

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

5-2022

RumorLens: interactive analysis and validation of suspected rumors on social media

Ran WANG

Kehan DU

Qianhe CHEN

Yifei ZHAO

Mojie TANG

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

WANG, Ran; DU, Kehan; CHEN, Qianhe; ZHAO, Yifei; TANG, Mojie; TAO, Hongxi; WANG, Shipan; LI, Yiyao; and Yong WANG. RumorLens: interactive analysis and validation of suspected rumors on social media. (2022). *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems Extended Abstracts, New Orleans, United States, April 29-May 5*. 1-7.

Available at: https://ink.library.smu.edu.sg/sis_research/7685

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Author

Ran WANG, Kehan DU, Qianhe CHEN, Yifei ZHAO, Mojie TANG, Hongxi TAO, Shipan WANG, Yiyao LI, and Yong WANG

RumorLens: Interactive Analysis and Validation of Suspected Rumors on Social Media

Ran Wang

School of Journalism and Information
Communication, Huazhong
University of Science and Technology
China

Kehan Du

School of Computer Science, United
Kingdom University of Bristol
United Kingdom

Qianhe Chen

School of Computer Science and
Technology, Huazhong University of
Science and Technology
China

Yifei Zhao

School of Artificial Intelligence and
Automation, Huazhong University of
Science and Technology
China

Mojie Tang

School of Architecture and Urban
Planning, Huazhong University of
Science and Technology
China

Hongxi Tao

Computer Science Department,
University College London
United Kingdom

Shipan Wang

School of Journalism and Information
Communication, Huazhong
University of Science and Technology
China

Yiyao Li

School of Journalism and Information
Communication, Huazhong
University of Science and Technology
China

Yong Wang

School of Computing and Information
Systems, Singapore Management
University
Singapore

ABSTRACT

With the development of social media, various rumors can be easily spread on the Internet and such rumors can have serious negative effects on society. Thus, it has become a critical task for social media platforms to deal with suspected rumors. However, due to the lack of effective tools, it is often difficult for platform administrators to analyze and validate rumors from a large volume of information on a social media platform efficiently. We have worked closely with social media platform administrators for four months to summarize their requirements of identifying and analyzing rumors, and further proposed an interactive visual analytics system, *RumorLens*, to help them deal with the rumor efficiently and gain an in-depth understanding of the patterns of rumor spreading. *RumorLens* integrates natural language processing (NLP) and other data processing techniques with visualization techniques to facilitate interactive analysis and validation of suspected rumors. We propose well-coordinated visualizations to provide users with three levels of details of suspected rumors: an overview displays both spatial distribution and temporal evolution of suspected rumors; a projection view leverages a metaphor-based glyph to represent each suspected rumor and further enable users to gain a quick understanding of their overall characteristics and similarity with each other; a propagation view visualizes the dynamic spreading details of a suspected rumor with a novel circular visualization design, and facilitates interactive analysis and validation of rumors

in a compact manner. By using a real-world dataset collected from Sina Weibo, one case study with a domain expert is conducted to evaluate *RumorLens*. The results demonstrated the usefulness and effectiveness of our approach.

CCS CONCEPTS

• **Human-centered computing** → **Visualization toolkits.**

KEYWORDS

Suspected Rumor, Social Media, Human-Computer Collaboration, Visualization Design, Location Distribution, Topic Evolution, Feature Projection, Propagation View, Circular Design.

ACM Reference Format:

Ran Wang, Kehan Du, Qianhe Chen, Yifei Zhao, Mojie Tang, Hongxi Tao, Shipan Wang, Yiyao Li, and Yong Wang. 2022. *RumorLens: Interactive Analysis and Validation of Suspected Rumors on Social Media*. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '22 Extended Abstracts)*, April 29-May 5, 2022, New Orleans, LA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3491101.3519712>

1 INTRODUCTION

Social media has been widely used in our everyday life and made information sharing and communication very convenient. However, it also provides an easy and fast way to generate and spread various rumors [33, 34]. Rumors often refer to the information that has already been spread among people, but has not been verified or is intentionally spread to mislead the audience [4, 27, 29, 47]. Such rumors on social media can have serious negative effects on both individuals [14, 16] and the society [3, 8, 24], for example, causing public panic, disrupting the social order, decreasing the government credibility, and even endangering national security [18]. Service providers of social media have tried to identify the rumor spread on social media platforms in the past few years [12].

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.

