

Fake News Detection Based on Subjective Opinions

Danchen Zhang and Vladimir I. Zadorozhny

School of Computing and Information, University of Pittsburgh, Pittsburgh, USA
{daz45, viz}@pitt.edu

Abstract. Fake news fluctuates social media, leading to harmful consequences. Several types of information could be utilized to detect fake news, such as news content features and news propagation features. In this study, we focus on the user spreading news behaviors on social media platforms and aim to detect fake news more effectively with more accurate data reliability assessment. We introduce Subjective Opinions into reliability evaluation and proposed two new methods. Experiments on two popular real-world datasets, BuzzFeed and PolitiFact, validates that our proposed Subjective Opinions based method can detect fake news more accurately than all existing methods, and another proposed probability based method achieves state-of-art performance.

Keywords: Fake news detection · Data reliability assessment · Subjective Opinions

1 Introduction

In recent years, social media, such as Facebook, Twitter, and WeChat, provides users with direct and instant communications, and people spend more and more time on these platforms. Increasingly, social media become a primary source for people to read the news. A survey from Pew Research Center, an internet research company, shows that 62% of United States adults get news from social media [1]. However, there are no rules for people to produce and spread the news in social media, and hence misinformation could spread more quickly and widely in our society. Consequently, fake news detection draws researchers' increasing attention. According to statistic data in [2], more and more works are devoted to this area.

When detecting fake news, many features and information could be considered, such as news content, news publisher profiles, and data generated in news propagation [2–6]. In this study, we focus on news user interaction in social media. More specifically, the network data constructed by the news and the users who spread the news. In such a network, some news is fake, and other news is real. Researchers could manually collect a set of news and have experts assign labels for the news veracity, and then build models to automatically predict the other unlabeled news.

Since the users may spread both real and fake news, the unlabeled news may be real or maybe fake. Reliability is a quantified measure of uncertainty about an event [8]. In this study, we define **user reliability** as the reliability of a user spreading real news, i.e., more reliable users spread more real news than fake news; less reliable users spread more fake news than real news. Besides, we define **news reliability** as the reliability of the news being real and will propose models to automatically learn the news reliability based on its related users’ reliability.

It is natural to use probability to evaluate the user reliability, i.e., the reliability of a user is the probability that she/he spreads the real news. However, the effectiveness of probability depends on the observed sample size. For example, if we observe the user shares 100 real news and 0 fake news, we could say this user is 100% reliable; if we only observe the user spreads a piece of real news and the veracity of the other 99 news that the user shared remains unknown, we cannot describe this user as 100% reliable.

In reliability assessment, users spreading more unlabeled news should be more uncertain than those spreading more labeled news, and past studies [9–11] fail to consider such uncertainty in their model comprehensively. Therefore, in this study, we propose a new reliability representation and correspondingly a new fake news detection model, which could deal with uncertainty in a better way, and thus detect the fake news more accurately.

Subjective Opinions, proposed by Jøsang, describe the probability affected by the degrees of uncertainty, and Subjective Logic is a calculus for Subjective Opinions [12]. We propose the Subjective Opinions based reliability representation. A subjective opinion from a person p towards a statement s can be represented by a triple $\omega_s^p = \{t, d, u\}$, with $t, d, u \in [0, 1]$, and $t + d + u = 1$. In this triple, t means trust, d means distrust, and u means uncertainty. Given the above example, when we observe the user sharing a lot of real news and no fake news, our opinion towards the statement “this user is reliable” or “this user spreads real news” can be represented as $\{0.98, 0.0, 0.02\}$. Then, when we observe the user sharing a few real news but majority news veracity stays unknown, we are highly uncertain about the user, and our opinion could be $\{0.1, 0.0, 0.9\}$. In this way, with Subjective Opinions based reliability, we could see the difference between two users.

The main contributions of this paper include:

- We are first to introduce Subjective Opinions into the fake news detection area.
- We propose a Subjective Opinion based model and a Probability based model to detect fake news automatically.
- We conduct extensive experiments on two real-world datasets, validate our proposed methods’ effectiveness, and compare them in several aspects.

We introduced the background and motivation in Section 1 and discuss related works in Section 2. Next, we present our two proposed methods in Section 3. The experiments, results, and discussions are shown in Section 4 and 5. Lastly, the paper is concluded in Section 6.

2 Related works

Several recent survey papers encompass the full range of research devoted to fake news [2–6]. Fake news detection mainly uses three types of information: (1) the content of news articles, including word-level, syntactic level, and semantic level information, (2) news propagation data on social networks, such as user profiles, news profiles, spreading data, and (3) network structure extracted from news articles and social media. Past works usually learn a classification model with these features to detect fake news automatically.

