Sticky fake news: A study on Diffusion and Sentiment

Rakshitha Gnanasekaran, Rejnald Lleshi, Saranda Rama, Kanaga Vaishnava Lakshmi Jayasankar

Institute of WeSt, Universität Koblenz-Landau, Germany
{rgnanasekaran, rlleshi, srama, kjayasankar}@uni-koblenz.de

Abstract : We investigate work already done in the topic of fake news diffusion and sentiment analysis and in particular perform a study on the differential diffusion of fake political and gossip news stories distributed on Twitter as well as perform opinion mining on the social context of these news pieces. The data comprises FakeNewsNet, a versatile dataset containing fake and true news pieces. We find that fake political stories diffuse farther and significantly faster than fake gossip stories. However, when it comes to the total number of users involved (including followers) fake gossip stories seem to have higher numbers than fake political stories. We also found that the general sentiment of direct Tweets for political news is more negative whereas for gossip news it was neutral. This changed when examining the replies to these Tweets where political news had an overall neutral sentiment, whereas gossip news was related with a more positive sentiment. We interpret this as a possible enjoyment of fake gossip stories.

Keywords - Fake News, Gossip Story, FakeNewsNet, Cascade, VADER

1. Introduction

In today's era, social media has become a mainstream source of information largely owning to its low cost and easy to access nature. Thus, micro-blogs like Twitter or Facebook are utilized as a major source of news information. However, due to the absence of a regulation authority controlling the flow of information, the quality of the news on social media is generally lower than in the traditional news organizations (Shu, Mahudeswaran and Liu, 2019). Studies show that 65% of US adult population accesses news through their social media (Anderson and Caumont, 2014) while the time spend overall on these platforms continues to increase[1].

Needless to say, truth and accuracy in small and great matters is central to the proper functioning and prosperity of a great majority of useful human endeavors. This has been established since times immemorial but recent foundational theories of decision-making, cooperation, communication and markets stress such a view as well (Vosoughi, Roy and Aral, 2018). Defining what is true and false has thus become a common political strategy. In 2017 a spokesman for the German government stated that they were dealing with a phenomenon of a dimension that they had not seen before when referring to the proliferation of fake news (Ruchansky, Seo and Liu, 2017). Even our economies are not immune by this phenomenon and this can be best illustrated by the loss of $130 billion in stock after a false tweet claimed that, Barack Obama was injured in an explosion (Rapoza, 2017). Especially in the last US presidential elections we witnessed a growing epidemic of fake news. By the end, approximately 1 Million tweets were related to the "Pizza gate" conspiracy alone (Robb, 2017).

Therefore, it seems to be a matter of the utmost pertinent attention that fake news be a foremost topic in research so as to contain its negative consequences and find relevant ways of ameliorating them.

There has been extensive work on the topic of fake news. The most comprehensive study so far on fake and true news is that of Vosoughi (2018) where the dissipation of false and true news, their effects, novelty and sentiment analysis is studied. Deepak and Shu have performed a series of studies (Shu *et al.*, 2018; Shu, Mahudeswaran and Liu, 2019) on fake and true news as well studying their dissipation and sentiment. They have also compiled a comprehensive dataset named FakeNewsNet (Shu *et al.*, 2018), which contains news content, social context and dynamic information on two categories of fake news, namely political and gossip.

This study focuses on two renown categories of fake news: political and gossip. After examining various datasets (detailed in chapter 3) in the literature, we found that FakeNewsNet[2] was the most relevant for our purposes. Already labelled fake news stories are fetched from two fact-checking websites, namely PolitiFact[3] for political and Gossip Cop[4] for gossip stories. Thereafter we study the differential diffusion of these two categories by employing cascades (Vosoughi, Roy and Aral, 2018) and peek into the timestamp distribution as well. Finally, we turn our attention to opinion mining. We pose the question whether or not people actually enjoy fake news content. How is it possible, for example for sly politicians to employ fake news so effectively if people do not enjoy it? Are they not exploiting people's feelings and emotions and appealing to their sentiments? Or maybe owning to the confirmation bias people accept what is fake without questioning it simply because it fits their worldview? How does this translate in the domain of gossip fake news? We do not find previous works (Vosoughi, Roy and Aral, 2018; Shu *et al.*, 2018) answering these questions. Therefore, we study the sentiment to see if it is overall positive. If yes, then this might be an expression of

[1] https://www.socialmediatoday.com/marketing/how-much- time-do-people-spend-social-media-infographic