CHI '22 Extended Abstracts, April 29-May 5, 2022, New Orleans, LA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9156-6/22/04...\$15.00

<https://doi.org/10.1145/3491101.3519712>

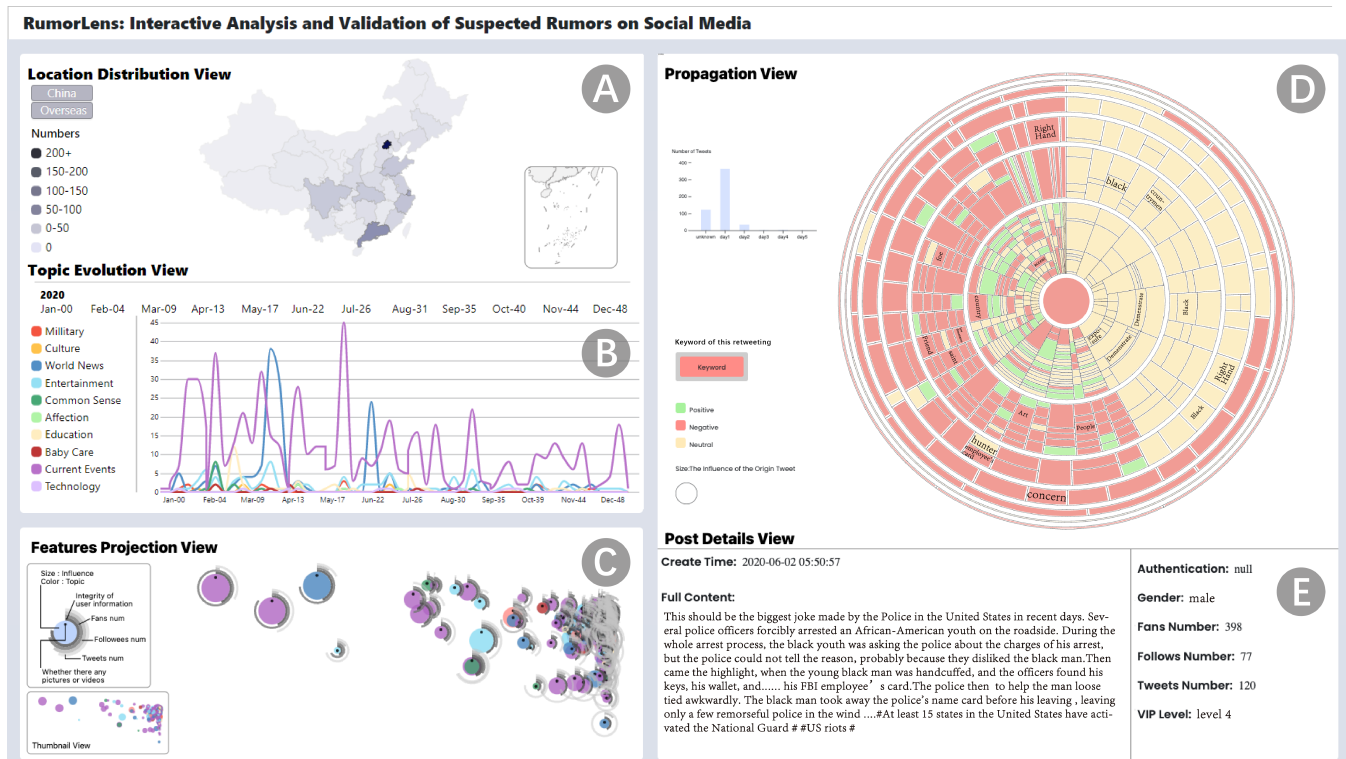


Figure 1: RumorLens, a multi-level visual analytics system to help users analyze and validate suspected rumors on social media in an interactive way. A) Location Distribution View provides a summary of suspected rumors' spatial distribution; B) Topic Evolution View shows changes of suspected rumors with different topics over time; C) Features Projection View reveals overall characteristics of suspected rumors and similarity with each other; D) Propagation View visualizes the dynamic spreading details of a suspected rumor with a novel circular visualization design; E) Post Details View displays the details of user information and post contents.

Conventional approaches for rumor identification rely on the personal experience of content specialists. Though high accuracy can be guaranteed by this kind of manual approach, it is quite time-consuming and labor-intensive to deal with massive of information. In addition, there have been more and more automatic rumor identification methods [13, 22, 26], where rumor identification is defined as a binary classification problem. Such automatic rumor identification methods are often more efficient than manual approaches. However, most existing automatic methods can not ensure accuracy and are often not robust in practical applications [9]. Compared with those automatic methods, the experience of human beings can provide a richer output that includes the reason for the decision [9]. Currently, the approaches for rumor identification based on human-computer collaboration are more preferred [18]. For example, computer algorithms [19] or crowdsourcing techniques are first employed to detect suspected rumors on social media, which will be validated by a team of journalists in a manual way subsequently [49].

These approaches enable platform administrators of social media to improve their efficiency for rumor identification. However, after several rounds of in-depth interviews with domain experts from social media, we identified several major issues and concerns of using

those rumor identification approaches. Above all, suspected rumors have been identified by crowdsourcing or automatic detector at first [9]. Nevertheless, due to the lack of effective tools, it is still a challenge for social media platform administrators to deal with a large number of suspected rumors, especially for those regarding emergencies with considerable impact. Second, it is necessary for platform administrators to gain an in-depth understanding of suspected rumor features in terms of content, user, topic, and propagation, which makes rumor validation solid and convincing. In particular, it is important to trace the dynamic propagation details of a suspected rumor on social media [41], which, however, is difficult to achieve now. In this study, we aim to develop novel visual analytic techniques to address these challenges.

