



2020

## Detecting Fake News Over Online Social Media via Domain Reputations and Content Understanding

Kuai Xu

*School of Mathematical and Natural Sciences, Arizona State University, Glendale, AZ 85306, USA.*

Feng Wang

*School of Mathematical and Natural Sciences, Arizona State University, Glendale, AZ 85306, USA.*

Haiyan Wang

*School of Mathematical and Natural Sciences, Arizona State University, Glendale, AZ 85306, USA.*

Bo Yang

*Jiangxi University of Finance and Economics, Nanchang 330013, China.*

Follow this and additional works at: <https://tsinghuauniversitypress.researchcommons.org/tsinghua-science-and-technology>



Part of the [Computer Sciences Commons](#), and the [Electrical and Computer Engineering Commons](#)

### Recommended Citation

Kuai Xu, Feng Wang, Haiyan Wang et al. Detecting Fake News Over Online Social Media via Domain Reputations and Content Understanding. *Tsinghua Science and Technology* 2020, 25(1): 20-27.

This Research Article is brought to you for free and open access by Tsinghua University Press: Journals Publishing. It has been accepted for inclusion in *Tsinghua Science and Technology* by an authorized editor of Tsinghua University Press: Journals Publishing.

# Detecting Fake News Over Online Social Media via Domain Reputations and Content Understanding

Kuai Xu\*, Feng Wang, Haiyan Wang, and Bo Yang

**Abstract:** Fake news has recently leveraged the power and scale of online social media to effectively spread misinformation which not only erodes the trust of people on traditional presses and journalism, but also manipulates the opinions and sentiments of the public. Detecting fake news is a daunting challenge due to subtle difference between real and fake news. As a first step of fighting with fake news, this paper characterizes hundreds of popular fake and real news measured by shares, reactions, and comments on Facebook from two perspectives: domain reputations and content understanding. Our domain reputation analysis reveals that the Web sites of the fake and real news publishers exhibit diverse registration behaviors, registration timing, domain rankings, and domain popularity. In addition, fake news tends to disappear from the Web after a certain amount of time. The content characterizations on the fake and real news corpus suggest that simply applying term frequency-inverse document frequency (tf-idf) and Latent Dirichlet Allocation (LDA) topic modeling is inefficient in detecting fake news, while exploring document similarity with the term and word vectors is a very promising direction for predicting fake and real news. To the best of our knowledge, this is the first effort to systematically study domain reputations and content characteristics of fake and real news, which will provide key insights for effectively detecting fake news on social media.

**Key words:** social media; fake news detection; content modeling; domain reputations

## 1 Introduction

The last decade has witnessed the rapid growth and success of online social networks, which has disrupted traditional media by fundamentally changing how, who, when, and where on the distribution of the latest news stories. Unlike traditional newspapers or magazines, anyone can spread any information at any time on many open and always-on social media platforms without real-world authentications and

accountability, which has resulted in unprecedented circulation and spreadings of fake news, social spams, and misinformation<sup>[1–5]</sup>.

Driven by the political or financial incentives, the creators of fake news generate and submit these well-crafted news stories on online social media, and subsequently recruit social bots or paid spammers to push the news to a certain popularity<sup>[6–8]</sup>. The recommendation and ranking algorithms on social media, if failed to immediately detect such fake news, likely surface such news to many other innocent users who are interested in the similar topics and content of the news, thus leading to a viral spreading process on social media. These rising social spams<sup>[9]</sup>, click baits<sup>[10]</sup>, and fake news<sup>[11]</sup>, mixed with real news and credible content, create challenges and difficulties for regular Internet users to distinguish credible and fake content.

• Kuai Xu, Feng Wang, and Haiyan Wang are with School of Mathematical and Natural Sciences, Arizona State University, Glendale, AZ 85306, USA. E-mail: {kuai.xu, fwang25, haiyan.wang}@asu.edu.

• Bo Yang is with Jiangxi University of Finance and Economics, Nanchang 330013, China.

\* To whom correspondence should be addressed.

Manuscript received: 2018-10-07; revised: 2018-11-10; accepted: 2018-11-15

Towards effectively detecting, characterizing, and modeling Internet fake news on online social media<sup>[11]</sup>, this paper proposes a new framework which systematically characterizes the Web sites and reputations of the publishers of the fake and real news articles, analyzes the similarity and dissimilarity of the fake and real news on the most important terms of the news articles via term frequency-inverse document frequency (tf-idf) and Latent Dirichlet Allocation (LDA) topic modeling, as well as explores document similarity analysis via Jaccard similarity measures between fake, real, and hybrid news articles.