[2] https://github.com/KaiDMML/FakeNewsNet

[3] https://www.politifact.com/

[4] https://www.gossipcop.com/

enjoyment. The other dimension of this study is diffusion of fake news. While Shu (2018) also performs a similar study, they are focusing more on a comparison between true and fake news and in relevant ways of discerning true from fake news content. This paper focuses on political and gossip news and in addition to the study of sentiment, we also study their diffusion. In this study we pose the following questions: Which kind (political or gossip) of fake news is more infectious? How much do people enjoy each?
This paper consists of six chapters. The first chapter deals with related work on various opinion mining techniques and argues why the chosen method is best for our study. The next section examines various datasets and elaborates more on the FakeNewsNet dataset. Chapter four explains how the research question is dealt with and chapter five discusses the results. Finally, we conclude the paper by a discussion, point out various limitations during our work and discuss future work.

2. Related Work

Here we examine various approaches on Opining Mining and elaborate on our chosen method. Opinion Mining, also called as sentiment analysis is a study of natural language processing which examines people's feedback, sentiments and opinions in blogs, tweets, Facebook posts and other social platforms. It is helpful in many fields like sociology, marketing and advertising, political science and psychology. Contents provided by micro-blogs like Facebook and Twitter pose major issues in sentiment analysis because of the shortness in text. Another possible major challenge is the increase in the volume of content. Hence it is necessity that a high adequate sentiment analysis tool is devised. Methodologies of studying sentiment are divided into three categories: Linguistic based, Lexicon based and Machine Learning. In this section we provide a brief overview on various sentiment analysis methodologies.

2.1 Linguistic Based

Linguistic study in fake news detection is defined as the study of an emotion and cognition experienced by liars that is reflected through non conscious word types, i.e., the study of structure and language of a context using methodologies like bag-of-word model and n-gram model (Conroy, Rubin and Chen, 2015). The bag-of-word model is the simplest method of data representation where each word is a single and equally significant unit. The N-gram model is a continuous sequence of n words from a text or a content (Broder, 1997). Cues of fake negatives statements are identified by summing up frequencies of words using n-gram or bag of word model (Larcker and Zakolyukina, 2012).
However, n-gram or bag of word model is not a good methodology to find sentiments in micro-blogs as content length and volume of content generated is large compared to other domains such as reviews or advertisement. Moreover, when using the bag-of-word model it would neglect the emotion of the content as it splits each word into a single significant unit. Additionally, in n-gram model no new vocabulary can be added to the model which is a major shortcoming in micro-blogs as the content of data keeps changing at a faster rate (Zhang and Ghorbani, 2019). Therefore, we avoid this kind of methodology for our research.

2.2 Lexicon Based

The lexicon-based approach is a way of excerpting sentiment from a content. According to Taboada (2011), lexicon based method is "a measure of subjectivity and opinion in text. It usually captures an evaluative factor (positive or negative) and potency or strength (degree to which the word, phrase, sentence, or document in question is positive or negative)." The lexicon-based approach is divided into two sub-categories: polarity based, and valence based.

2.2.1 Polarity based

Polarity-based approaches provide a single value as an output: either the content has a positive emotion or otherwise it has a negative emotion. In these lexicon-based approaches, the dictionaries to which the words are compared are created manually. Automatic dictionaries are also complied by adding seeds words to the list of existing dictionaries. Lemmatization is performed on the content and the adjectives are matched to the dictionary scores. Finally, an aggregated single score is commuted for a content (Taboada *et al.*, 2011). The popular lexicon-based approaches are Linguistic Inquiry and Word Count (LIWC) as well as General Inquiry (GI). However, studies have found that lexicon-based approaches are domain oriented and they don't consider acronyms and initialisms. Moreover, other important attributes of expressing emotions (e.g., slag) are not taken into consideration. Above all, polarity-based approach doesn't perform so well in sentiment intensity. For example: "The food here is exceptional" is evaluated with the same intensity as "The food here is okay", though we can understand that 'exceptional' has more sentiment than 'okay' (Hutto and Gilbert, 2014). For these reasons we also didn't employ polarity-based methodologies in our research.

