

FAKE NEWS DETECTION WITH SEMANTIC FEATURES AND TEXT MINING

Pranav Bharadwaj¹ and Zongru Shao²

¹South Brunswick High School, Monmouth Junction, NJ

²Spectromn , USA

ABSTRACT

Nearly 70% of people are concerned about the propagation of fake news. This paper aims to detect fake news in online articles through the use of semantic features and various machine learning techniques. In this research, we investigated recurrent neural networks vs. the naive bayes classifier and random forest classifiers using five groups of linguistic features. Evaluated with real or fake dataset from kaggle.com, the best performing model achieved an accuracy of 95.66% using bigram features with the random forest classifier. The fact that bigrams outperform unigrams, trigrams, and quadgrams show that word pairs as opposed to single words or phrases best indicate the authenticity of news.

KEYWORDS

Text Mining, Fake News, Machine Learning, Semantic Features, Natural Language Processing (NLP)

1. INTRODUCTION

Nearly 70% of the population is concerned about malicious use of fake news [3]. Fake news detection is a problem that has been taken on by large social-networking companies such as Facebook and Twitter to inhibit the propagation of misinformation across their online platforms. Some fake news articles have targeted major political events such as the 2016 US Presidential Election and Brexit [4]. However, the scope of fake news extends beyond globally significant political events. Individuals falsely reported that a golden asteroid on target to hit Earth contains \$10 quadrillion worth of precious metals in an attempt to increase the value of Bitcoin [1]. With fake news infiltrating multiple facets of public information, many are rightly concerned. According to a Pew Research Center survey, fake news and misinformation has had a significant impact on 68% of Americans' confidence in government, 54% of Americans' confidence in each other, and 51% of Americans' confidence in their political leaders to get work done [5]. Additionally, the previously mentioned survey also states that 79% of US adults believe that measures should be taken to inhibit the propagation of misinformation [5]. Residents of a Macedonian town named Veles use Google AdSense to distribute fake news around the internet, run politically manipulative Facebook pages and websites in order to make a living [12]. One of their Facebook pages had garnered over 1.5 million followers [12]. The rising problem that is fake news has become increasingly important because of the vulnerability of the massive readers and its widespread malicious influence. As these invalid sources of information have gained traction and have established themselves as credible informants to many individuals, preventing this category of content from spreading and detecting it at the source has become increasingly crucial. Consequently, automated identification of fake news has been studied by Facebook and Twitter as well as other researchers [9].

In this paper, we present a comparison between recurrent neural networks, naive bayes, and
DOI: 10.5121/ijnlc.2019.8302

random forest algorithms using various linguistic features for fake-news-detection. We use the *real or fake* dataset from kaggle.com evaluate these models. The remainder of this paper is structured into three sections. Section 2 details related works, how they approached detecting fake news, and Section 3 describes the semantic features and machine learning algorithms in our experiment. Section 4 presents the evaluation results, in which random forest with bigram features achieved the best accuracy of 95.66%. Section 5 presents the conclusions and future work.

2. RELATED WORKS

Several solutions were proposed for this problem. Prior studies employed logistic regression and “boolean crowd-sourcing algorithms” to detect fake news in social networking websites [13]. However, this research assumed that agents who post misinformation can be detected by the users who have prior contact of the content [13]. Another study used convolutional neural networks (CNNs), with a long short term memory (LSTM) layer to detect fake news by the text context and additional metadata [14]. Shu et al. studied linguistic features such as word count, frequency of words, character count, and similarity and clarity score for videos and images while proposing rumor classification, truth discovery, click-bait detection, and spammer and bot detection [11]. Rubin et al. proposed to classify fake news as one of three types: (a) serious fabrications, (b) large-scale hoaxes, (c) humorous fakes. It also discussed the challenges that each variant of fake news presents to its detection [11].

However, none of the prior studies had utilized recurrent neural networks (RNNs), naive bayes, or random forest.

3. MATERIALS AND METHODS

In this section, we describe the material dataset, text preprocessing, semantic features including term frequency (TF), term frequency–inverse document frequency (TFIDF), bigrams, trigram, quadgram, vectorized word representations, and machine learning algorithms such as naive bayes classifier, random forest, recurrent neural networks (RNN). The process of making each model has been detailed below in Figure 1.



Figure 1

3.1 Dataset

We used *real-or-fake news* dataset from kaggle.com in our experiments to evaluate semantic features. It contains 6256 articles including their titles. 50% of the articles are labeled as *FAKE* and the remaining as *REAL*. Therefore, detecting the *FAKE* articles is a binary classification problem. We split the dataset into 80% for training and 20% for testing.

3.2 Text Pre-Processing

We pre-process the raw text to extract semantic features for machine learning. We use n-grams as semantic features. We first tokenize the title and body of each article. Then, each token is transformed into lower cases and proper nouns lose their capital-letter information. Next, we remove stopwords and numbers for unigrams since they carry less meaning in the context. As a

result, the remaining tokens are semantic representations from linguistic perspective. Stopwords and numbers are reserved for n-grams other than unigrams. Then, we extract TF and TFIDF numerical features using the semantic representations.