At the early stage of fake news detection, word level and syntactic level features were the most useful features. Several studies [14–17] show that bag-of-words features give a very effective way to reveal the linguistic cues of deception. Then, the syntactic features, such as the distribution of POS tagger and n-grams, are also able to capture the unique features of fake news [18–21]. Later, the development of neural networks in natural language processing (NLP) allows the semantic features to be extracted, further improving the detection performance, such as works in [22–24]. These studies show that both word and syntactic information are essential, and the extraction of high-quality semantic features can further increase detection effectiveness.

Meanwhile, many researchers explore news propagation features, such as the number of likes, propagation times, and related user profiles. However, many studies have found that systems cannot accurately detect fake news if they only use social network features. These features are usually used together with content mentioned above features, such as works in [17, 25, 26].

Besides, multiple network structures can be extracted from the news related data, such as user-follow-user networks, news-agree/conflict-news networks, publisher-publish-news network, and user-spread-news networks [9–11, 27–30]. In this study, we focus on the user-spread-news network, and hence we chose the state-of-art methods Harmonic [9], HC-CB-3 [10], and TriFN [11] as baselines, which will be discussed more in Section 4.

To the best of our knowledge, this work is the first to apply Subjective Opinions to the area of fake news detection. Before this, it is widely used in many other areas, including trust network analysis [31], conditional inference [32], information provider reliability assessment [33], trust management in sensor networks [34]. In one of our earlier works [35], Subjective Opinions is introduced for Truth Discovery. With the data represented and recorded more comprehensively, our proposed Subjective Opinions based method is expected to have a better performance than the probability based method.

3 Proposed Methods

Consider a dataset that contains a set of news $News = \{news_1, news_2, \dots, news_n\}$, and a set of users $Users = \{user_1, user_2, \dots, user_m\}$. Users could share the news on social media. Such data could be represented as a matrix shown in Table 1. Each row corresponds to a piece of news, each column corresponds to a user,

Table 1: The dataset is represented by a matrix, with n news and m users.

	user₁	user₂	...	user_m
news₁	<i>share</i>	—	...	<i>share</i>
news₂	—	<i>share</i>	...	—
...
news_n	—	—	...	<i>share</i>

and the cell values show whether the user shares news or not. Also, we assume that the veracity labels for a subset of news are known, and the veracity labels for the rest news need to be predicted.

In this study, we propose two semi-supervised models to evaluate the reliability for users and news, and the news reliability scores decide the news veracity. In the first model, news and user reliability is evaluated with probability and is named as **Prob_fnd**. In the second model, the reliability is evaluated with Subjective Opinions and is named as **SO_fnd**. Both two models consist of three steps: (1) update user reliability based on news labels; (2) update news reliability based on user reliability; (3) predict and update labels for news that are not in training data. These three steps iteratively run until the reliability scores converge. We find that our methods usually converge within three iterations. Table 2 compares each step of Prob_fnd and SO_fnd, and we can see that the main difference is whether probability or Subjective Opinion is used, and correspondingly whether the unknown data are ignored or recorded.

Table 2: Comparison of Prob_fnd and SO_fnd

Steps	Prob_fnd	SO_fnd
Step 1: Update <i>user_reliability</i> .	Use Probability to describe our belief in a user sharing real news. Ignore unknown cases.	Use Subjective Opinions to describe our belief in a user sharing real news. Record unknown cases with uncertainty.
Step 2: Update <i>news_reliability</i> .	Use Average to fuse related users' reliability.	Use Consensus to fuse related users' reliability.
Step 3: Update news veracity.	Predict with <i>news_reliability</i> .	Predict with <i>news_reliability</i> .

3.1 Method 1: Probability based fake news detection (Prob_fnd)

Step-1: evaluate user reliability. For each user $user_i$, the reliability is defined as the probability that this user shares the real news, and hence we have the following formula:

$$user_reliability(user_i) = \frac{\#(r_n)}{\#(r_n) + \#(f_n)}, \quad (1)$$

where $\#(r_n)$ is the count of real news that $user_i$ shared, and $\#(f_n)$ is the count of fake ones. In early iterations, if the veracity of the news that the $user_i$ shared is not predicted yet, this news is not counted in this formula; if all news veracity that $user_i$ shared are unknown, reliability cannot be judged and $user_i$ is temporarily classified as unknown users.