2.2.2 Valence based

Valence based methods on the other hand give both a positive and negative score. Some valence-based methods are SentiStrength, Valence Aware Dictionary for sentiment Reasoning (VADER), and Affective Norm of English Word (ANEW). These valence-based methodologies satisfy all the drawbacks of polarity-based approaches. Another advantage of valence-based approaches is that they also account for words having extra letters or punctuation. For example: 'happpyyy!!!!' will have a higher score than 'happy!!' which will have a higher score than 'happy'. SentiStrength methodology has a downside in that they cannot understand sarcasm or jokes. While ANEW is a method that provide three types of score: Arousal, Dominance and Pleasure. For example: the valence for the word "betray" would be 1.68, "bland" 4.01, "dream" 6.73, and "delight" 8.26. However ANEW is insensitive to common sentiment related lexical feature (Hutto and Gilbert, 2014). On the other hand, VADER is sensitive to sentiment expression and provide good score for sarcasm, abbreviations, and slags. Researchers have found that VADER even outperforms human raters (F1 score = 0.84) with an accuracy of 0.84 and F1 score of 0.96 . (Hutto and Gilbert, 2014)

2.3 Machine Learning:

It is possible to use Naïve Bayes Classifier, Support Vector Machine or Logistic Regression for performing sentiment analysis. However, to perform Machine Learning techniques it is necessary to have a large training dataset. There is a scarcity of dataset as detecting fake content in social blogs is a tedious process, as fake content is written intentionally to mis-guide people (Shu *et al.*, 2018). Machine learning algorithms also requires a large computational power and they are domain specific, for instance if you train the algorithm with news content they won't provide a good result with gossip content (Hutto and Gilbert, 2014).

Considering our dataset (see chapter 3, Data Set) and what was discussed so far, we therefore choose VADER[5] to be an optimal method for our research. VADER can score emoticons, has different intensity for multiple letters in a word. as well as all caps words, can understand jokes and sarcasm, works well with slag, acronyms, initialisms and abbreviations and is moreover specifically attuned to social media which is exactly what we need. Finally, VADER outperforms polarity based methods, machine learning and even humans. (Hutto and Gilbert, 2014).

3. Data Set

Fake news content analysis is a growing field in research and in-order to study fake content analysis there is a requirement of a comprehensive multidimensional dataset. There are several such existing datasets for the study of fake news content like BuzzFeedNews[6], LIAR[7], BS Detector[8], CREDBANK[9], BuzzFace[10], Facebook Hoax[11]. However, most of them contain only one or two aspects/dimensions (linguistic or social features). In this chapter the popular datasets are compared with the chosen dataset (FakeNewsNet).

BuzzFeedNews: This dataset comprises of 1627 article which contains news article from September 19th to 23rd and September 26th and 27th (a week during US presidential election of 2016). These posts broadcasted in Facebook from 9 news agencies, were fact- checked by 5 BuzzFeed journalists. This dataset comprehends 826 mainstream, 356 left-wing and 545 right wing articles. However, BuzzFeedNews dataset only includes linguistic news content. Its dataset does not provide information on user's profile like the time and location in which their profile was created, or on the number of followers or followees they have, or on the post content. And since we utilize these dimensions of social context for our study, we have therefore avoided this dataset.

LIAR: This dataset comprises about 12.8 thousand linguistic news articles, which are fact-checked through the fact-checking website Politifact. Moreover, these articles are human labelled into 6 categories ranging from completely true to completely false (Wang, 2017). However, this dataset misses a lot of useful social context which is necessary for our study.

BS Detector: BS Detector comprises of linguistic news content which are collected by using a browser extension, bs detector, which checks all the links in a webpage for unreliable sources. These unreliable sources are compiled manually with a list of domains. Similar to BuzzFeedNews and LIAR, BS Detector also lacks social context.

CREDBANK: This data repository is a collection of 60 million tweets for 96 days, which are related to 1000 news events fact-checked using 30 annotators from Amazon Mechanical Turk (Mitra and Gilbert, 2015). CREDBANK provides data about linguistic news content, information of user's profile, their post and their followers and followees count. However, it lacks data on retweets which we use for sentiment analysis on our study.

BuzzFace: This dataset is an extension of BuzzFeedNews, where the comments in Facebook are collected. BuzzFace comprises of 2263 articles with 1.6 Million comments on these articles (Santia and Williams, 2018). Like CREDBANK, BuzzFace provide linguistic news article, user's profile and post. They also provide information on second order i.e. their likes, their comments and their retweets.

FacebookHoax: This linguistic news content, data repository contains 15,500 posts from 32 pages which has more than 2,300,000 likes from Facebook collected by Facebook Graph API. It comprises of non-hoax article, which are scientific and hoax article, which are conspiracy (Tacchini *et al.*, 2017). FacebookHoax also provides information on user's profile and their post along with their likes and comments. It misses however information on friends and followers.