To this end, we propose an interactive visual analytic system, *RumorLens*, to help the administrators of social media platforms to analyze and validate suspected rumors efficiently and gain deep insights into the rumor spreading. In particular, we develop well-coordinated visualization views to provide users with details of the suspected rumors at different levels: (1) an overview of suspected rumors is displayed for filtering suspected rumors rapidly, which visualizes the spatial distribution and temporal evolution of rumors; (2) a projection view is designed for quick selection based

on their overall features and similarity with each other, and each suspected rumor is represented with a metaphor-based glyph; (3) a propagation view is provided to analyze and validate suspected rumors interactively, where we propose a novel circular visualization design to visualize the dynamic spreading details of the selected suspected rumor in a compact manner. Furthermore, *RumorLens* enables linked exploration and smooth interactions across the three major views so that our target users can explore the suspected rumors from multiple perspectives and different levels of detail. To extensively evaluate *RumorLens*, we conducted a case study to demonstrate the usefulness and effectiveness of a dataset collected from Sina Weibo, a popular social media in China. The main contributions of this paper are summarized as follows:

- An interactive visualization system, *RumorLens*, to assist social media platform administrators in analyzing and validating suspected rumors efficiently with multiple levels of details.
- A novel circular design to show the dynamic spreading details of a suspected rumor, facilitating interactive analysis and validation of rumors in a compact manner.
- A case study conducted by a domain expert to demonstrate the usefulness and effectiveness of *RumorLens* in helping users gain deep insights into suspected rumors on social media.

2 RELATED WORK

Our related work can be categorized into three parts: rumor detection in social media, visual analytics of social media, and visualization of information diffusion.

Rumor Detection in Social Media: Depending on whether the human is involved in the process of rumor identification, rumor detection can generally be classified into two types: manual methods and automatic methods. Manual methods include (1) expert-based method and (2) crowd-sourced based method [45]. For the expert-based method, experts who have a deep knowledge of certain domains will manually annotate tweets as ‘true’, ‘false’, or ‘unverified’ [1]. They validate rumors with the help of fact-checking websites [12] and inquiry retweets [47]. For crowd-sourced fact-checking, any user on Sina Weibo can report suspicious rumors, and the final declaration would be made by a group of experts [18]. However, manual methods only rely on human judgment and efforts. Thus, they cannot handle a large volume of data [49]. To address the scalability, automatic rumor detection techniques have been developed, leveraging natural language processing and machine learning techniques [13, 27]. Features are first extracted and then identified by various kinds of classifiers [2, 21, 31, 39, 47, 50]. The key issue focus on integrating multiple rumor features and are committed to selecting the most effective features [11, 12] which are divided into four types [11]: (1) content features [22]; (2) user features, such as fans number, followers number [13]; (3) topic features, which are aggregates computed from content and user features [32]; and (4) propagation features, which consider attributes related to the propagation path [14, 25]. However, none of them can deal with all rumors with different topics and dynamic propagation effectively. Meanwhile, manual labeling is a very challenging

and time-consuming task. In this paper, we propose a visual analyzing approach that combines human domain knowledge and automated machine learning techniques to make a rumor analysis and validation reliable and explainable.

Visual Analytics of Social Media: Prior studies in visual analytics of social media most focus on information collection and behaviors understanding [42]. In the field of information collection, it can be generally summarized into three categories: keyword-based [5, 40], topic-based [7, 46] and multi-faceted approaches [43]. Research studies about behaviors understanding are proposed to help understand user’s behaviors, including collective behaviors [17, 37], cooperation and competition [35, 38]. These approaches enable a multi-aspect analysis of a large collection of tweets through data visualization and rich interactions, such as filtering, zooming, and clustering, etc. However, few studies have been conducted to visually analyze and validate suspected rumors on social media in the tight cooperation of theoretical research and experience experts from social media.

Visualizations of Information Diffusion: Two main types of data visualization techniques have been proposed to represent information diffusion on social media, i.e., network-based (e.g., node-link diagram) and space-filling-based (e.g., Voronoi, treemap) approaches [42]. For network-based methods, nodes and links are employed to visualize a social network, allowing users to understand opinion flow or explore propagation patterns [10, 20, 24]. However, such network-based approaches suffer from scalability issues when the network on social media scales to tens of thousands of nodes. To address this issue, space-filling-based methods have been proposed. Adjacency diagrams encode connectivity between nodes with an adjacency matrix rather than individually drawn edges [23, 28]. Nonetheless, the time series and hierarchical relationship among tweets are usually ignored, which are quite important for rumor validation. Inspired by such visualizations, we propose a novel circular design to show the dynamic spreading of a suspected rumor in a compact manner.

3 DATA ABSTRACTION AND PROCESSING

Sina community management center is an online platform in charge of collecting, recording and dealing with illegal messages including suspected rumors published on Sina Weibo, one of the most famous social media in China. As reported by Sina Weibo, tens of thousands of rumors in various fields have to be dealt with in time every year, or else they may result in a serious impact on both the public and the government [41]. We collected a dataset of suspected rumors between 2019/12/27 and 2020/12/14 from the center, with 936 suspected rumors, approximately 80000 corresponding retweets and comments, and 53843 user profiles.