Step-2: evaluate news reliability. For each news $news_j$, we collect all users who shared it and define the news reliability as the average $user_reliability$ of these related users, and we have the following formula:

$$news_reliability(news_j) = Average(user_reliability(related\ users)). \quad (2)$$

In early iterations, if a user's reliability is unknown yet, it is temporarily ignored; if all related users' reliability is unknown, the reliability of this news will be calculated in the future iteration and temporarily classified as unknown news. Please note that we also tried to use the Median to replace Average, and experiment results show no significant difference.

Step-3: predict news veracity. Then, $\{news_reliability\}$ is used as the feature and is put into Support Vector Machine (SVM) to learn a classification model to automatically determine the veracity label for the news. Only the news labels in training data are used as initial labels, and the model will predict veracity labels for the news in testing data.

3.2 Method 2: Subjective Opinions based fake news detection (SO_fnd)

In the second method, we replace probability with Subjective Opinions. As described in Section 1, people's subjective opinion towards a statement can be represented as a triple $\{trust, distrust, uncertainty\}$. Information of unknown users and news is not kept in Prob_fnd, while we record such information with the *uncertainty* in SO_fnd, and hope that reliability assessment could be more accurate with more comprehensive information considered. SO_fnd also consists of three steps, as shown below.

Step-1: evaluate user reliability. User reliability is defined as our opinion towards the statement " $user_i$ is reliable", which is defined as:

$$user_reliability(user_i) = \{user_trust, user_distrust, user_uncertainty\}, \quad (3)$$

where

$$\begin{cases} user_trust = \frac{\#(r_n)}{\#(r_n) + \#(f_n) + \#(u_n)} * (1 - \alpha), \\ user_distrust = \frac{\#(f_n)}{\#(r_n) + \#(f_n) + \#(u_n)} * (1 - \alpha), \\ user_uncertainty = \frac{\#(u_n)}{\#(r_n) + \#(f_n) + \#(u_n)} * (1 - \alpha) + \alpha. \end{cases} \quad (4)$$

α is a constant, representing people's natural uncertainty. $\#(u_n)$ is the count of unknown news that this user shared. If all news veracity that a user shared are unknown, reliability of this user is $\{0, 0, 1\}$.

Step-2: evaluate news reliability. For each news, we collect all users who shared it and cumulatively fuse these users' reliability with Subjective Logic consensus operation (See more information about consensus operation in [13]). The fused opinion represents our belief for the statement " $news_j$ is real".

$$news_reliability(news_j) = user_1_reliability \oplus user_2_reliability \oplus \dots \oplus user_k_reliability, \quad (5)$$

where $user_1, user_2, \dots, user_k$ are the users that shared $news_j$. The obtained $news_reliability$ is a triple, as shown in:

$$news_reliability(news_j) = \{news_trust, news_distrust, news_uncertainty\}. \quad (6)$$

Step-3: predict news veracity based on reliability. There are a lot of unknown users and news in the first iteration, leading to high *uncertainty* in reliability assessment, making the news veracity prediction more challenging than that in later iterations. If in such an uncertain situation, given a piece of news, the model could accumulate high enough *news_distrust*, we are confident to predict that this news is fake. Thus, we use $\{news_distrust\}$ as the feature and put it into SVM to learn a classification model to predict the news veracity automatically (see more discussion about classifier and feature selection in Section 4 and 5).

Next, from the second iteration, with the number of unknown users and news decreasing, the situation is less uncertain, and the reliability assessment is hence more convincing. Then the news veracity is decided by comparing *news_trust* and *news_distrust*. I.e., the news with higher *news_trust* are predicted to be real, and ones with higher *news_distrust* are predicted to be fake.

4 Experiments

4.1 Data

FakeNewsNet [36] is selected as the validation dataset in this study. It consists of two real-world datasets, BuzzFeed and PolitiFact. BuzzFeed contains 90 real news and 90 fake news, with 15,257 users interacting with (re-tweet or like) the news. PolitiFact contains 120 real news, 120 fake news, and 23,865 users interacting with the news.

In the preprocessing procedure, if a user ($user_i$) shares only one piece of news ($news_j$), this user is removed. The reason is that (1) if the $news_j$ veracity is unknown, we cannot infer $user_i$ reliability, and hence $user_i$ cannot provide information; (2) if the $news_j$ veracity is known, as $user_i$ does not share other news, model needs no information from it. Please note that, after preprocessing, if the news no longer has related users, we directly label it as fake news and do not update its label in iterations. We do know such a straightforward strategy

may lead to failure, but such loss on unpopular news is acceptable in real-world scenarios.

After preprocessing, 3,002 users are left in BuzzFeed, and 4,139 users are left in PolitiFact. Also, we find the PolitiFact data is denser than BuzzFeed data, i.e., people share more news in PolitiFact.