3.3 Linguistic Features

3.3.1 TF and TFIDF

Note that a text subject (e.g., an article) is called a document in natural language processing. TF computes how frequently a term appears in a document. Given a document d with a set of terms $T = \{t_1, t_2, \dots, t_M\}$, and the document length is N (the total occurrence of all terms); suppose term t_i appeared x_i times; then, TF of t_i is denoted as

$$TF(t_i, d) = \frac{x_i}{N}$$

As a result, $[TF(1), TF(2), \dots, TF(i), \dots, TF(M)]$, $i \in [1, M]$ is a semantic representation for the document. Inverse term frequency (IDF) denotes the popularity of a term across documents. Given a set of documents $D = \{d_1, d_2, \dots, d_k\}$ as the subjects of interest, and TF(i) for term t_i is calculated for each document; suppose $C(i)$ denotes the number of documents in which $x_i \neq 0$; then,

$$IDF(t_i, D) = \frac{K}{C_i}$$

Note that each term appears at least once in D . Meanwhile, TF and IDF are calculated in logarithmically scaled:

$$TF(t_i, d_j) = \log \frac{x_i}{N}$$

$$IDF(t_i, D) = \log \frac{K}{C_i}$$

Where $i \in [1, M]$ and $j \in [1, K]$. Then, TFIDF is the product of TF and IDF:

$$TFIDF(t_i, d_j) = TF(t_i, d_j) \times IDF(t_i, D)$$

3.3.2 N-grams

N-grams are continuous chunks of n items from a tokenized sequence for a document. Especially, uni-grams are terms where $n = 1$. Bigrams are pairs of adjacent terms where $n = 2$. Trigrams and quad-grams are three and four continuous terms, respectively. For example, the sentence “Your destination is 3 miles away” is tokenized into “your”, “destination”, “is”, “3”, “miles”, and “away”, where each term is a unigram. The bigrams are two-term strings: “your destination”, “destination is”, “is 3”, “3 miles”, and “miles away”. Trigrams are three-term strings: “your destination is”, “destination is 3”, “is 3 miles”, and “3 miles away”. And quadgrams are four-term strings: “your destination is 3”, “destination is 3 miles”, and “is 3 miles away”. In our experiment, we use unigrams, bigrams, trigrams, and quadgrams to calculate the correlated TF and TFIDF features.

3.3.3 Naive Bayes Classifier

The naive Bayes classifier is a classifier based on Bayes' Theorem:

$$P(A|B) = \frac{P(B|A) \times P(A)}{P(B)}$$

Where A and B are two conditions. The naive Bayes classifier takes each semantic feature as a condition and classify the samples with the highest occurring probability. Note that it assumes that the semantic features are independent [8].

3.3.4 Random Forest Classifier

A decision tree is a “tree” where different conditions branch off from their parents and each node represents a class for classification. The random forest classifier is an ensemble method that operates a multitude of decision trees and thus improves the accuracy. We adjust parameters such as *max depth*, *min samples split*, *n estimators*, and *random state* to achieve the best performance; where *Max depth* is the maximum depth of a decision tree; *Min samples split* is the minimum amount of samples to split an internal node, and *N estimators* is the number of decision trees in the random forest [2].

3.4 GLOVE

GloVe is an unsupervised learning algorithm that parallels the closeness of two words with their distance in a vector space [7]. The generated vector representations are called word embeddings. We use word embeddings as semantic features in addition to n-grams because they represent the semantic distances between the words in the context.

3.5 RNN

Recurrent neural networks (RNNs) utilize “memory” to process inputs and are widely used in text generation and natural language processing [6]. Long short-term memory (LSTM) is a RNN architecture that uses “gates” to “forget” the input at a condition. In our model we use 100 LSTM cells in one layer and a softmax activation function. We trained the model with 22 epochs and a batch size of 64.

4. EXPERIMENTAL RESULTS

This section presents the experimental results using naive bayes, random forest, and RNN with six groups of features: TF, TFIDF, frequency of bigrams, trigrams, and quadgrams, and GloVe word embeddings.

	TF	TFIDF	Bigram	Trigram	Quadgram	GloVe
Naive Bayes	88.08%	89.90%	90.77%	90.06%	89.74%	N/A
Random Forest	89.03%	89.34%	95.66%	94.71%	89.60%	N/A
RNN	N/A	N/A	N/A	N/A	N/A	92.70%

Table 1 shows the accuracy using each method. Observe that random forest results in better accuracy than the naive bayes classifier with TF, bigrams, trigrams, and quadgrams. Meanwhile,

bigrams outperform TF, TFIDF, trigrams, and quadgrams. The RNN model with GloVe features outperform TF, TFIDF, and quadgrams but not bigrams and trigrams.