Followed by suggestions of platform administrators and previous studies mentioned above, multiple features regarding user, content, topic and propagation of suspected rumors are important for rumor analysis and validation. Hence, in addition to the original data of tweets, TF-IDF [30], sentiment recognition [44], topic classification [15], influence calculation [6] are employed to further extract keywords, sentiment, topic and influence of suspected rumors respectively. In particular, we adopt t-SNE [36] to reduce the feature

dimension of suspected rumors and project onto a 2D map to show relationships among them.

4 REQUIREMENTS ANALYSIS

To better understand the major challenges and design requirements for suspected rumor analysis and validation, we have worked closely with three platform administrators (E1-E3) from social media for four months. They were experts from three famous internet companies (Sohu, NetEase, and TouTiao) respectively. E1 is in charge of rumor collection and validation in Sohu for two years. E2 is responsible for content production in NetEase and familiar with features of rumors. E3 is a platform administrator who is mainly checking and publishing content in TouTiao. All the experts have experience in dealing with rumors on the Internet. By conducting a series of interviews and discussions with them, we collected their feedback and summarized the major design requirements **R1-R3** as follows.

R1. Explore the Overall Space-Time Distribution of Suspected Rumors. Based on the feedback of our experts, all of them agreed that it is almost impossible to manually review all the messages within a limited time. According to their work principles, those recent suspected rumors concerning people's livelihood are supposed to be dealt with with higher priorities. Therefore, an effective approach should help easily and quickly filter suspected rumors with priorities, especially those could have negative effects. E1 suggested that spatial-temporal attributes and the topic of the message should be considered.

R2. Inspect Suspected Rumor Cases Through Characteristics Comparisons. All the experts agreed that the characteristics analysis of rumors can provide significant clues for rumor validation. E2 pointed out that the messages whose characteristics are different from others are more likely rumors, especially for those ones with significant influence. In addition, according to E3's work experience, pictures and videos are usually added for rumors to enhance their credibility. Thus, visual inspection of suspected rumors in feature space can help platform administrators select a case with the most suspicious.

R3. Explore the Propagation Details of Individual Suspected Rumors. As E3 suggested, a functionality of exploring the propagation details of a suspected rumor is essential. It can help platform administrators further confirm suspected rumor cases, which may be wrongly labeled by some automated detection methods or reported by users from social media. E1 and E2 agreed that propagation path, contents of tweets and retweets are important factors for rumor validation. For example, rumors with negative sentiment diffused significantly faster and deeper than the truth. Also, some users from social media are eager to refute a rumor in comments. For such kind of exploration, a convenient and compact way to validate a suspected rumor is highly appreciated.

5 VISUAL DESIGN

As mentioned above, we also proposed straightforward and intuitive visual designs to help platform administrators analyze and validate suspected rumors based on their distribution, features and propagation.

5.1 Suspected Rumors Overview

Suspected Rumors Overview (Figure 1(A) and (B)) is designed to provide platform administrators with an overview of all suspected rumors in terms of spatial-temporal distribution (**R1**). The Suspected Rumors Overview is composed of two main parts, a choropleth map and a line chart. The choropleth map (Figure 1(A)) displays numbers of suspected rumors across different parts of China and overseas as a whole. The line chart (Figure 1(B)) displays the number of changes of suspected rumors with different topics over time.

In the choropleth map on the top, the location with more suspected rumors is shown in a darker color and the exact number of which can be observed by hovering on it. When clicking on the map, corresponding locations are highlighted and selected. The line chart below shows the topic evolution of suspected rumors selected on the choropleth map, with different colored lines representing different topics. Once a topic is chosen, the keywords of suspected rumors under the topic are shown upon the corresponding line. A time axis is also provided to filter suspected rumors further by choosing a start and end time.

5.2 Projection View

Projection View (Figure 1(C)) is designed to help platform administrators quickly inspect and locate the most suspicious features, which may belong to a rumor, for further validation (**R2**). The Projection View consists of glyphs representing filtered suspected rumors projected onto a 2D feature map, where the distance between glyphs indicates the similarity between suspected rumor messages.

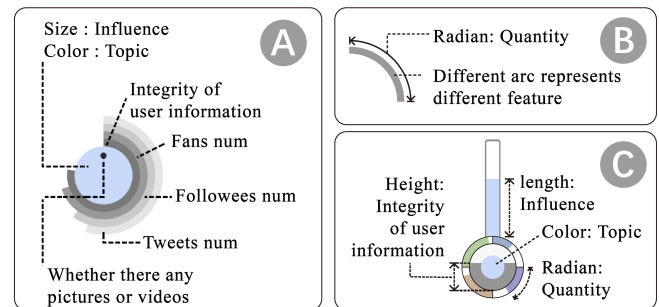


Figure 2: Glyph designs for features of each suspected rumor. (a) round glyph design; (b) arc glyph design; (c) thermograph shape glyph design.