The contributions of this paper can be summarized as follows:

- We systematically characterize the Web sites and reputations of the publishers of the fake and real news articles on their registration patterns, Web site ages, domain rankings, domain popularity, and the probabilities of news disappearance from the Internet.
- We analyze the similarity and dissimilarity of the fake and real news on the most important terms of the news articles via tf-idf and LDA topic modeling.
- We explore document similarity between fake, real, or hybrid news articles via Jaccard similarity to distinguish, classify, and predict fake and real news.

The remainder of this paper is organized as follows. Section 2 describes the background of the fake news problem over online social media and describes datasets used in this study. Section 3 characterizes the Web sites and reputations of the publishers of the fake and real news articles, while Section 4 focuses on analyzing the similarity and dissimilarity of the fake and real news on the most important terms of the news articles. In Section 5, we show the promising direction of leveraging document similarity to distinguish fake and real news by measuring their document similarity. Section 6 summarizes related work in detecting and analyzing fake news and highlights the difference between this effort with existing studies. Finally, Section 7 concludes this paper and outlines our future work.

## 2 Background and Datasets

As online social media such as Facebook and Twitter continue to play a central role in disseminate news articles to billions of Internet users, fake and real news shares the same distribution channels and diffusion networks. The creators of fake news, motivated by a variety of reasons including financial benefits and

political campaigns, are very innovative in writing the news stories and attractive titles that convince thousands of regular people to read, like, comment, or forward. Such high engagement in a short time period can make the news go viral with little challenges or doubts on authenticity, verification, or fact checking.

In this paper, we explore the research data shared from a recent study in Ref. [12]. The data consists three datasets, each of which includes 40 very popular fake and real news stories over a 3-month time-span from dozens of fake news sites as well as well-respected major news outlets including *New York Times*, *Washington Post*, *NBC News*, *USA Today*, and *Wall Street Journal*. These three datasets are referred to as *dataset 1*, *dataset 2*, and *dataset 3* throughout the rest of this paper. The *dataset 1*, *dataset 2*, and *dataset 3* cover news stories from February 2016–April 2016, May 2016–July 2016, and August 2016–November 2016, respectively. Rather than merging these datasets into a single dataset, this study intends to separate these datasets since the previous analysis discovers that fake and real news stories exhibit different characteristics over a long time period<sup>[12]</sup>. For each fake or real news article, the data includes the title of the story, the Web URL of the news story, the publisher of the news, and the total engagement, measured by the total number of shares, likes, comments, and other reactions of the news received on Facebook.

## 3 Characterizing Fake and Real News

In this section, we study a variety of subjective features on the publishers of real and fake news such as the registration behaviors of publishers' Web sites, the sites ages of the publishers, and the probability of the news disappearance on the Internet.

### 3.1 Web site registration behavior of the publishers

The real or fake news publishers typically have to go through the domain registration process, which allows anonymous domain registrar to serve as a proxy for publishers who prefer to hide their identities. If a publisher chooses to remain anonymous, the Internet whois database will show the proxy, e.g., Domains By Proxy, Limited Liability Company (LLC) as the registration organization. Most popular and well known newspapers typically choose to use the real organization name during the registration process. For example, the registration organization for *wsj.com* is Dow Jones & Company, Inc, which owns Wall Street Journal

newspaper.

Our findings show that the majority of the fake news publishers register their Web sites via proxy services to remain anonymous, while all the real news publisher use their real identifies during the domain registration process. As shown in Table 1, over 78% of the domains publishing fake news are registered via proxy services to hide their true identities of the domain owners, while less than 2% of the domains publishing real news are registered in such a fashion. Thus we believe such patterns can become a powerful feature for machine learning models to distinguish fake and real news.