4.2 Experiment settings

Following procedures in [9–11], we learn and evaluate our models with 5-fold cross-validation, i.e., 20% of data is used as testing, while 80% of data is used to train the model. Each cross-validation is repeated 50 times, and the average performance with standard deviation is reported. Accuracy, Precision, Recall, and F1 of detecting fake news are selected as the evaluation metrics.

Prob_fnd has no parameters, while SO_fnd has one parameter α , which describes people’s natural uncertainty. Following [9–11], we select $\alpha = 0.9$, because it achieves the highest performance with both datasets in cross-validation.

Please note that, in Step-3 of both Prob_fnd and SO_fnd, SVM is used to predict news veracity. We tried different classification models, including SVM, Naive Bayes, Decision Tree, Random Forest, and Logistic Regression. We found that SVM could provide the best performance for two models on two datasets. Besides, with different α in SO_fnd, SVM provides more robust performance than other models. Therefore, in this study, SVM is selected.

4.3 Baselines

Harmonic from [9]. This method is very similar to our proposed Prob_fnd, iteratively evaluating the reliability scores of both users and news, and both methods ignore the unknown news and users in the calculation. The major difference is that Harmonic explicitly differentiate reliable users from unreliable users, and real news is those that accumulate more scores from reliable users, while fake news is those that accumulate more scores from unreliable users. On the other hand, in Prob_fnd, news reliability is defined as the average reliability of the users that shared it without explicitly different reliable and unreliable users.

HC-CB-3 from [10]. This method is developed based on Harmonic. It utilizes the word-level features of news content with a logistic regression model. If the news is shared by more than λ people, social-network based Harmonic is used; otherwise, content based classification is used.

TriFN from [11]. This method designed a Tri-Relationship embedding framework, which utilizes the information from news content, news-user interaction, and news-publisher relationship. TriFN shows much better performance than several other baselines, which use content based or social-network based features, and they are not included in this paper due to page limit.

4.4 Results

Experiment results are shown in Table 3, and their comparison is better illustrated in Figure 1. Best and second-best performed runs are labeled with ‘**’

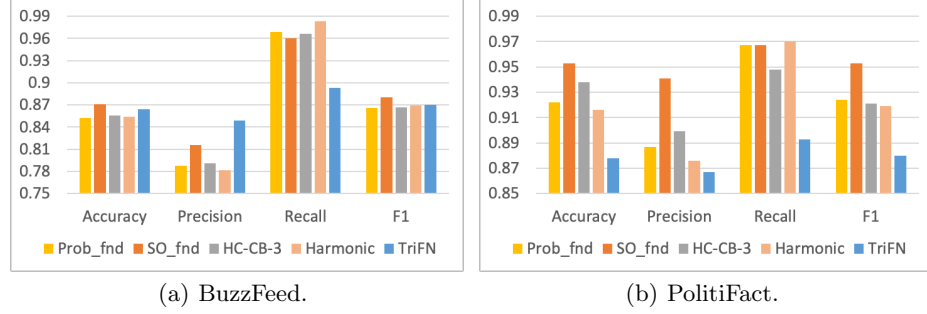


Fig. 1: Repeated 5-fold cross-validation results on two real-world datasets.

and '*'. In this study, the Wilcoxon Signed Ranks test is only conducted among the results of Prob_fnd and SO_fnd. As baseline papers only report the averaged results with standard deviation, the significance test cannot be conducted. Since all standard deviation values are relatively small (< 0.1), when comparing the model performance in this section, we are comparing the reported average values.

Table 3: Repeated 5-fold cross-validation results on two real-world datasets.

		Prob_fnd	SO_fnd	HC-CB-3	Harmonic	TriFN
BuzzFeed	accuracy	.852 ± .055	.871 ± .051**	.856 ± .052	.854 ± .052	.864 ± .026*
	precision	.788 ± .086	.816 ± .079*	.791 ± .076	.782 ± .075	.849 ± .040**
	recall	.969 ± .043*	.960 ± .004	.966 ± .045	.983 ± .041**	.893 ± .013
	F1	.866 ± .052	.880 ± .050**	.867 ± .050	.869 ± .050	.870 ± .019*
PolitiFact	accuracy	.922 ± .036	.953 ± .029**	.938 ± .029*	.916 ± .042	.878 ± .020
	precision	.887 ± .056	.941 ± .048**	.899 ± .057*	.876 ± .074	.867 ± .034
	recall	.967 ± .034*	.967 ± .034*	.948 ± .046	.970 ± .030**	.893 ± .023
	F1	.924 ± .035*	.953 ± .030**	.921 ± .041	.919 ± .044	.880 ± .017

x^{**} : the run with the best performance.

x^* : the run with the second best performance.

