Proceedings of the 2021 Pre-ICIS SIGDSA Symposium
Special Interest Group on Decision Support and Analytics (SIGDSA)

12-2021

Exploring Roles of Emotion in Fake News Detection

Kenan Xiao Auburn University

Ashish Gupta

Wenting Jiang

Xiao Qin

Exploring Roles of Emotion in Fake News Detection

Research-in-Progress

Kenan Xiao
Auburn University
kzx0010@auburn.edu

Ashish Gupta
Auburn University
azg0074@auburn.edu

Wenting Jiang
Auburn University
wzj0027@auburn.edu

Xiao Qin
Auburn University
xqin@auburn.edu

Abstract

Detecting fake news is becoming widely acknowledged as a critical activity with significant implications for social impact. As fake news tends to evoke high-activating emotions from audiences, the role of emotions in identifying fake news is still under-explored. Existing research made efforts in examining effective representations of emotions conveyed in the news content to help discern the veracity of the news. However, the aroused emotions from the audience are usually ignored. This paper first demonstrates effective representations of emotions within both news content and users' comments. Furthermore, we propose an emotion-aware fake news detection framework that seamlessly incorporates emotion features to enhance the accuracy of identifying fake news. Future work will include thorough experiments to prove that the proposed framework with the emotions expressed in news and users' comments improves fake news detection performance.

Keywords

Fake news detection, misinformation, emotion, deep learning.

Introduction

Nowadays, there is significant evidence that fake news can negatively affect both individuals and society. To begin with, individuals who are deceived by fake news form misguided opinions. Second, fake news is designed to influence how people react to legitimate news. Thirdly, widespread dissemination of fake news has potentials to destroy the legitimacy of the entire news ecosystem. Consequently, identifying fake news on social media is crucial and time-consuming.

Recent research has explored a variety of methods to ramp up fake news detection techniques. A primary approach is to design classifiers for discerning news as real or fake news. A majority of existing work on fake news detection relies on news content (Castillo et al. 2011; Pérez-Rosas et al. 2017; Qazvinian et al. 2011), social context information (Guo et al. 2018; Ma et al. 2016; Ma et al. 2018), or multimodal information such as images (Jin et al. 2016; Qi et al. 2019). The existing research has advanced the fake news detection landscape to such a high level; however, the role of emotion conveyed in news content and social media users' comments are under-explored.

Tracing back to why people fall for fake news needs extensive efforts from cognitive science and psychology communities. Pennycook and Rand have examined the psychology of fake news (Pennycook and Rand 2021), especially why people fall for misinformation. They enlightened that fake news headlines and contents are often emotionally evocative (Pennycook and Rand 2021). This finding is important because the audience experiencing more emotion initially tends to believe false news. More times than not, fake news guides the audience to judge the veracity of news by emotions (Martel et al. 2020). Therefore, examining the role of emotion amid fake news detection must be placed under spotlight to optimize the fake news detection performance.

Examining the role of emotion in fake news detection is still in its infancy. Dey et al. (Dey et al. 2018) investigated 200 tweets related to the 2016 US Presidential Election and demonstrated that credible tweets were predominantly positive or neutral in sentiment, whereas fake tweets exhibited a greater inclination toward the negative sentiment. Ajao et al. (Ajao et al. 2019) proved that the veracity of news is related to the sentiment of news content, thereby creating a sentimental feature to help classifiers identify fake news. Cui et al. (Cui et al. 2019) first demonstrated that the sentiment polarity of social users' comments of fake news was greater than real ones. The users' latent sentiments feature was incorporated into a deep embedding framework to detect fake news. Zhang et al. (Zhang et al. 2021) explored the dual emotions aroused in news content and users' responses and designed compelling emotional features that are deployed to arbitrary fake news detectors. Despite the promising results from the preceding research, the role of emotion in fake news detection has challenging and intriguing open issues to be explored.

In order to investigate the role of emotion in fake news detection, we posit the following research questions that guide the design of our study.

- **RQ 1:** What are the effective representations of emotions in the fake news or misinformation literature?
- **RQ2:** Do the emotion representations have distinct patterns between fake news and real ones?
- **RQ3:** Does the fake news arouse the emotion of the audience different from the real news?