### 3.2 Internet site ages of the publishers

Beside the domain registration behavior, we also study the ages of the domains for the fake and real news in three datasets. For each dataset, we characterize the domain age distribution for the fake and real news,

**Table 1 Domain registration with proxy service for hiding domain owners' identify.** (%)


Category	Dataset 1	Dataset 2	Dataset 3	Average
Fake news	90	65	80	78.3
Real news	5	0	0	1.7

respectively. As illustrated in Fig. 1, all datasets exhibit consistent observations which reveal the very short domain ages for fake news, and the long domain ages for real news. This result is not surprising in that the credible newspapers registered their domains in early 1990s when Internet and Web start to attract attentions, while the fake news driven publishers often temporarily register the sites for the purpose of spreading fake news in a very short time of period.


### 3.3 Domain rankings

In this study, we also measure the domain rankings of Web sites which publish fake or real news in the datasets via the widely-used Alexa top 1 million sites<sup>[13]</sup>. If a Web site is not included in the top 1 million list, we assign the site a ranking of 1 000 001. As shown in Fig. 2, the credible Web sites publishing real news have much higher domain rankings than sites publishing fake news across all three datasets.

The medium ranking among all Web sites publishing fake news is 987 298, while the medium ranking among all Web sites publishing real news is 158. This observation confirms that the Web sites with influential real news are typically very credible and top sites, while



**Fig. 1 Web site age distribution of the fake news publishers vs. real news publishers.**



**Fig. 2 Domain rankings of the fake and real news publishers.**

those with fake news are likely unknown sites with low or no domain rankings.

### 3.4 Domain popularity

A key metric of the domain popularity on Alexa and other domain reputation measurement services is the average daily page view per visitor. Intuitively a popular site has a higher daily page view per visitor, since the visitor tends to spend more time on browsing the site and read content from different pages. As shown in Fig. 3, the credible Web sites publishing real news have much higher domain popularity than sites publishing fake news across all three datasets.

The medium page views per visitor among all Web sites publishing fake news is 1.0, while the medium page views per visitor among all Web sites publishing real news is 1.86. As a reference data point, the free and collaborative Internet encyclopedia, i.e., wikipedia.org has an average daily page view per visitor of 3.11 as of September 2018. Thus, these results show that the Web sites with influential real news are typically very popular sites, while in general those with fake news have much less page views. However, two Web sites publishing fake news, i.e., breitbart.com and donaldrumpnews.co, have surprising higher page views per visitor than most of the sites publishing real or fake news, which suggests that a few Web sites publishing fake news are very successful in attracting Internet users, as evidenced by substantial daily page views per visitor.

### 3.5 Probability of news disappearance

Credible news agency tends to maintain high quality sites that keep the published news for a long time. However, fake news sites often take the news offline after achieving the short-term goals of misleading the readers. Our analyses on the fake and real news corpus

confirm such common practice.

As shown in Table 2, the three datasets of fake news corpus exhibit consistent news disappearing patterns, while the real news corpus has zero news that is taken offline. Thus we believe news disappearance could become a valuable feature for differentiating or modeling fake and real news. On the other hand, this feature has limited value for distinguishing the latest and emerging fake news, since this feature is derived once the news is taken offline after a certain time period.

In summary, our preliminary results on these popular fake and real news reveal substantial difference between fake and real news on the quality of the news pages, as well as the reputations of the publishing domains reflected by domain ages, domain rankings, domain popularity, and the interesting usage of the registration proxies.

## 4 Topics and Content of Fake and Real News


In this section, we first identify the most important topics of each fake or real news article via tf-idf analysis<sup>[14]</sup>. Subsequently, we explore the probabilistic LDA topic model to understand the difference or similarity of topics between labeled fake and real news.

### 4.1 Important topics identifications via tf-idf analysis

In information retrieval and text mining, tf-idf is a widely used statistical technique for measuring the importance of a term  $t$  in a given document  $d$  from a document corpus  $D$ . The tf-idf value of a term  $t$

**Table 2** Page not found due to news disappearing. (%)

Category	Dataset 1	Dataset 2	Dataset 3	Average
Fake news	40	70	55	55
Real news	0	0	0	0



**Fig. 3** Average domain page views per visitor of the fake and real news publishers.

in a given document  $d$  is a product of term frequency (tf), i.e.,  $\text{tf}(t, d)$ , and inverse document frequency (idf),  $\text{idf}(t, D)$ , which quantify the frequency of the term  $t$  in the document  $d$  and measure the commonness or rarity of the term  $t$  across all documents in the corpus  $D$ , respectively.