FakeNewsNet: It is a multidimensional dataset consisting of fake and true stories in two categories, political and gossip. True and Fake News are collected from fact-checking websites, namely PolitiFact and Gossip Cop. Moreover, it also collects the social engagement of the users and thus this dataset is rich in features, having information on user's behavior towards each article like number of retweet, number of likes and comments, users' profile, their timeline, their network - followers and followees count as well spatial-temporal information - the time and location in which the user's account is created. In the context of our research, we only collected fake news from

[5] https://github.com/cjhutto/vaderSentiment

[6] https://github.com/BuzzFeedNews/

[7] https://github.com/thiagorainmaker77/liar_dataset

[8] https://github.com/selfagency/bs-detector

[9] http://compsocial.github.io/CREDBANK-data/

[10] https://github.com/gsantia/BuzzFace

[11] https://github.com/gabll/some-like-it-hoax

PolitiFact and Gossip Cop using the FakeNewsNet repository. The data comprise about 5800 news stories tweeted by approximately 300,000 people more than 650,000 times. (Shu *et al.*, 2018)

	Linguistic (news content)	User	Network (followers and followees count)	Second Order (likes, retweet, comments)	Post
BuzzFeedNews	Yes	No	No	No	No
LIAR	Yes	No	No	No	No
BSDetector	Yes	No	No	No	No
CREDBANK	Yes	Yes	Yes	No	No
BuzzFeed	Yes	No	No	Yes	Yes
FacebookHoax	Yes	Yes	No	Yes	Yes
FakeNewsNet	Yes	Yes	Yes	Yes	Yes

Figure 1 Comparison of dataset

From Figure 1 we can see the comparison of the abovementioned datasets. Research question one, which studies which of the fake news, political or gossip dissipates more, needs data about user, network and second order. Moreover, FakeNewsNet also has spatial temporal information which enables us to perform a timeline analysis. And research question two, which performs sentiment analysis on the social context needs linguistic data, second order and post (content of tweets/retweets). Looking at the various dimensions of the datasets from Figure 1 and considering what was discussed so far, it is clear that FakeNewsNet is our best option.

Lastly, for the purpose of studying sentiment analysis we independently collected the replies of each tweet (these data were not part of the FakeNewsNet dataset). This was done using NASTY[12] which is a library for retrieving Tweets via the Twitter Web UI.

4. Methods

We study the diffusion of fake political and gossip news on two levels: first by studying cascades and then by performing a study on the timestamp/timeline. To this end we utilize a comprehensive dataset that includes ~5800 news story pieces, tweeted by ~300,000 people more than 650,000 times.

A cascade can be started by an individual by tweeting something about one of the 5800 pieces of news articles that we study. Another individual could yet start another cascade by tweeting something else about that same news piece. Thus, these two tweets would represent two cascades of the same news piece. The smallest size of a cascade is one (this means that no one else retweeted the tweet). So, the number of cascades pertaining to a news piece is equal to the number of times the news piece was independently tweeted by a user (not, however retweeted) (Vosoughi, Roy and Aral, 2018).

For example, if news piece "A" is tweeted by 10 people separately, but not retweeted, it would then have 10 cascades each of size one. But if a news piece "B" is independently tweeted by 10 people and each of those tweets is retweeted 100 times, then the news piece would have 10 cascades each of size 100.

We therefore quantify the cascades' size (the number of users involved in the cascade over time, i.e. number of retweets), maximum breadth (the maximum number of users involved in the cascade, i.e. including -unique- followers) and overall the total number of cascades (Vosoughi, Roy and Aral, 2018).

Second, we perform a timeline analysis of both political and gossip fake news in order to examine which category spreads faster.

We then performed opinion mining by analyzing the tone of replies on two levels:

- On the direct tweets of each news piece.
- On the replies to these tweets.

Our hypothesis is that if the tone of the replies to the tweets connected to a particular news piece is overall positive, then users might be enjoying the fake news piece. We obtain the content of the tweets from FakeNewsNet dataset and their replies using NASTY. Then using VADER, we calculate the ratio of positive, negative and neutral sentiment on each of the two levels.

5. Discussion

A greater fraction of fake gossip stories experienced between 1 and 1000 cascades (Figure 2, Number of Cascades), whereas a greater fraction of fake political stories experienced more than 1000 cascades. Moreover, fake political stories experienced more cascades than fake gossip stories.