To explore the aforementioned research questions, we first target effective representations of emotion expressed in the text of news and users' responses. We scrutinize three indicative features, namely, emotion distribution, emotion intensity, and overall sentiments. Furthermore, we propose an emotion-aware fake news detection framework, which seamlessly incorporates the emotion features. Lastly, we will conduct experiments on the real-world dataset to demonstrate the effectiveness of these emotion features in identifying fake news and help answer the posited research questions.

Related Work

Comprehending the presence of emotions and its effectiveness in fake news detection has not yet received much attention. In this section, we review two folds of research that are highly relevant to our task: general fake news detection and emotion analysis in fake news.

Fake News Detection

The definition of "fake news" has been extensively studied; however, there is no universally accepted definition. The concept of "fake news" is inextricably related to "false news", "rumor", "misinformation", etc. (Pierri and Ceri 2019). Early attempts in fake news detection are highly related to the area of information credibility assessment. More recent research in fake news detection has advanced into two types of approaches: content-based approaches and social context-based approaches.

The content-based approach seeks to classify news according to the content of the information to be verified. The cue and feature-based approach aims to identify indicative cues and features that can guide fake news detection (Driscoll 1994; Zhao et al. 2015; Zhou et al. 2004). Driscoll (Driscoll 1994) utilized a scientific content analysis (SCAN) scheme that included cues for detecting deception. Zhou et al. (Zhou et al. 2004) constructed a cue set consisting of 14 linguistic cues that are effective at detecting deception. Zhao et al. (Zhao et al. 2015) provided a range of regular expressions for capturing inquiry and correction patterns in social media posts. More recent research considers linguistic analysis to detect fake news as such techniques does not require task-specific cue sets and features. These studies also employ N-grams and Part-of-Speech (POS) tags to extract the linguistic patterns of fake news (Ott et al. 2011). Feng et al. (Feng et al. 2012) investigate the use of Probabilistic Context-Free Grammars (PCFG) to encode more detailed syntactic features for deception detection. Zhang et al. (Zhang et al. 2019) proposed a two-phase analytical approach for detecting fake news, namely fake topics detection and fake events detection. Recent research in fake news detection has been proposed to utilize the power of deep learning models, such as CNN-based models and RNN-based models (Ma et al. 2016; Yu et al. 2017).

The social context-based approach identifies fake news by leveraging rich secondary data from user responses, user characteristics, and news propagation patterns on social media. Castillo et al. (Castillo et al. 2011) crafted a feature set that included user-based, text-based, and propagation-based features and classified fake news by a decision tree model. The propagation patterns of fake news on social media have been well explored. Wu et al. (Wu et al. 2015) explored the use of random walk graph kernel over propagation trees to detect fake news. Ma et al. (Ma et al. 2017) considered the propagation patterns by evaluating the similarity between propagation trees using tree kernels. Another recent study conducted with the user responses perspective extracted the user responses feature using doc2vec word embeddings (Ruchansky et al. 2017). Chen et al. (Chen et al. 2018) collected textual information of user responses by LSTM architecture coupled with an attention mechanism.

Emotion Analysis in Fake News Detection

Limited existing research has explored the role of emotions in fake news detection. Ajao et al. (Ajao et al. 2019) provide evidence that the veracity of news is related to the sentiment of news content and create a sentimental feature to help classifiers identify the fake news. Guo et al. (Guo et al. 2019) modeled the emotion signals by emotional category, emotional intensity, and emotional expression and proposed an end-to-end emotion-based fake news detection framework based on Bi-GRU. Giachanou et al. (Giachanou et al. 2019) proposed an EmoCred framework that modeled the emotional signals from emotional lexicons, emotional intensity, and emotional reactions. Cui et al. (Cui et al. 2019) designed the users' latent sentiments feature and incorporated it into a deep embedding framework to detect fake news. Zhang et al. (Zhang et al. 2021) explored the dual emotions aroused in news content and users' responses and designed compelling emotional features that can work on arbitrary fake news detectors. Anoop et al. (Anoop et al. 2020) designed a principal way to alter the news content to emotionalized news content and applied multiple classifiers to identify fake news. A recent study analyzed different types of emotions such as anger, joy, fear and sadness in the context of COVID-19 infodemic (Gupta et al. 2021). Despite the promising results from the research above, most work fails to prove the effectiveness of the emotion representations and distinctiveness between fake news and real ones. A detailed exploration of the role of emotion is much needed in fake news detection.

Methodologies

We first introduce the emotion representations of the text information within news content and users' comments from three perspectives: emotion distribution, emotion intensity, and the overall sentiment. We then demonstrate how these representations can be integrated into an emotion-aware fake news detection framework to help improve fake news detection.