The term frequency  $\text{tf}(t, d)$  measures the frequency of the term  $t$  appearing in the document  $d$ , reflecting the importance and relevance of the term to the document. Let  $f_{t,d}$  denote the absolute frequency of the term  $t$  in the document  $d$ , then  $\text{tf}(t, d)$  is calculated as

$$\text{tf}(t, d) = \frac{f_{t,d}}{w_d} \quad (1)$$

where  $w_d$  represents the total number of terms in the document  $d$ .

To measure the commonness and rarity of a term  $t$  across all documents in  $D$ ,  $\text{idf}(t, D)$  is calculated as the inverse fraction of the number of documents containing the term  $t$ , i.e.,  $|d \in D, t \in d|$ , over the total number of document  $|D|$  on a logarithmical scale, i.e.,

$$\text{idf}(t, D) = \frac{|D|}{|d \in D, t \in d|} \quad (2)$$

Thus the tf-idf value of a term  $t$  in the document  $d$ ,  $\text{tf-idf}(t, d, D)$ , then becomes

$$\text{tf-idf}(t, d, D) = \text{tf}(t, d) \cdot \text{idf}(t, D) \quad (3)$$

In this study, we identify the most important and relevant terms from each document via calculating and sorting tf-idf values of all terms in a non-increasing order. For each document, we select the top 20 terms, and subsequently extract the top 10 terms across all documents based on the number of documents the terms appearing as the top terms. Table 3 shows that the most frequent 10 terms extracted from fake and real news corpora, which aggregate all fake and real news across three datasets, as well as the hybrid fake and real news corpus. As shown in Table 3, these terms are very

**Table 3 Most frequent terms ranked by tf-idf values in fake, real, and hybrid news corpora.**

Fake news corpus	Real news corpus	Hybrid news corpus
violence	Trump	Comey
trade	nation	transgender
Palin	Melania	Putin
nuclear	intelligence	fraud
Mexico	FBI	Obama
isis	corrupt	nuclear
goods	Conway	corrupt
country	conservative	Melania
Comey	Hillary	isis
Canada	wikileaks	Trump

similar, thus relying on these terms alone is inefficient for detecting or distinguishing fake news.

## 4.2 Latent Dirichlet allocation topic modeling

Topic models are widely used for understanding the content of documents based on word usage. For example, LDA<sup>[15]</sup>, a generative probabilistic topic model, represents documents in a corpus as a random mixture of latent topics. Each latent topic in the LDA model is characterized by a probability distribution over a vocabulary of words or terms extracted from all documents in the corpus.

In this paper, we explore LDA topic modeling to capture the topics of fake, real, and hybrid news corpus, respectively. The goal of LDA topic modeling on fake and real news is to understand the difference or similarity of topics between labeled fake and real news.

For each latent topic, we measure the topic quality with the widely used coherence score<sup>[16]</sup> to characterize as follows:

$$\text{coherence}(\mathcal{W}_t) = \sum_{\{w_i, w_j\} \in \mathcal{W}_t} \text{score}(w_i, w_j) \quad (4)$$

where  $\mathcal{W}_t$  represents all the words included in the latent topic  $t$ , and  $w_i, w_j \in \mathcal{W}_t$ .

The coherence score,  $\text{score}(w_i, w_j)$ , between two words  $w_i$  and  $w_j$  is defined as

$$\text{score}(w_i, w_j) = \log \frac{\mathcal{D}(w_i, w_j) + \epsilon}{\mathcal{D}(w_j)} \quad (5)$$

where  $\mathcal{D}(w_i, w_j)$  is the size of documents with both words  $w_i$  and  $w_j$  and  $\mathcal{D}(w_j)$  is the size of documents with  $w_j$ . Considering the scenarios of  $\mathcal{D}(w_i, w_j) = 0$ , we set  $\epsilon$  to 1.

Tables 4–6 illustrate the 5 most frequent terms for the top three latent topics with the highest coherence scores for each corpus. As shown in these tables, the fake and real news share strong similarity in the overall topics, thus the LDA topic model alone is not an effective approach to detect or differentiate fake or real news in the real world.