Except for two outliers (particularly significant fake Gossip stories, one of them concerning Kim Kardashian), fake political stories experienced a greater cascade size than fake gossip stories (Figure 3, Cascade Size). This means that more users were generally involved with fake political stories compared to gossip stories. However, when considering the maximum-breadth (Figure 4), which also includes followers (since when a user tweets or retweets an article it will appear on the timeline of his followers as well) fake Gossip stories reached more people than fake Political stories. This might be due to the fact than more influencers and celebrities are participating in gossip than in political news.

When timeline was analyzed (Figure 5) we found that fake political stories diffused significantly faster in particular within the first 24-hour interval.

In sentiment analysis, when considering the content of the tweets alone, it seems that fake political news elicits a more negative response whilst fake gossip news elicits a rather neutral response. One explanation given for this is that there is a greater difficulty in identifying the fake news in gossip by common people than there is with political news (Shu *et al.*, 2018).

[12] https://github.com/lschmelzeisen/nasty

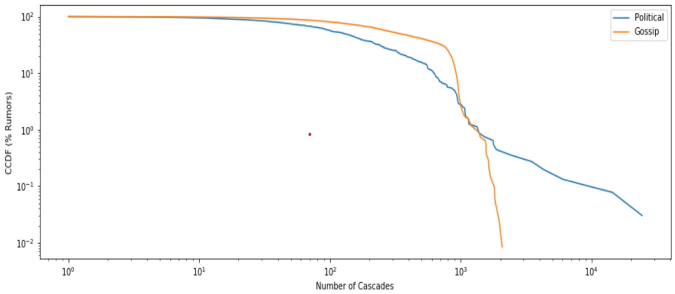

Figure 2 Number of Cascades

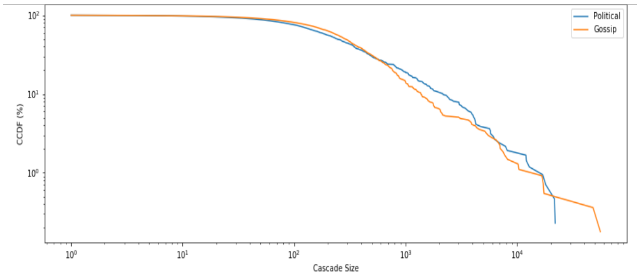

Figure 3 Cascade Size

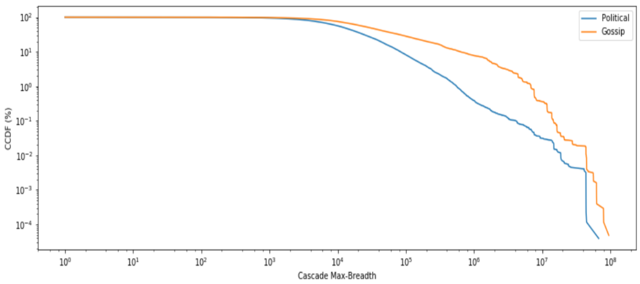

Figure 4 Maximum Breadth

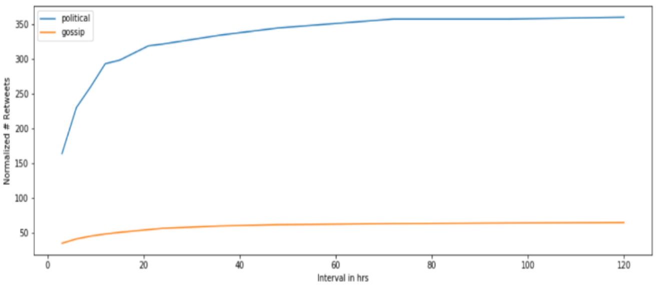

Figure 5 Timeline analysis

We also performed sentiment analysis on the replies of Tweets using the same technique. If a Tweet comment has for example the sentiment distribution of [5, 5, 5] it occurs in the middle of the triangle in Figure 6. Here we observe that fake political news pieces elicit a neutral sentiment, but fake gossip news pieces interestingly elicit a positive sentiment. We also computed a normalized weighted composite score (single unidimensional measure of sentiment) of the whole corpus of Tweet's replies using the threshold of 0,05 for classifying the sentiment as overall positive (Hutto and Gilbert, 2014). These Tweet comments scored a 0.1, which is double the threshold needed to classify it as containing positive sentiments. This could mean that the users engaging with the original Tweet, which points to a fake news piece, do enjoy this fake gossip story at some level at least.

