076-ASU) and ONR (N00014-21-1-4002).

## REFERENCES

- [1] Gordon W Allport and Leo Postman. 1947. The psychology of rumor. (1947).
- [2] Shlomo Argamon, Moshe Koppel, Jonathan Fine, and Anat Rachel Shimoni. 2003. Gender, genre, and writing style in formal written texts. *Text & Talk* 23, 3 (2003), 321–346.
- [3] Susan J Ashford and J Stewart Black. 1996. Proactivity during organizational entry: The role of desire for control. *Journal of Applied psychology* 81, 2 (1996), 199.
- [4] Prashant Bordia, Elizabeth Hobman, Elizabeth Jones, Cindy Gallois, and Victor J Callan. 2004. Uncertainty during organizational change: Types, consequences, and management strategies. *Journal of business and psychology* 18, 4 (2004), 507–532.
- [5] Tom Buchanan. 2020. Why do people spread false information online? The effects of message and viewer characteristics on self-reported likelihood of sharing social media disinformation. *Plos one* 15, 10 (2020), e0239666.
- [6] Matteo Cardaioli, Stefano Ceconello, Mauro Conti, Luca Pajola, and Federico Turrin. 2020. Fake news spreaders profiling through behavioural analysis. In *CLEF*.
- [7] Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on twitter. In *Proceedings of the WWW conference*. ACM, 675–684.
- [8] Robert B Cialdini. 2007. *Influence: The psychology of persuasion*. Vol. 55. Collins New York.
- [9] Right Be Credible. 2014. CERC: Psychology of a Crisis. (2014).
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers. *arXiv:1810.04805* (2018).
- [11] Nicholas DiFonzo and Prashant Bordia. 2007. *Rumor psychology: Social and organizational approaches*. American Psychological Association.
- [12] Manny Fernandez, Lizette Alvarez, and Ron Nixon. 2017. Still Waiting for FEMA in Texas and Florida After Hurricanes. <https://www.nytimes.com/2017/10/22/us/fema-texas-florida-delays-.html> Accessed 11 March 2019.
- [13] Anastasia Giachanou, Paolo Rosso, and Fabio Crestani. 2019. Leveraging emotional signals for credibility detection. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 877–880.
- [14] Ryan Grenoble. 2017. Hurricane Harvey Is Just The Latest In Facebook’s Fake News Problem. [https://www.huffingtonpost.com/entry/facebook-hurricane-harvey-fake-news\\_us\\_59b17900e4b0354e441021fb](https://www.huffingtonpost.com/entry/facebook-hurricane-harvey-fake-news_us_59b17900e4b0354e441021fb) Accessed 11 March 2019.
- [15] Aditi Gupta, Hemank Lamba, Ponnurangam Kumaraguru, and Anupam Joshi. 2013. Faking sandy: characterizing and identifying fake images on twitter during