## 5 Document Similarity Analysis for News Predictions

As the LDA topics are inefficient in distinguishing

**Table 4 LDA topics for fake news corpus.**

Top 1	Top 2	Top 3
Trump	president	candidate
Clinton	news	black
Comey	state	will
Hillary	American	one
Donald	time	said

**Table 5** LDA topics for real news corpus.

Top 1	Top 2	Top 3
Trump	Facebook	republican
Clinton	Romney	democratic
Donald	people	authoritarian
president	source	politician
people	see	party

**Table 6** LDA topics for hybrid fake and real news corpora.

Top 1	Top 2	Top 3
people	Trump	Trump
authoritarian	Clinton	Donald
politician	republican	people
party	democratic	make
American	president	will

fake and real news, our followup analysis to explore document similarity between fake, real, or hybrid news articles. First, we randomly divide the labeled fake and real news into training sets and test sets with a split ratio of 67% for training and 33% for test.


For each fake or real news  $n$  in the test corpus, we measure the document similarity between  $n$  and every news in the fake news training set  $\mathcal{F}$  and the real news training set  $\mathcal{R}$ . In particular, we calculate Jaccard similarity  $J(\text{doc}_1, \text{doc}_2)$ , a widely used similarity measure between two documents  $\text{doc}_1$  and  $\text{doc}_2$  with the following equation:

$$J(\text{doc}_1, \text{doc}_2) = \frac{\text{doc}_1 \cap \text{doc}_2}{\text{doc}_1 \cup \text{doc}_2} \quad (6)$$

where  $\text{doc}_1$  and  $\text{doc}_2$  are represented with the vectors, typically sparse, of terms in the documents.

Figure 4a shows the fake news in the test set has a much higher average document similarity with the news in the fake news training set  $\mathcal{F}$  than with those in  $\mathcal{R}$ . However, Fig. 4b shows the real news in the test set has surprising similar document similarity with the news in the real news training set  $\mathcal{R}$  and with those in  $\mathcal{F}$ . Thus as shown in Fig. 4a, document similarity can potentially detect fake news. One of our future work is to systematically quantify the precision and recall of detecting both fake and real news in a large-scale news corpus. Although the analysis in this paper relies on the static and offline training datasets for training and evaluation, we believe our proposed prediction method can be extended to real-time model training and classifications with dynamic training datasets which continuously include the latest news stories with fake or real labels.

In summary, our preliminary analysis on the topics


**Fig. 4** Prediction on fake and real news based on labeled fake and real news corpus.

and content of fake and real news reveals that it is very challenging to simply exploring the tf-idf and LDA topic modeling to effectively detecting fake news. Our study also shows the promising aspect of leveraging document similarity to distinguish fake and real news by measuring the document similarity of the news under tests with the known fake and real news corpus.

## 6 Related Work

In recent years, several algorithms<sup>[1,2,8,17-24]</sup> have been proposed to detect the dissemination of information, misinformation, or fake news. For example, Ref. [2] exploits the diffusion patterns of information to automatically classify and detect misinformation, hoaxes, or fake news, while Ref. [8] proposes linguistic approaches, network approaches, and a hybrid approach combining linguistic cues and network-based behavior insights for identifying fake news. In addition, Ref. [1] reviews the data mining literature on characterizing and detecting fake news on social media.

Similarly, Ref. [18] proposes an SVM-based algorithm for predicting misleading news with

predictive features such as absurdity, grammar, punctuation, humor, and negative affect, and Ref. [19] uses logistic regression to distinguish credible news from fake news based on n-gram linguistic, embedding, capitalization, punctuation, pronoun use, and sentiment polarity features. A recent effort in Ref. [20] formulates the fake news mitigation as the problem of optimal point process intervention in a network, and combines reinforcement learning with a point process network activity model for mitigating fake news in social networks.

In addition, Ref. [25] classifies the task of fake news detection into three different types: serious fabrications, large-scale hoaxes, and humorous fakes, and discusses the challenges of detecting each type of fake news. To address the lack of labeled datasets for fake news detection, Ref. [26] introduces a real-world dataset consisting of 12 836 statements with real or fake labels. In Ref. [21], the authors located the hidden paid posters who get paid for posting fake news via modeling the behavioral patterns of paid posters.