Our proposed SO_fnd has the best performance on reported Accuracy and F1 in both datasets. To be more specific, on the BuzzFeed dataset, SO_fnd got the best Accuracy 87.1% and F1 score 88.0%, and second-best precision 81.6%; on the PolitiFact dataset, SO_fnd has the best Accuracy 95.3%, F1 score 95.3%, precision 94.1%, and second-best recall 96.7%. It shows that SO_fnd can differentiate fake news from real news much more accurately than other methods. Besides, on both datasets, we find that SO_fnd significantly outperforms Prob_fnd on Accuracy, Precision, and F1 ($p\text{-value} < 0.01$). It indicates that keeping a record of unknown cases as an uncertainty value is essential, which can highly improve the fake news detection accuracy.

We can also see that on two datasets, both the baselines and our proposed methods have lower precision and higher recall. It indicates that most fake news could be correctly detected, but many pieces of real news are wrongly classified as fake news. Such a model is better than one with higher precision and lower recall. The reasons are: (1) the broad propagation of fake news may lead to inestimable damages, and hence we could sacrifice some real news to get a high recall on fake news detection; (2) we can hire experts to check these identified fake news manually, and since the data size decreases a lot, the manual effort cost is lowered down; (3) if a piece of real news is falsely recognized as fake news and removed, people could still search online or ask experts if they really need it.

Besides, we can find that Prob.fnd has very similar performance with Harmonic, implying that whether or not explicitly differentiate reliable users from unreliable users does not make a big difference in these two datasets.

5 Discussion

5.1 Is SO.fnd performance sensitive to natural uncertainty parameter?

SO.fnd has one parameter α , which describes people’s natural/basic uncertainty. Above reported results are obtained when $\alpha = 0.9$, which is the highest performance with both datasets in cross-validation. We report the Accuracy and F1 of SO.fnd when α changes in range $[0.1, 0.9]$. We repeat the cross-validation procedure 50 times, and the average performance is reported in Figure 2.

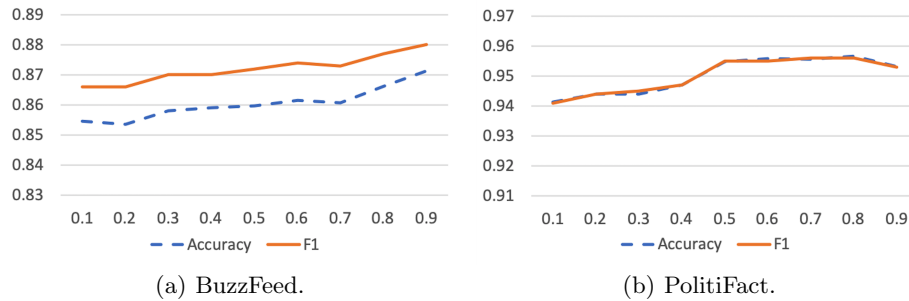


Fig. 2: Accuracy and F1 of SO.fnd varying with different α on BuzzFeed and PolitiFact.

From Figure 2, we can observe that though SO.fnd prefers larger α in both datasets when α varies across $[0.1, 0.9]$, the accuracy and F1 do not change a lot, with the increments less than 2%. It shows SO.fnd is relatively robust to the parameter α in range $[0.1, 0.9]$.

5.2 How does the training data size affect the performance of the methods?

In this subsection, we explored the performance of Prob_fnd and SO_fnd when they are trained by different sizes of data. As shown in Figure 3, the size of training data increases from 10% to 90%, and the accuracy is evaluated for each model on both BuzzFeed and PolitiFact. We repeat the training and testing procedure 50 times for each run, and the average performance is reported.

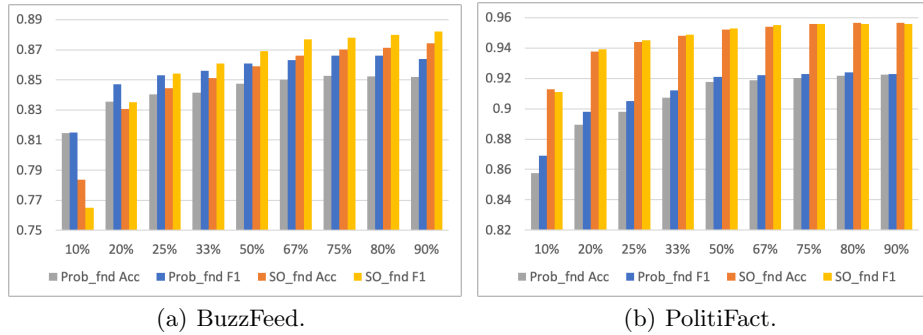


Fig. 3: Accuracy and F1 of Prob_fnd and SO_fnd varying with different training data size on BuzzFeed and PolitiFact.