For facilitating visual inspection and comparison of suspicious features, we encode each suspected rumor as a round glyph which is composed of two parts: the inner circle and the outer sectors, as shown in Figure 2 (A). The color of the inner circle represents the topic of the suspected rumor, while the size indicates its influence. The outer four arcs reveal the number of fans, followees, tweets, and integrity of user information respectively, and the glyph is illustrated in Figure 2 (B). In particular, since the number of fans and followees varies greatly among different users, they are calculated by the logarithmic approach for convenient comparison. Compared with the thermograph shape (Figure 2 (C)), our glyph design can

provide a more concise and compact way to show each suspected rumor. Besides, under the guidance of thumbnails on the left side, the position and size of the view displayed can be adjusted by dragging and zooming. Once selected, the glyph will be magnified to highlight the differences from others for close observation.

5.3 Propagation View

Propagation View (Figure 1 (D) and (E)) provides a detailed understanding of how a suspected rumor message diffuses on social media, which enables platform administrators to make a final decision (R3). It is of great importance since rich information in propagation mentioned by experts' experience and previous studies can be displayed and explored for suspected rumor validation. The view contains two parts, a novel circular design to visualize the suspected rumor propagation at the top and a table showing corresponding content details at the bottom.

We propose a novel circular design and it consists of the original tweet in the center and its retweets of multi-levels in the surrounding. As shown in Figure 3, the original tweet represented by a circle node is located in the center, where size indicates its influence. Around it, there are several concentric circular rings with different widths, each denoting the retweets at the same hierarchy. Every concentric circular ring contains hands of sectors separated by radial white lines, which are arranged clockwise by dates. Each sector is fully filled with retweets encoded by a number of cells, the size of which indicates the number of words contained. Also, the cell's color represents whether the sentiment of a retweet is neutral (yellowish), positive (green) or negative (red). In particular, keywords able to provide rough meanings of the retweets are shown within some large cells.

Rich interactions are supported in Propagation View. First, a panel would appear to display some brief user information when hovering on the corresponding cell, and more detailed information can be displayed in Figure 1(E) through clicking, such as user attributes, create time and full content. Second, a comparative analysis is also supported to facilitate a careful inspection. By clicking two cells of retweets successively, information in detail of both can be viewed and checked side by side, as shown in Figure 4. Third, there is a time histogram at the top left corner showing changes in the retweeting number over time. Platform administrators can click on it and highlight cells retweeted on the same day at different hierarchies.

During our design process, we initially considered using a node-link diagram, tree-map, spiral timeline or sunburst chart to visualize a tweet propagation on social media. However, none of them can present retweeting hierarchies and time series from the propagation path simultaneously. Besides, a node-link diagram is not space-efficient enough when a considerable amount of data needs visualization. In comparison, our proposed circular design is able to show the dynamic spreading details of a suspected rumor, facilitating interactive analysis and validation of rumors in a compact manner.

6 CASE STUDY

We conducted a case study on a suspected rumor dataset of Sina Weibo to demonstrate the effectiveness of our proposed visual

design. In this case study, E1, who has rich experience in dealing with rumors, was invited to find a suspected rumor of his interest and further validate whether it is actually a rumor or not.

Overall Exploration. At first, E1 loaded all the suspected rumors into *RumorLens*, and started to interact with the choropleth map in Figure 1(A). He first clicked the button of *overseas* and filtered all the suspected rumors from overseas into the topic evolution view shown in Figure 1(B). As the number of suspected rumors with different topics was changed and displayed over time, E1 paid attention to a blue curve with a relatively high peak, which belonged to the topic of *World News*. After checking through the keywords along the curve, he clicked the label representing *World News* on the left for further exploration.

Feature Inspection. Then he moved to the Projection View to inspect the suspected rumors of *World News* on a feature map. Because the glyphs are too crowded to choose, as shown in Figure 1(C), E1 dragged and zoomed in for better observation until a comparatively bigger glyph than its surroundings is found. According to figure legends on the left, this suspected rumor was of significant influence, yet with an unauthenticated user account and a relatively small number of fans. Since the features mentioned above highly match his work experience, E1 selected the suspected rumor for further validation in the propagation view.

Instance Validation. In Figure 1(D), he found that there were 6 white circular rings, demonstrating that the depth of retweets reached as many as 6, which was deeper than most normal circumstances. Moreover, most of the colors on the view are red and green, which means the retweets showed strong sentiment inclinations, especially for the negative. Then he figured out what happened by clicking the node at the center, as shown in Figure 1(E) and checked keywords floating on some cells for more detail. Most of them were related to street demonstrations, indicating that some social security incidents may have happened yet not focused on the incident itself directly. He continued to explore until a keyword called ID card was found, which may be connected tightly to the incident. When clicking on the cell, a table with detailed information showed the full list of comments saying that the original tweet was misleading. Several retweets supported the viewpoint.