## 7 Conclusion and Future Work

As fake news and disinformation continue to grow in online social media, it becomes imperative to gain in-depth understanding on the characteristics of fake and real news articles for better detecting and filtering fake news. Towards effectively combating fake news, this paper characterizes hundreds of very popular fake and real news from a variety of perspectives including the domains and reputations of the news publishers, as well as the important terms of each news and their word embeddings. Our analysis shows that the fake and real news exhibit substantial differences on the reputations and domain characteristics of the news publishers. On the other hands, the difference on the topics and word embeddings shows little or subtle difference between fake and real news. Our future work is centered on exploring the word2vec algorithm<sup>[27]</sup>, a computationally-efficient predictive model based on neural networks for learning the representations of words in the high-dimensional vector space, to learn word embedding of the important words or terms discovered via the aforementioned tf-idf analysis. Rather than comparing the few important words of each new article, word2vec will allow us to compare the entire vector and embeddings of each word for broadly capturing the similarity and dissimilarity of the content

in the fake or real news.


## Acknowledgment

This work was supported in part by National Science Foundation (NSF) Algorithms for Threat Detection (ATD) Program (No. DMS #1737861) and NSF Computer and Network Systems (CNS) Program (No. CNS #1816995).

## References

- [1] K. Shu, A. Sliva, S. H. Wang, J. L. Tang, and H. Liu, Fake news detection on social media: A data mining perspective, *ACM SIGKDD Explorati.*, vol. 19, no. 1, pp. 22–36, 2017.
- [2] E. Tacchini, G. Ballarin, M. L. D. Vedova, S. Moret, and L. de Alfaro, Some like it hoax: Automated fake news detection in social networks, Tech. Rep. UCSC-SOE-17-05, School of Engineering, University of California, Santa Cruz, CA, USA, 2017.
- [3] M. M. Waldrop, News feature: The genuine problem of fake news, *Proc. Natl. Acad. Sci. USA*, vol. 114, no. 48, pp. 12631–12634, 2017.
- [4] Z. B. He, Z. P. Cai, and X. M. Wang, Modeling propagation dynamics and developing optimized countermeasures for rumor spreading in online social networks, in *Proc. IEEE 35<sup>th</sup> Int. Conf. on Distributed Computing Systems*, Columbus, OH, USA, 2015.
- [5] Z. B. He, Z. P. Cai, J. G. Yu, X. M. Wang, Y. C. Sun, and Y. S. Li, Cost-efficient strategies for restraining rumor spreading in mobile social networks, *IEEE Trans. Veh. Technol.*, vol. 66, no. 3, pp. 2789–2800, 2017.
- [6] C. C. Shao, G. L. Ciampaglia, O. Varol, A. Flammini, and F. Menczer, The spread of fake news by social bots, arXiv preprint arXiv: 1707.07592v1, 2017.
- [7] J. Thorne, M. J. Chen, G. Myrianthous, J. S. Pu, X. X. Wang, and A. Vlachos, Fake news detection using stacked ensemble of classifiers, in *Proc. EMNLP Workshop on Natural Language Processing Meets Journalism*, Copenhagen, Denmark, 2017.
- [8] N. J. Conroy, V. L. Rubin, and Y. M. Chen, Automatic deception detection: Methods for finding fake news, *Proc. Assoc. Inf. Sci. Technol.*, vol. 52, no. 1, pp. 1–4, 2015.
- [9] B. Markines, C. Cattuto, and F. Menczer, Social spam detection, in *Proc. 5<sup>th</sup> Int. Workshop on Adversarial Information Retrieval on the Web*, Madrid, Spain, 2009.
- [10] Y. M. Chen, N. J. Conroy, and V. L. Rubin, Misleading online content: Recognizing clickbait as “false news”, in *Proc. 2015 ACM on Workshop on Multimodal Deception Detection*, Seattle, WA, USA, 2015.
- [11] N. Ruchansky, S. Seo, and Y. Liu, CSI: A hybrid deep model for fake news detection, in *Proc. 2017 ACM on Conf. on Information and Knowledge Management*, Singapore, 2017.
- [12] BuzzFeed News, This analysis shows how viral fake election news stories outperformed real news On Facebook, <https://www.buzzfeednews.com/article/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook>, 2016.