From Figure 3, we can observe that, Accuracy and F1 score of Prob_fnd and SO_fnd all increase in both datasets when training data size rises. Also, SO_fnd’s performance outperforms Prob_fnd in most cases, except when they are trained with 20% or less data on BuzzFeed.

Shu et al. reported TriFN’s performance with the different training set sizes in [11]. On BuzzFeed, when training data is 40% and less, TriFN’s Accuracy and F1 is less than 80%; however, Prob_fnd’s Accuracy and F1 are above 80% even with 10% training data, and SO_fnd’s Accuracy and F1 are above 80% with 20% or more training data. On PolitiFact, TriFN’s Accuracy and F1 are above 80% with 40% or more training data; however, SO_fnd’s Accuracy and F1 are above 90% even with 10% training data, and Prob_fnd’s Accuracy and F1 are above 90% with 33% or more training data. It shows that, compared to TriFN, our proposed two models are able to achieve a similar or even better performance with much less labeled training data.

5.3 In what situations, will Prob_fnd and SO_fnd win and fail?

In this subsection, we explore in what situations two proposed methods shall win and shall fail. This experiment is conducted on PolitiFact because two methods’ performance difference is larger on it than BuzzFeed. First, we use all news labels

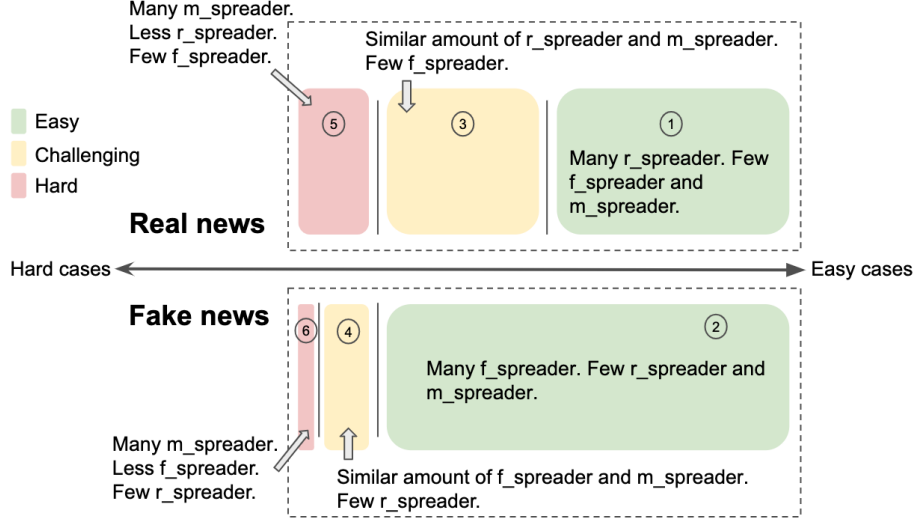


Fig. 4: Easy classified, challenging, and hard classified news in PolitiFact.

in PolitiFact to calculate *user_reliability* in Formula 1. Second, we mark users whose *user_reliability* > 0.8 as **real news spreader (r_spreader)**, mark users whose *user_reliability* < 0.2 as **fake news spreader (f_spreader)**, and mark other users as the **mixed news spreader (m_spreader)**. Among 4,132 users, there are 994 r_spreader, 2,619 f_spreader and 569 m_spreader. Surprisingly, 2,594 users are spreading only fake news, and 979 users are spreading only real news. Third, as shown in Figure 4, based on the distribution of three types of users, we identify three levels of difficulty for news classification:

- **Easily classified news.** News that is shared mainly by r_spreader, as shown in area 1 in Figure 4, are very likely to be real news; news that is spread mainly by f_spreader, as shown in area 2, are very likely to be fake news. Usually, r_spreader and f_spreader could be easily identified with training data, and hence models can easily classify them.
- **Challenging news.** If the news is shared by a similar amount of m_spreader and r_/f_spreader, as shown in areas 3 and 4 in Figure 4, the classification performance is affected by the reliability assessment of m_spreader. The reliability assessment accuracy for m_spreader vary in different models, and the classification of such news is challenging. We found that our proposed two models were able to identify them successfully.
- **Hard classified news.** If the news is shared mainly by m_spreader, as shown in area 5 and 6 in Figure 4, it is hard to classify them. The classification performance is directly decided by (1) the reliability assessment accuracy for the m_spreader, and (2) how the news reliability assessment is designed. We checked the failure cases for both methods in repetitive experiments and found that there were 15 out of 17 frequently appearing failure cases in

Prob_fnd, and 7 out of 9 frequent failure cases in SO_fnd, can be attributed to the large number of m_spreader.