- hurricane sandy. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, 729–736.
- [16] Lynn Hasher, David Goldstein, and Thomas Toppino. 1977. Frequency and the conference of referential validity. *Journal of verbal learning and verbal behavior* 16, 1 (1977), 107–112.
- [17] Bohan Jiang, Mansoorreh Karami, Lu Cheng, Tyler Black, and Huan Liu. 2021. Mechanisms and Attributes of Echo Chambers in Social Media. *arXiv preprint arXiv:2106.05401* (2021).
- [18] Zhiwei Jin, Juan Cao, Yongdong Zhang, and Jiebo Luo. 2016. News Verification by Exploiting Conflicting Social Viewpoints in Microblogs. In *AAAI*. 2972–2978.
- [19] Mansoorreh Karami, Ahmadrza Mosallanezhad, Michelle V Mancenido, and Huan Liu. 2020. "Let's Eat Grandma": When Punctuation Matters in Sentence Representation for Sentiment Analysis. *arXiv preprint arXiv:2101.03029* (2020).
- [20] Yubo Kou, Xinning Gui, Yunan Chen, and Kathleen Pine. 2017. Conspiracy Talk on Social Media: Collective Sensemaking During a Public Health Crisis. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW, Article 61 (Dec. 2017), 21 pages. <https://doi.org/10.1145/3134696>
- [21] Sejeong Kwon, Meeyoung Cha, Kyomin Jung, Wei Chen, and Yajun Wang. 2013. Prominent features of rumor propagation in online social media. In *Data Mining (ICDM), 2013 IEEE 13th International Conference on*. IEEE, 1103–1108.
- [22] Kathleen G Larson, Kelly Kaleda, and Laurie Fenstermacher. 2019. Applying cognitive psychology principles to the (dis) information environment: an examination of discourse comprehension, memory, and fusion of news articles. In *Signal Processing, Sensor/Information Fusion, and Target Recognition XXVIII*, Vol. 11018. International Society for Optics and Photonics, 1101805.
- [23] David Lazer, Matthew Baum, Katherine Ognyanova, and John Della Volpe. 2020. The State of the Nation: A 50-State COVID-19 Survey Report. (2020).
- [24] Yang Liu and Yi-Fang Brook Wu. 2018. Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [25] Michael P. Lynch. 2016. Fake News and the Internet Shell Game. <https://www.nytimes.com/2016/11/28/opinion/fake-news-and-the-internet-shell-game.html> Accessed 6 March 2019.
- [26] Jing Ma, Wei Gao, Zhongyu Wei, Yueming Lu, and Kam-Fai Wong. 2015. Detect rumors using time series of social context information on microblogging websites. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*. ACM, 1751–1754.
- [27] Michael MacKuen, Jennifer Wolak, Luke Keele, and George E Marcus. 2010. Civic engagements: Resolute partisanship or reflective deliberation. *American Journal of Political Science* 54, 2 (2010), 440–458.
- [28] George E. Marcus. 2017. How Affective Intelligence Theory Can Help Us Understand Politics. <https://emotionresearcher.com/how-affective-intelligence-theory-can-help-us-understand-politics/> Accessed 18 November 2019.
- [29] Amy Mitchell and Mark Jurkowitz. 2020. Americans Who Mainly Get Their News on Social Media Are Less Engaged, Less Knowledgeable. <https://www.journalism.org/2020/07/30/americans-who-mainly-get-their-news-on-social-media-are-less-engaged-less-knowledgeable/>
- [30] Tanushree Mitra and Eric Gilbert. 2012. Have You Heard?: How Gossip Flows Through Workplace Email. In *ICWSM*.
- [31] Federico Monti, Fabrizio Frasca, Davide Eynard, Damon Mannion, and Michael M Bronstein. 2019. Fake News Detection on Social Media using Geometric Deep Learning. *arXiv preprint arXiv:1902.06673* (2019).
- [32] Raha Moraffah, Mansoorreh Karami, Ruocheng Guo, Adrienne Raglin, and Huan Liu. 2020. Causal interpretability for machine learning-problems, methods and evaluation. *ACM SIGKDD Explorations Newsletter* 22, 1 (2020), 18–33.
- [33] Anne Oeldorf-Hirsch and S. Shyam Sundar. 2015. Posting, commenting, and tagging: Effects of sharing news stories on Facebook. *Computers in Human Behavior* 44 (2015), 240–249. <https://doi.org/10.1016/j.chb.2014.11.024>
- [34] World Health Organization. 2020. COVID-19 Mythbusters. <https://www.who.int/emergencies/diseases/novel-coronavirus-2019/advice-for-public/myth-busters>.
- [35] Gordon Pennycook, Jonathon McPhetres, Yunhao Zhang, Jackson G Lu, and David G Rand. 2020. Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy-nudge intervention. *Psychological science* 31, 7 (2020), 770–780.
- [36] Martin Potthast, Johannes Kiesel, Kevin Reinartz, Janek Bevendorff, and Benno Stein. 2017. A Stylometric Inquiry into Hyperpartisan and Fake News. *arXiv preprint arXiv:1702.05638* (2017).
- [37] Ralph L Rosnow and Gary A Fine. 1976. *Rumor and gossip: The social psychology of hearsay*. Elsevier.
- [38] Natali Ruchansky, Sungyong Seo, and Yan Liu. 2017. Csi: A hybrid deep model for fake news detection. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. ACM, 797–806.
- [39] H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dziurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin EP Seligman, et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PLoS one* 8, 9 (2013), e73791.
- [40] Kai Shu, Amrita Bhattacharjee, Faisal Alatawi, Tahora H Nazer, Kaize Ding, Mansoorreh Karami, and Huan Liu. 2020. Combating disinformation in a social media age. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 10, 6 (2020), e1385.
- [41] Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 2020. Fakenewsnet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big data* 8, 3 (2020), 171–188.
- [42] Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake News Detection on Social Media: A Data Mining Perspective. *KDD exploration newsletter* (2017).
- [43] Kai Shu, Suhang Wang, and Huan Liu. 2018. Understanding User Profiles on Social Media for Fake News Detection. *Proceedings of the 1st international workshop on Fake MultiMedia* (2018).
- [44] Kai Shu, Suhang Wang, and Huan Liu. 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*. ACM, 312–320.
- [45] Shalini Talwar, Amandeep Dhir, Puneet Kaur, Nida Zafar, and Melfi Alrasheedy. 2019. Why do people share fake news? Associations between the dark side of social media use and fake news sharing behavior. *Journal of Retailing and Consumer Services* 51 (2019), 72–82.
- [46] Yla R Tausczik and James W Pennebaker. 2010. The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of language and social psychology* 29, 1 (2010), 24–54.
- [47] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [48] Kathryn Waddington. 2012. *Gossip and organizations*. Routledge.
- [49] Charles J Walker and Bruce Blaine. 1991. The virulence of dread rumors: A field experiment. *Language & Communication* (1991).
- [50] William Yang Wang. 2017. "liar, liar pants on fire": A new benchmark dataset for fake news detection. *arXiv preprint arXiv:1705.00648* (2017).
- [51] Brian E Weeks. 2015. Emotions, partisanship, and misperceptions: How anger and anxiety moderate the effect of partisan bias on susceptibility to political misinformation. *Journal of Communication* 65, 4 (2015), 699–719.
- [52] Mienke WH Weenig, Arieneke CWJ Groenenboom, and Henk AM Wilke. 2001. Bad news transmission as a function of the definitiveness of consequences and the relationship between communicator and recipient. *Journal of personality and social psychology* 80, 3 (2001), 449.
- [53] Kirsten Weir. 2017. Why We Believe Alternative Facts. *Monitor on Psychology* 48, 5 (May 2017), 34–39. <https://www.apa.org/monitor/2017/05/alternative-facts>
- [54] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*. 1480–1489.
- [55] Xinyi Zhou and Reza Zafarani. 2020. A survey of fake news: Fundamental theories, detection methods, and opportunities. *ACM Computing Surveys (CSUR)* 53, 5 (2020), 1–40.