- [13] Alexa, <https://www.alex.com/topsites>.
- [14] G. Salton and M. J. McGill, *Introduction to Modern Information Retrieval*. New York, NY, USA: McGraw-Hill, Inc, 1986.
- [15] D. M. Blei, A. Y. Ng, and M. I. Jordan, Latent Dirichlet allocation, *J. Mach. Learn. Res.*, vol. 3, no. 3, pp. 993–1022, 2003.
- [16] K. Stevens, P. Kegelmeyer, D. Andrzejewski, and D. Buttler, Exploring topic coherence over many models and many topics, in *Proc. Joint Conf. on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, Jeju Island, Korea, 2012.
- [17] F. Wang, H. Y. Wang, K. Xu, J. H. Wu, and X. H. Jia, Characterizing information diffusion in online social networks with linear diffusive model, in *Proc. IEEE Int. Conf. on Distributed Computing Systems*, Philadelphia, PA, USA, 2013.
- [18] V. L. Rubin, N. J. Conroy, Y. M. Chen, and S. Cornwell, Fake news or truth? Using satirical cues to detect potentially misleading news, in *Proc. Ann. Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, San Diego, CA, USA, 2016.
- [19] M. Hardalov, I. Koychev, and P. Nakov, In search of credible news, in *Proc. 17<sup>th</sup> Int. Conf. on Artificial Intelligence: Methodology, Systems, and Applications*, Varna, Bulgaria, 2016.
- [20] M. Farajtabar, J. C. Yang, X. J. Ye, H. Xu, R. Trivedi, E. Khalil, S. Li, L. Song, and H. Y. Zha, Fake news mitigation via point process based intervention, in *Proc. 34<sup>th</sup> Int. Conf. on Machine Learning*, Sydney, Australia, 2017.
- [21] C. Chen, K. Wu, V. Srinivasan, and X. D. Zhang, Battling the internet water army: Detection of hidden paid posters, in *Proc. IEEE/ACM Int. Conf. on Advances in Social Networks Analysis and Mining*, Niagara, Canada, 2013.
- [22] F. Wang, K. Orton, P. Wagenseller, and K. Xu, Towards understanding community interests with topic modeling, *IEEE Acc.*, vol. 6, pp. 24660–24668, 2018.
- [23] K. Lei, Y. Liu, S. R. Zhong, Y. B. Liu, K. Xu, Y. Shen, and M. Yang, Understanding user behavior in Sina Weibo online social network: A community approach, *IEEE Acc.*, vol. 6, pp. 13302–13316, 2018.
- [24] K. Xu, F. Wang, H. Y. Wang, and B. Yang, A first step towards combating fake news over online social media, in *Proc. 13<sup>th</sup> Int. Conf. on Wireless Algorithms, Systems, and Applications*, Tianjin, China, 2018.
- [25] V. L. Rubin, Y. M. Chen, and N. J. Conroy, Deception detection for news: Three types of fakes, in *Proc. 78<sup>th</sup> ASIS&T Ann. Meeting: Information Science with Impact: Research in and for the Community*, Saint Louis, MO, USA, 2015.
- [26] W. Y. Wang, “Liar, liar pants on fire”: A new benchmark dataset for fake news detection, in *Proc. 55<sup>th</sup> Ann. Meeting of the Association for Computational Linguistics*, Vancouver, Canada, 2017.
- [27] T. Mikolov, K. Chen, G. S. Corrado, and J. Dean, Efficient estimation of word representations in vector space, in *Proc. Int. Conf. on Learning Representations*, Scottsdale, AZ, USA, 2013.



**Kuai Xu** is an associate professor at Arizona State University. He received the BS and MS degrees from Peking University, China in 1998 and 2001, respectively, and the PhD degree from the University of Minnesota in 2006. His research interests include network security, Internet measurement, big data, data mining, and machine learning.



**Feng Wang** is an associate professor of Applied Computing Program at School of Mathematical and Natural Sciences in Arizona State University. She received the PhD degree from University of Minnesota Twin Cities in 2005, the MS degree from Peking University in 1999, and the BS degree from Wuhan University in 1996. She is interested in graph theory, algorithm design and optimization, social network analysis, and network security.



**Haiyan Wang** is a professor at Arizona State University. He received the BS degree from Northwest Normal University, China in 1985, the MS degree from Ocean University of China in 1988, and the PhD degree from Michigan State University in 1997. His research interests include differential equations, mathematical biology, online social networks, machine learning, and data science.



**Bo Yang** is an associate professor at Jiangxi University of Finance and Economics, China. His research interests include cryptography and network security. He received the BS degree from Huazhong University of Science and Technology in 1996, the MS degree from Nanchang University in 2002, and the PhD degree from Jiangxi University of Finance and Economics in 2012.