Two other failure cases for Prob_fnd and SO_fnd are: (1) a piece of real news (news id 84) is spread by more f_spreader and hence is wrongly classified. (2) a piece of real news (news 22) is wrongly classified because its related user information (reliability) cannot be assessed from the data, and it is directly classified as fake news.

From Figure 4, we can also observe that fake news is mainly spread by f_spreader (area 2 is very large), but only half real news is mainly spread by r_spreader (area 1), with the other half real news spread by a lot of m_spreaders (area 3 and 5). It indicates that most fake news in PolitiFact is easily classified news, and only a few are challenging or hard classified fake news, while nearly half of real news is challenging or hard to classify. With assessed news reliability of SO_fnd, when *news_distrust* is high, the news is likely to be fake; when *news_distrust* is low, the news is likely to be real (because challenging and hard classified fake news are few). However, when *news_trust* is high, the news is likely to be real; when *news_trust* is low, the news may be hard classified real news, or maybe fake news. It explains why *news_distrust* has a better performance than *news_trust* when used as the SVM classification feature in the first iteration of the SO_fnd. Hence, *news_distrust* is selected instead of *news_trust*.

6 Conclusion and future works

In this study, to detect the fake news that spreads on the social network, we propose two new methods. Validated on two real-world datasets, our proposed probability based method, Prob_fnd, can achieve state-of-art performance. Furthermore, our proposed Subjective Opinions based method, SO_fnd, can record the model’s uncertainty more accurately in reliability evaluation, and hence has a much better performance than all existing models. Also, we explored our model’s sensitivity to the parameters, training data size’s effect on model performance, and also our models’ winning and losing cases.

In this process, we do notice that there is a set of news lacking related user information, and we plan to work on the news content mining to predict the veracity of the news. As a next step, we will also consider introducing the user profile features, such as age, account verification, registration time, and follower count.

References

1. Sihombing, Sabrina Oktaria. "Predicting intention to share news through social media: An empirical analysis in Indonesian youth context." *Business and Economic Horizons* 13.4 (2017): 468-477.
2. Bondielli, Alessandro, and Francesco Marcelloni. "A survey on fake news and rumour detection techniques." *Information Sciences* 497 (2019): 38-55.

3. Kumar, Srijan, and Neil Shah. "False information on web and social media: A survey." arXiv preprint arXiv:1804.08559 (2018).
4. Shu, Kai, et al. "Fake news detection on social media: A data mining perspective." ACM SIGKDD Explorations Newsletter 19.1 (2017): 22-36.
5. Conroy, Niall J., Victoria L. Rubin, and Yimin Chen. "Automatic deception detection: Methods for finding fake news." Proceedings of the Association for Information Science and Technology 52.1 (2015): 1-4.
6. Chen, Yimin, Niall J. Conroy, and Victoria L. Rubin. "Misleading online content: recognizing clickbait as" false news"." Proceedings of the 2015 ACM on workshop on multimodal deception detection. 2015.
7. Walker, Warren E., et al. "Defining uncertainty: a conceptual basis for uncertainty management in model-based decision support." Integrated assessment 4.1 (2003): 5-17.
8. Wang, Yijie, et al. "A survey of queries over uncertain data." Knowledge and information systems 37.3 (2013): 485-530.
9. acchini, E., Ballarin, G., Della Vedova, M. L., Moret, S., de Alfaro, L. (2017). Some like it hoax: Automated fake news detection in social networks. arXiv preprint arXiv:1704.07506.
10. Della Vedova, M. L., Tacchini, E., Moret, S., Ballarin, G., DiPierro, M., de Alfaro, L. (2018, May). Automatic online fake news detection combining content and social signals. In 2018 22nd Conference of Open Innovations Association (FRUCT) (pp. 272-279). IEEE.
11. Shu, Kai, Suhan Wang, and Huan Liu. "Exploiting tri-relationship for fake news detection." arXiv preprint arXiv:1712.07709 (2017).
12. Jøsang, Audun. Subjective logic. Heidelberg: Springer, 2016.
13. Jøsang, Audun. "The consensus operator for combining beliefs." Artificial Intelligence 141.1-2 (2002): 157-170.
14. Rubin, V. L., Lukoianova, T. (2015). Truth and deception at the rhetorical structure level. Journal of the Association for Information Science and Technology, 66(5), 905-917.
15. Wang, William Yang. "'liar, liar pants on fire': A new benchmark dataset for fake news detection." arXiv preprint arXiv:1705.00648 (2017).
16. Potthast, M., Kiesel, J., Reinartz, K., Bevendorff, J., Stein, B. (2017). A stylistic inquiry into hyperpartisan and fake news. arXiv preprint arXiv:1702.05638.
17. Kumar, S., West, R., Leskovec, J. (2016, April). Disinformation on the web: Impact, characteristics, and detection of wikipedia hoaxes. In Proceedings of the 25th international conference on World Wide Web (pp. 591-602). International World Wide Web Conferences Steering Committee.
18. Hassan, N., Li, C., Tremayne, M. (2015, October). Detecting check-worthy factual claims in presidential debates. In Proceedings of the 24th acm international on conference on information and knowledge management (pp. 1835-1838). ACM.
19. Ott, M., Choi, Y., Cardie, C., Hancock, J. T. (2011, June). Finding deceptive opinion spam by any stretch of the imagination. In Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies-volume 1 (pp. 309-319). Association for Computational Linguistics.
20. Feng, S., Banerjee, R., Choi, Y. (2012, July). Syntactic stylometry for deception detection. In Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2 (pp. 171-175). Association for Computational Linguistics.