Note that unigrams represent words; bigrams represent words and their one-to-one connections; trigrams carry level-two connections for words if consider a one-to-one connection between two words as level-one. As a result, bigrams carry more information than unigrams; trigrams more than bigrams; and quadgrams more than trigrams. Also, more information for training suppose to provide better accuracy. Therefore, we would expect quadgrams to result in higher accuracy than trigrams, trigrams higher than bigrams, and bigrams higher than unigrams. This assumption contradicts the data displayed in Table 1 as quadgrams do not result in the highest accuracy in the table. The reason is when information increases, the training process takes specific details and the model is “over-fitted”, when a model predicts the training set too well that it impairs its ability to classify examples that are not within its training set. However, bigrams do outperform unigrams because they carry more information. For the same reason, TF and TFIDF result in similar accuracies because they both are derived from unigrams. Meanwhile, GloVe with RNN outperforms unigrams and quadgrams but results in lower accuracy than bigrams and trigrams. This is caused by single layer LSTM cells and word embeddings represent unigrams. Therefore, the RNN model outperforms random forest if disregard the difference between word embeddings and unigrams.

Some implications of the success of these models are that they can be used by readers to filter through the content they consume to be wary of what articles may contain misinformation. These models can also be used by agencies, organizations, corporations, campaigns, or any other formal group to filter through news to find any false claims made about them or their actions. Additionally, publishing houses or news agencies can employ these methods in order to fact check pieces their writers compose to avoid any fake news from being produced under their name.

5. CONCLUSION

In this paper, we applied semantic features including unigram TF & TFIDF, bigrams, trigrams, quad- grams, and GloVe word embeddings along with naive bayes, random forest, and RNN classifiers to detect fake news. The performance is promising as bigrams and random forest achieved an accuracy of 95.66%. It implies that semantic features are useful for fake news detection. As the next step, semantic features may be combined with other linguistic cues and meta data to improve the detection performance.

REFERENCES

- [1] Babayan, Davit. “Bitcoin Bulls Spreading Fake News About Golden Asteroid: Peter Schiff.” NewsBTC, NewsBTC, 29 June 2019, www.newsbtc.com/2019/06/29/bitcoin-bulls-spreading-fake-news-about-golden-asteroid-peter-schiff/.
- [2] Breiman, L. (2001). Random forests machine learning. 45: 5–32. View Article PubMed/NCBI Google Scholar.
- [3] Handley, Lucy. “Nearly 70 Percent of People Are Worried about Fake News as a 'Weapon,' Survey Says.” CNBC, CNBC, 22 Jan. 2018, www.cnbc.com/2018/01/22/nearly-70-percent-of-people-are-worried-about-fake-news-as-a-weapon-survey-says.html.
- [4] Kowalewski, J. (2017). The impact of fake news: Politics. Lexology.
- [5] Mitchell, Amy, et al. “Many Americans Say Made-Up News Is a Critical Problem That Needs To Be Fixed.” Pew Research Center's Journalism Project, Pew Research Center, 28 June 2019,

www.journalism.org/2019/06/05/many-americans-say-made-up-news-is-a-critical-problem-that-needs-to-be-fixed/.

- [6] Mikolov, T., M. Karafiát, L. Burget, J. Černocký, and S. Khudanpur (2010). Recurrent neural network based language model. In Eleventh annual conference of the international speech communication association.
- [7] Patel, S. (2017). Chapter 1 : Supervised learning and naive bayes classification - part 1 (theory). Medium.
- [8] Pennington, J., R. Socher, and C. Manning (2014). Glove: Global vectors for word representation. In Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), pp. 1532–1543.
- [9] Romano, A. (2018). Mark zuckerberg lays out facebook's 3-pronged approach to fake news.
- [10] Rubin, V. L., Y. Chen, and N. J. Conroy (2015). Deception detection for news: three types of fakes. In Proceedings of the 78th ASIS&T Annual Meeting: Information Science with Impact: Research in and for the Community, pp. 83. American Society for Information Science.
- [11] Shu, K., A. Sliva, S. Wang, J. Tang, and H. Liu (2017). Fake news detection on social media: A data mining perspective. ACM SIGKDD Explorations Newsletter 19(1), 22–36.
- [12] Soares, Isa, and Florence Davey-Atlee. “The Fake News Machine: Inside a Town Gearing up for 2020.” CNNMoney, Cable News Network, money.cnn.com/interactive/media/the-macedonia-story/.
- [13] Tacchini, E., G. Ballarin, M. L. Della Vedova, S. Moret, and L. de Alfaro (2017). Some like it hoax: Automated fake news detection in social networks. arXiv preprint arXiv:1704.07506.
- [14] Wang, W. Y. (2017). “liar, liar pants on fire”: A new benchmark dataset for fake news detection. arXiv preprint arXiv:1705.00648.