Emotion Representations

To comprehend the emotional aspects of news and their impact on social media users, we propose to model the emotions contained in the news itself as well as emotions aroused among social media users as result of consuming fake news. We construct our emotion representations for the aforementioned two emotion sources from various aspects of emotions, including emotion lexicons, emotional intensity, and the sentiments of text information.

Emotion Distribution

The emotions conveyed by a piece of text are often leveraged by several emotion-indicating words that are annotated in the emotion lexicons such as EmoLex (Mohammad and Turney 2013). Each emotion-indicating word associates with one specific emotion. For example, "sad" expresses sadness, whereas the word "angry" indicates anger. By examining these words throughout the news content, we can extract the emotion distribution of the news content. The distribution can serve as one of the effective representations of the emotions contained in the news contents. We refer to the "NRC Word-Emotion Association Lexicon", also known as "EmoLex", as our emotional lexicons (Mohammad and Turney 2013). EmoLex is a collection of English words associated with eight fundamental emotions, namely, anger, fear, anticipation, trust, surprise, sadness, joy, and disgust.

Emotion Intensity

We presented the distinctive emotion distribution representation within the text in the previous subsection. However, the intensity of the emotion words is not described in the previous description. Words in the same emotion category may express different intensities (Guo et al. 2019; Zhang et al. 2021). For example, the word "furious" is much stronger than the word "angry" when describing the "anger' emotion. In order to capture the full characteristics of emotions, we propose to model the emotion intensity feature. The extraction of emotion intensity feature is similar to the emotion distribution extraction. The assumptions still hold when extracting the emotion intensity feature.

Overall Sentiment

Fake news often arouses the high-activating emotions of social media users. Thus, the overall sentiment of the news and social users' comments is yet another effective emotion-related feature to help identify fake news. The overall sentiment of a piece of text is a value between -1 and 1. As the sentiment value goes to -1, the text presents an overall stronger negative sentiment whereas the sentiment value goes to 1, it indicates a stronger positive sentiment within the text. Available public toolkits, such as VADER (Hutto and Gilbert 2014), can calculate the overall sentiment of the text.

Fake News Detection Framework

In this subsection, we present how the emotion representation can be integrated into a fake news detection framework. To test the power of emotion representations, we did not choose complex models. Instead, we select the model that are suitable for capturing the text information conveyed in news and audience's responses. Standard LSTMs are capable of recursively processing input sequences of variable length and capturing long-term dependencies. We employ the Bi-directional long short-term memory (Bi-LSTM), a variant of LSTM, as the basis of our framework. Bi-LSTM is capable of capturing bidirectional semantic dependencies and is suitable for grasping contextual information. For example, one may write a "I love movies. Not at all!". Bi-LSTM can capture the meaning that the person does not love movie at all from the backwards direction.

Emotion Feature Extraction

To empower the full capability of emotion representations, we extract the emotion representations from three aspects: the headline of the news, the news content, and social media users' responses. Pennycook and Rand demonstrated in their recent studies that fake news headlines are often emotionally evocative (Pennycook and Rand 2021). Therefore, it is worthy to mine the emotion expressed in the news headlines.

Emotion-Aware Fake News Detection Framework

The emotion-aware fake news detection framework consists of two collaborative modules, namely, news content module and emotion module. The details of the framework are shown in Figure 1. The news content module aims to capture the textual information within the news content, while the emotion module extracts the emotional signals conveyed in the news headlines, news content, and social media users' comments. We aggregate those two sources of information to jointly help identify the veracity of the news.

News Content Module: The news content module can capture the textual information conveyed in the news content. We apply the start-of-the-art Bidirectional Encoder Representations from Transformer (Bert) to extract word embeddings. Compared to word2vec, Bert word embeddings has the advantage of considering contextual information expressed in the text. After we extract word embeddings, we employ Bi-directional long short-term memory (Bi-LSTM) model to help detect fake news. Bi-LSTM is composed of two LSTMs: one that receives input going forward and another that receives input going backward. It is beneficial to examine the textual information from two directions because it captures the bidirectional semantic dependencies. The hidden state h_c is achieved by concatenating the hidden states from two LSTMs.