Finally, E1 checked the details of the original tweet again for further validation. It was mainly about an American African who was accidentally arrested and then turned out to be an FBI detective. He found the content of the tweet contains some apparent ambiguous words like *approximately*, which was used to excuse the activity of police officers. According to the social support theory [48], rumors with persuasive and uncertain words are more likely to be disseminated on social media. Individuals with similar experiences are inclined to obtain social support and start to trust those kinds of rumormongers and even nurtured to retweet the rumors. At last, After taking into consideration all the above factors, E1 confidently concluded that the tweet is probably a rumor and should be further dealt with.

7 CONCLUSION AND FUTURE WORK

We have proposed *RumorLens*, an interactive visual analytics system to help administrators of social media platforms deal with suspected rumors efficiently. To tackle the challenges of validating suspected




Figure 3: A glyph design for Propagation View. The graph is used to show the specific encoding and interaction demonstration.

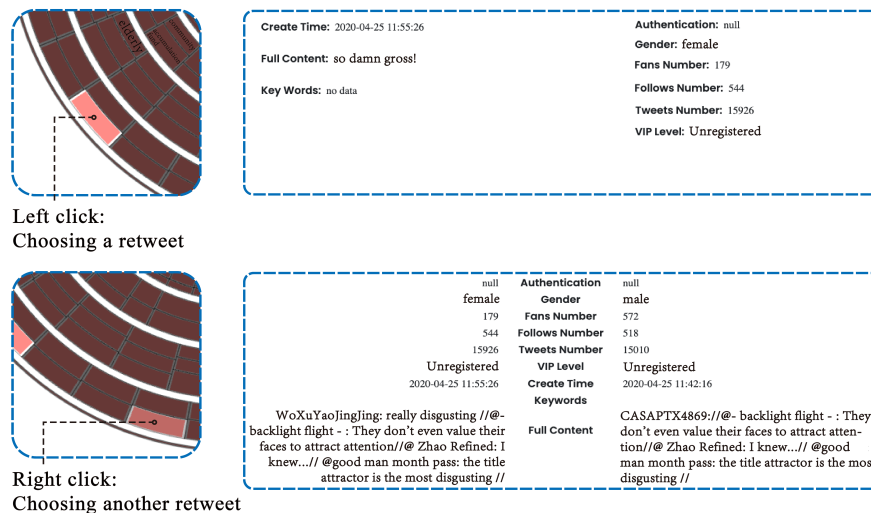


Figure 4: Interaction Design of the Propagation View. It allows users to select two retweets through left and right clicks. The Detail View will show the detailed information of the selected two retweets for a detailed comparison.

rumors, a novel visual interactive system is proposed to provide deep insights into multiple levels of details of the suspected rumors. In addition, a novel circular glyph design is proposed to show the dynamic spreading details of a suspected rumor, facilitating interactive analysis and validation of rumors in a compact manner. A case study conducted by a domain expert provides support for the effectiveness of *RumorLens* in helping users analyze and validate the suspected rumors.

However, *RumorLens* targeting on interactive analysis and validation of suspected rumors still needs further improvement. First, by working with domain experts, we realize the importance of user information for rumor identification. For example, the message has a high risk of being a rumor if the account of the user is undefined and there are identified rumors tweeted several times before. Hence, it is necessary to provide more information on historical complaints related to the user. Second, rumors can be identified by various kinds of features, and thus how to choose and evaluate their effects on rumor validation is still a problem to be solved.

RumorLens is currently designed for and evaluated through the dataset collected from Sina Weibo, where the language is Chinese. In future work, we would like to extend *RumorLens* to other languages and further evaluate its effectiveness with datasets of other

languages. Also, our approach is mainly focusing on the text information of rumor messages, where the image information of some rumor messages is ignored. It will be interesting to further incorporate the image information in the analysis and validation of suspected rumors on social media.

ACKNOWLEDGMENTS

We thank all the domain experts for their valuable feedback, and the anonymous reviewers for their constructive comments. This research is funded by National Social Science Fund of China award number(s): 19CXW032. The computation is completed in the HPC platform of Huazhong University of Science and Technology. Yong Wang is the corresponding author (yongwang@smu.edu.sg).

REFERENCES

- [1] Mohammad Ahsan, Madhu Kumari, and TP Sharma. 2019. Rumors detection, verification and controlling mechanisms in online social networks: A survey. *Online Social Networks and Media* 14 (2019), 100050.
- [2] A. Aker, L. Derczynski, and K. Bontcheva. 2017. Simple Open Stance Classification for Rumour Analysis. In *RANLP 2017 - Recent Advances in Natural Language Processing Meet Deep Learning*. INCOMA Ltd. Shoumen, Varna, Bulgaria, 31–39.
- [3] Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of Economic Perspectives* 31, 2 (2017), 211–36.