21. Afroz, S., Brennan, M., Greenstadt, R. (2012, May). Detecting hoaxes, frauds, and deception in writing style online. In 2012 IEEE Symposium on Security and Privacy (pp. 461-475). IEEE.
22. Ajao, O., Bhowmik, D., Zargari, S. (2018, July). Fake news identification on twitter with hybrid cnn and rnn models. In Proceedings of the 9th International Conference on Social Media and Society (pp. 226-230). ACM.
23. Song, C., Tu, C., Yang, C., Liu, Z., Sun, M. (2018). CED: Credible Early Detection of Social Media Rumors. arXiv preprint arXiv:1811.04175.
24. Zubiaga, A., Kochkina, E., Liakata, M., Procter, R., Lukasik, M., Bontcheva, K., ... Augenstein, I. (2018). Discourse-aware rumour stance classification in social media using sequential classifiers. *Information Processing Management*, 54(2), 273-290.
25. Liu, Y., Jin, X., Shen, H. (2019). Towards early identification of online rumors based on long short-term memory networks. *Information Processing Management*, 56(4), 1457-1467.
26. Li, Q., Zhang, Q., Si, L. (2019, July). Rumor Detection by Exploiting User Credibility Information, Attention and Multi-task Learning. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (pp. 1173-1179).
27. Gupta, M., Zhao, P., Han, J. (2012, April). Evaluating event credibility on twitter. In Proceedings of the 2012 SIAM International Conference on Data Mining (pp. 153-164). Society for Industrial and Applied Mathematics.
28. Jin, Z., Cao, J., Zhang, Y., Luo, J. (2016, March). News verification by exploiting conflicting social viewpoints in microblogs. In Thirtieth AAAI Conference on Artificial Intelligence.
29. Ruchansky, N., Seo, S., Liu, Y. (2017, November). Csi: A hybrid deep model for fake news detection. In Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (pp. 797-806). ACM.
30. Guacho, G. B., Abdali, S., Shah, N., Papalexakis, E. E. (2018, August). Semi-supervised Content-based Detection of Misinformation via Tensor Embeddings. In 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM) (pp. 322-325). IEEE.
31. Jøsang, Audun, Ross Hayward, and Simon Pope. "Trust network analysis with subjective logic." Proceedings of the 29th Australasian Computer Science Conference-Volume 48. Australian Computer Society, Inc., 2006.
32. Josang, Audun. "Conditional reasoning with subjective logic." *Journal of Multiple-Valued Logic and Soft Computing* 15.1 (2008): 5-38.
33. Pelechrinis, Konstantinos, et al. "Automatic evaluation of information provider reliability and expertise." *World Wide Web* 18.1 (2015): 33-72.
34. Oleshchuk, Vladimir, and Vladimir Zadorozhny. "Trust-aware query processing in data intensive sensor networks." *sensorcomm*. IEEE, 2007.
35. Zhang, Danchen, Vladimir I. Zadorozhny, and Vladimir A. Oleshchuk. "SLFTD: A Subjective Logic Based Framework for Truth Discovery." *European Conference on Advances in Databases and Information Systems*. Springer, Cham, 2019.
36. Shu, Kai, et al. "Fakenewsnet: A data repository with news content, social context and dynamic information for studying fake news on social media." arXiv preprint arXiv:1809.01286 (2018).