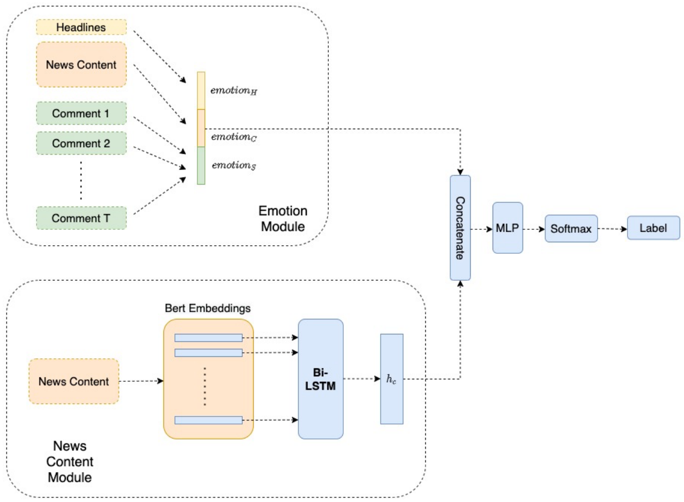

Figure 1. Emotion-Aware Fake news Detection Framework

Emotion Module: The emotion module aims to extract effective representations of emotion conveyed in the news headlines, news content, and news comments on social media. The details of emotion feature extractions are well demonstrated in the above subsection. After we obtain $emotion_H$, $emotion_C$, and $emotion_S$, we concatenate these three vectors as our final emotion feature vector.

The last hidden state of Bi-LSTM h_c and emotion feature vector are concatenated together and fed into a multi-layer perceptron (MLP) layer and a softmax layer for the final prediction of news veracity.

Conclusion and Future Work

This study explores how emotion can be utilized in enhancing fake news detection performance. We model the emotion of text information from three distinctive aspects between fake news and real ones. A novel emotion-aware fake news detection framework is presented, incorporating emotions from news headlines, news content and users' comments. Future work will include collecting a suitably large dataset and conduct a statistical significance test to prove the distinctiveness of emotions between fake news and real news. Moreover, we will compare our proposed framework with the state-of-the-art fake news detectors to demonstrate the advantages of including emotions in fake news detectors.

References

Ajao, O., Bhowmik, D., and Zargari, S. 2019. "Sentiment Aware Fake News Detection on Online Social Networks," *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*: IEEE, pp. 2507-2511.

Anoop, K., Deepak, P., and Lajish, V. 2020. "Emotion Cognizance Improves Health Fake News Identification," *IDEAS*, p. 24th.

Castillo, C., Mendoza, M., and Poblete, B. 2011. "Information Credibility on Twitter," *Proceedings of the 20th international conference on World wide web*, pp. 675-684.

Chen, T., Li, X., Yin, H., and Zhang, J. 2018. "Call Attention to Rumors: Deep Attention Based Recurrent Neural Networks for Early Rumor Detection," *Pacific-Asia conference on knowledge discovery and data mining*: Springer, pp. 40-52.

Cui, L., Wang, S., and Lee, D. 2019. "Same: Sentiment-Aware Multi-Modal Embedding for Detecting Fake News," *Proceedings of the 2019 IEEE/ACM international conference on advances in social networks analysis and mining*, pp. 41-48.

Dey, A., Rafi, R. Z., Parash, S. H., Arko, S. K., and Chakrabarty, A. 2018. "Fake News Pattern Recognition Using Linguistic Analysis," *2018 Joint 7th International Conference on Informatics, Electronics & Vision (ICIEV) and 2018 2nd International Conference on Imaging, Vision & Pattern Recognition (icIVPR)*: IEEE, pp. 305-309.

Driscoll, L. N. 1994. "A Validity Assessment of Written Statements from Suspects in Criminal Investigations Using the Scan Technique," *Police Stud.: Int'l Rev. Police Dev.* (17), p. 77.

Feng, S., Banerjee, R., and Choi, Y. 2012. "Syntactic Stylometry for Deception Detection," *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 171-175.

Giachanou, A., Rosso, P., and Crestani, F. 2019. "Leveraging Emotional Signals for Credibility Detection," *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 877-880.

Gupta, H., Li, H., Farnoush, A., Jiang, W. 2021 Understanding Patterns of COVID Infodemic: A Systematic and Pragmatic Approach to Curb Fake News, Journal of Business Research,2021, forthcoming, https://doi.org/10.1016/j.jbusres.2021.11.032

Guo, C., Cao, J., Zhang, X., Shu, K., and Yu, M. 2019. "Exploiting Emotions for Fake News Detection on Social Media," *arXiv preprint arXiv:1903.01728*).