- [4] Gordon W Allport and Leo Postman. 1947. The psychology of rumor. *Transactions of the New York Academy of Sciences* 8 (1947), 61–81.
- [5] Daniel Archambault, Derek Greene, Pádraig Cunningham, and Neil Hurley. 2011. ThemeCrowds: Multiresolution summaries of twitter usage. In *Proceedings of the 3rd International Workshop on Search and Mining User-generated Contents*. Elsevier, New York, 77–84.
- [6] Eytan Bakshy, Jake M Hofman, Winter A Mason, and Duncan J Watts. 2011. Everyone's an influencer: quantifying influence on twitter. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining*. ACM, New York, 65–74.
- [7] M. S. Bernstein, B. Suh, L. Hong, J. Chen, and E. H. Chi. 2010. Eddi: Interactive topic-based browsing of social status streams. In *Proceedings of the 23rd Annual ACM symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, 303–312.
- [8] C. Boididou, S. Papadopoulos, M. Zampoglou, L. Apostolidis, and Y. Kompatsiaris. 2017. Detection and visualization of misleading content on Twitter. *International Journal of Multimedia Information Retrieval* 7, 3 (2017), 71–86.
- [9] A. Bondielli and F. Marcelloni. 2019. A Survey on Fake News and Rumour Detection Techniques. *Information Sciences* 497 (2019), 38–55.
- [10] Nan Cao, Yu-Ru Lin, Xiaohua Sun, David Lazer, Shixia Liu, and Huamin Qu. 2012. Whisper: Tracing the spatiotemporal process of information diffusion in real time. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2649–2658.
- [11] Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on twitter. In *Proceedings of the 20th International Conference on World Wide Web*. Association for Computing Machinery, New York, 675–684.
- [12] Anh Dang, Abidrahman Moh'd, Aminul Islam, and Evangelos Milios. 2019. Early detection of rumor veracity in social media. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*. HICSS, Piscataway, 2355–2364.
- [13] Raveena Dayani, Nikita Chhabra, Taruna Kadian, and Rishabh Kaushal. 2015. Rumor detection in twitter: An analysis in retrospect. In *2015 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*. IEEE, Piscataway, 1–3.
- [14] Michela Del Vicario, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. 2016. The spreading of misinformation online. *Proceedings of the National Academy of Sciences* 113, 3 (2016), 554–559.
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *ArXiv abs/1810.04805* (2018), 4171–4186.
- [16] N. Difonzo and P. Bordia. 2007. Rumor, Gossip and Urban Legends. *Diogenes* 54, 1 (2007), 19–35.
- [17] Siwei Fu, Yong Wang, Yi Yang, Qingqing Bi, Fangzhou Guo, and Huamin Qu. 2018. Visforum: A visual analysis system for exploring user groups in online forums. *ACM Transactions on Interactive Intelligent Systems (TiIS)* 8, 1 (2018), 1–21.
- [18] Gang, Liang, Wenbo, He, Chun, Xu, Liangyin, and Chen. 2015. Rumor Identification in Microblogging Systems Based on Users Behavior. *IEEE Transactions on Computational Social Systems* 2, 3 (2015), 99–108.
- [19] A. Gupta, P. Kumaraguru, C. Castillo, and P. Meier. 2014. TweetCred: Real-Time Credibility Assessment of Content on Twitter. In *Springer International Publishing*. Springer, Berlin, 228–243.
- [20] Jeffrey Heer and Danah Boyd. 2005. Vizster: Visualizing online social networks. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005*. IEEE, USA, 32–39.
- [21] Jun Ito, Jing Song, Hiroyuki Toda, Yoshimasa Koike, and Satoshi Oyama. 2015. Assessment of tweet credibility with LDA features. In *Proceedings of the 24th International Conference on World Wide Web*. ACM, New York, 953–958.
- [22] Quanzhi Li, Qiong Zhang, Luo Si, and Yingchi Liu. 2019. Rumor detection on social media: datasets, methods and opportunities. *arXiv abs/1911.07199* (2019), 10 pp.
- [23] C. Y. Lin, C. Nan, X. L. Shi, S. Papadimitriou, and J. Sun. 2009. SmallBlue: Social Network Analysis for Expertise Search and Collective Intelligence. In *IEEE International Conference on Data Engineering*. IEEE, USA, 1483–6.
- [24] Chi Ma, Yi Zhang, Nicolas-Emmanuel Robert, and Yuze Li. 2019. Visual analysis of the opinion flow among multiple social groups. *Journal of Visualization* 23, 3 (2019), 1–15.
- [25] Jing Ma, Wei Gao, Prasenjit Mitra, Sejeong Kwon, Bernard J Jansen, Kam-Fai Wong, and Meeyoung Cha. 2016. Detecting rumors from microblogs with recurrent neural networks. *Proceedings of the 25th International Joint Conference on Artificial Intelligence, IJCAI 2016* 2016-January (2016), 3818–3824.
- [26] Jing Ma, Wei Gao, and Kam-Fai Wong. 2017. Detect rumors in microblog posts using propagation structure via kernel learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Stroudsburg, 708–717.
- [27] Jing Ma, Wei Gao, and Kam-Fai Wong. 2018. Rumor detection on twitter with tree-structured recursive neural networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Stroudsburg, 1980–1989.