Guo, H., Cao, J., Zhang, Y., Guo, J., and Li, J. 2018. "Rumor Detection with Hierarchical Social Attention Network," *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pp. 943-951.

Hutto, C., and Gilbert, E. 2014. "Vader: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text," *Proceedings of the International AAAI Conference on Web and Social Media*.

Jin, Z., Cao, J., Zhang, Y., Zhou, J., and Tian, Q. 2016. "Novel Visual and Statistical Image Features for Microblogs News Verification," *IEEE transactions on multimedia* (19:3), pp. 598-608.

Ma, J., Gao, W., Mitra, P., Kwon, S., Jansen, B. J., Wong, K.-F., and Cha, M. 2016. "Detecting Rumors from Microblogs with Recurrent Neural Networks,").

Ma, J., Gao, W., and Wong, K.-F. 2017. "Detect Rumors in Microblog Posts Using Propagation Structure Via Kernel Learning," Association for Computational Linguistics.

Ma, J., Gao, W., and Wong, K.-F. 2018. "Rumor Detection on Twitter with Tree-Structured Recursive Neural Networks," Association for Computational Linguistics.

Martel, C., Pennycook, G., and Rand, D. G. 2020. "Reliance on Emotion Promotes Belief in Fake News," *Cognitive research: principles and implications* (5:1), pp. 1-20.

Mohammad, S. M., and Turney, P. D. 2013. "Crowdsourcing a Word–Emotion Association Lexicon," *Computational intelligence* (29:3), pp. 436-465.

Ott, M., Choi, Y., Cardie, C., and Hancock, J. T. 2011. "Finding Deceptive Opinion Spam by Any Stretch of the Imagination," *arXiv preprint arXiv:1107.4557*).

Pennycook, G., and Rand, D. G. 2021. "The Psychology of Fake News," *Trends in cognitive sciences*).

Pérez-Rosas, V., Kleinberg, B., Lefevre, A., and Mihalcea, R. 2017. "Automatic Detection of Fake News," *arXiv preprint arXiv:1708.07104*).

Pierri, F., and Ceri, S. 2019. "False News on Social Media: A Data-Driven Survey," *ACM Sigmod Record* (48:2), pp. 18-27.

Qazvinian, V., Rosengren, E., Radev, D., and Mei, Q. 2011. "Rumor Has It: Identifying Misinformation in Microblogs," *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pp. 1589-1599.

Qi, P., Cao, J., Yang, T., Guo, J., and Li, J. 2019. "Exploiting Multi-Domain Visual Information for Fake News Detection," *2019 IEEE International Conference on Data Mining (ICDM)*: IEEE, pp. 518-527.

Ruchansky, N., Seo, S., and Liu, Y. 2017. "Csi: A Hybrid Deep Model for Fake News Detection," *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pp. 797-806.

Wu, K., Yang, S., and Zhu, K. Q. 2015. "False Rumors Detection on Sina Weibo by Propagation Structures," *2015 IEEE 31st international conference on data engineering*: IEEE, pp. 651-662.

Yu, F., Liu, Q., Wu, S., Wang, L., and Tan, T. 2017. "A Convolutional Approach for Misinformation Identification," *IJCAI*, pp. 3901-3907.

Zhang, C., Gupta, A., Kauten, C., Deokar, A. V., and Qin, X. 2019. "Detecting Fake News for Reducing Misinformation Risks Using Analytics Approaches," *European Journal of Operational Research* (279:3), pp. 1036-1052.

Zhang, X., Cao, J., Li, X., Sheng, Q., Zhong, L., and Shu, K. 2021. "Mining Dual Emotion for Fake News Detection," *Proceedings of the Web Conference 2021*, pp. 3465-3476.

Zhao, Z., Resnick, P., and Mei, Q. 2015. "Enquiring Minds: Early Detection of Rumors in Social Media from Enquiry Posts," *Proceedings of the 24th international conference on world wide web*, pp. 1395-1405.

Zhou, L., Burgoon, J. K., Nunamaker, J. F., and Twitchell, D. 2004. "Automating Linguistics-Based Cues for Detecting Deception in Text-Based Asynchronous Computer-Mediated Communications," *Group decision and negotiation* (13:1), pp. 81-106.