- [28] A. Perer and B. Shneiderman. 2006. Balancing Systematic and Flexible Exploration of Social Networks. *Proceedings of IEEE Transactions on Visualization & Computer Graphics* Sept 12 (2006), 693–700.
- [29] Vahed Qazvinian, Emily Rosengren, Dragomir Radev, and Qiaozhu Mei. 2011. Rumor has it: Identifying misinformation in microblogs. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Stroudsburg, 1589–1599.
- [30] Juan Ramos et al. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the First Instructional Conference on Machine Learning*, Vol. 242. Citeseer, Piscataway, 29–48.
- [31] Victoria Rubin, Niall J. Conroy, Yimin Chen, and Sarah Cornwell. 2016. Fake News or Truth? Using Satirical Cues to Detect Potentially Misleading News. In *NAACL-CADD2016: Workshop on Computational Approaches to Deception Detection at the 15th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-CADD2016)*. Association for Computational Linguistics, Stroudsburg, 7–17.
- [32] VP Sahana, Alwyn R Pias, Richa Shastri, and Shweta Mandloi. 2015. Automatic detection of rumoured tweets and finding its origin. In *2015 International Conference on Computing and Network Communications (CoCoNet)*. IEEE, Piscataway, 607–612.
- [33] S. Shelke and V. Attar. 2019. Source detection of rumor in social network – A review. *Online Social Networks and Media* 9 (2019), 30–42.
- [34] Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake news detection on social media: A data mining perspective. *ACM SIGKDD Explorations Newsletter* 19, 1 (2017), 22–36.
- [35] G. Sun, Y. Wu, S. Liu, T Peng, J. J. H Zhu, and R. Liang. 2014. EvoRiver: Visual Analysis of Topic Coepetition on Social Media. *IEEE Transactions on Visualization & Computer Graphics* 20, 12 (2014), 1753–1762.
- [36] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9, 11 (2008), 2579–2605.
- [37] Fernanda Viégas, Martin Wattenberg, Jack Hebert, Geoffrey Borggaard, Alison Cichowlas, Jonathan Feinberg, Jon Orwant, and Christopher Wren. 2013. Google+ripples: A native visualization of information flow. In *Proceedings of the 22nd International Conference on World Wide Web*. International World Wide Web Conference Committee, New York, 1389–1398.
- [38] FB Viégas, M. Wattenberg, and K. Dave. 2004. Studying cooperation and conflict between authors with history flow visualizations. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* 6, 1 (2004), 575–582.
- [39] Soroush Vosoughi. 2015. *Automatic detection and verification of rumors on Twitter*. Ph. D. Dissertation. Massachusetts Institute of Technology.
- [40] Yong Wang, Hammad Haleem, Conglei Shi, Yanhong Wu, Xun Zhao, Siwei Fu, and Huamin Qu. 2018. Towards easy comparison of local businesses using online reviews. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, Wiley, New Jersey, 63–74.
- [41] Ke Wu, Song Yang, and Kenny Q Zhu. 2015. False rumors detection on sina weibo by propagation structures. In *2015 IEEE 31st International Conference on Data Engineering*. IEEE, Piscataway, 651–662.
- [42] Yingcai Wu, Nan Cao, David Gotz, Yap-Peng Tan, and Daniel A Keim. 2016. A survey on visual analytics of social media data. *IEEE Transactions on Multimedia* 18, 11 (2016), 2135–2148.
- [43] Yingcai Wu, Shixia Liu, Kai Yan, Mengchen Liu, and Fangzhao Wu. 2014. Opinionflow: Visual analysis of opinion diffusion on social media. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 1763–1772.
- [44] Guixian Xu, Yueting Meng, Xiaoyu Qiu, Ziheng Yu, and Xu Wu. 2019. Sentiment Analysis of Comment Texts Based on BiLSTM. *IEEE Access* 7 (2019), 51522–51532. <https://doi.org/10.1109/ACCESS.2019.2909919>
- [45] Xichen Zhang and Ali A Ghorbani. 2020. An overview of online fake news: Characterization, detection, and discussion. *Information Processing & Management* 57, 2 (2020), 102025.
- [46] Jian Zhao, Nan Cao, Zhen Wen, Yale Song, Yu-Ru Lin, and Christopher Collins. 2014. # FluxFlow: Visual analysis of anomalous information spreading on social media. *IEEE transactions on Visualization and Computer Graphics* 20, 12 (2014), 1773–1782.
- [47] Zhe Zhao, Paul Resnick, and Qiaozhu Mei. 2015. Enquiring minds: Early detection of rumors in social media from enquiry posts. In *Proceedings of the 24th International Conference on World Wide Web*. ACM, New York, 1395–1405.
- [48] Cheng Zhou, Kai Li, and Yanhong Lu. 2021. Linguistic characteristics and the dissemination of misinformation in social media: The moderating effect of information richness. *Information Processing & Management* 58, 6 (2021), 102679.
- [49] Arkaitz Zubiaga, Ahmet Aker, Kalina Bontcheva, Maria Liakata, and Rob Procter. 2018. Detection and resolution of rumours in social media: A survey. *ACM Computing Surveys (CSUR)* 51, 2 (2018), 1–36.
- [50] A. Zubiaga, M. Liakata, and R. Procter. 2016. Learning Reporting Dynamics during Breaking News for Rumour Detection in Social Media. *ArXiv abs/1610.07363* (2016), 20 pp.