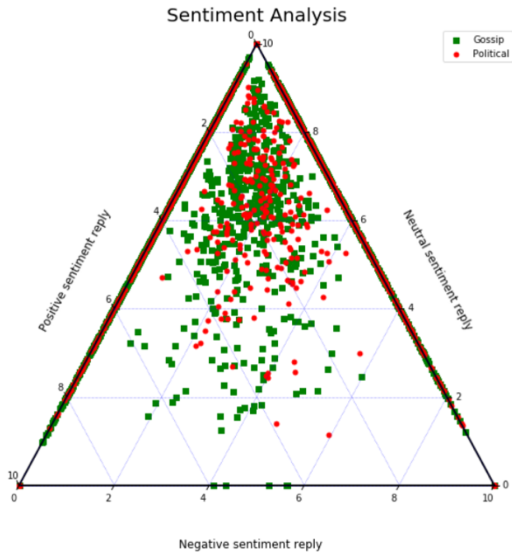

Figure 6 Sentiment Analysis

6. Conclusion

When we analyzed the diffusion dynamics of fake political and gossip news, we found that fake political news dissipates farther, and significantly faster than fake gossip news. On the other hand, fake gossip news spread slightly deeper than fake political news in that its maximum breadth is larger, it overall infects more users. When performing opinion mining, we got two sets of results. First, we studied the direct tweets that point directly to the news piece. As was claimed by Shu (2018) we also observed that political news elicited more of a negative sentiment while gossip news elicited a mostly neutral response. Second, we analyzed the sentiment of the comments on each of the tweets. In this case political replies had a neutral sentiment whilst gossip replies interestingly showed an overall positive sentiment. This might mean that people are enjoying it, when they are engaging with gossip news even though they are fake. Thus, one reason as to why fake gossip news dissipates might be that people enjoy it.

6.1 Limitations

Even though we meticulously tried to pick the best dataset for our research, during the course of it, it became apparent that the dataset had started to decay. In some cases, we were not able to properly harvest the social content related to the news pieces because it had either been deleted by the original poster or Twitter had removed it (Twitter has policies against fake news and tries to fight it). The original creators of the dataset had speculated upon this and therefore claimed to keep it updated but as of this time this is not yet the case. The results would have been better had we had fresh content to work upon.

6.2 Possible Future Work

The first takeaway from all this work surely is that research into fake news is fascinating topic. In the end one emerges with more questions when compared to the start. In the context of this research we examined a lot of academic work on the topic and although we do not want to judge the papers themselves here, we were left out with a little bit of disappointment: almost all of them seemed just to scratch the surface. And yet there is a plethora of question that one might pose on the topic of fake news.

Why is it that fake news spreads consistently deeper, faster and further than true news? Is it not that people actually enjoy fake news? If so, then why should people enjoy them? Is it not that fake news appeals to our rather primitive and crude senses much more than true news? What might be the role of confirmation bias in fake news? What is the relationship between Schadenfreude and fake news? Might one not use Schadenfreude as a catapultier for fake news spreading in the same way as our cognitive biases are exploited every day by almost every company for the purpose of boosting profit alone? Indeed, one of the reasons that helped Trump in the last presidential elections was the amount of Schadenfreude that he generated early on (which in the long term profited him alone of course).

Therefore future research could easily establish some Twitter Streaming APIs, gathering the content of renowned politicians like Trump, Johnson, Merkel etc. and then study in part the question posed above. This would be a very interesting avenue of research in the future.

7 References

1. Anderson, M. and Caumont, A. (2014) *How social media is reshaping news*, *Pew Research Center*.
2. Broder, A. Z. (1997) 'Syntactic clustering of the Web', *Computer Networks*. doi: 10.1016/s0169-7552(97)00031-7.
3. Conroy, N. J., Rubin, V. L. and Chen, Y. (2015) 'Automatic deception detection: Methods for finding fake news', *Proceedings of the Association for Information Science and Technology*, 52(1), pp. 1–4. doi: 10.1002/pra2.2015.145052010082.
4. Hutto, C. J. and Gilbert, E. (2014) 'VADER: A parsimonious rule-based model for sentiment analysis of social media text', in *Proceedings of the 8th International Conference on Weblogs and Social Media, ICWSM 2014*.
5. Larcker, D. F. and Zakolyukina, A. A. (2012) 'Detecting Deceptive Discussions in Conference Calls', *Journal of Accounting Research*. doi: 10.1111/j.1475-679X.2012.00450.x.
6. Mitra, T. and Gilbert, E. (2015) 'CREDBANK: A large-scale social media corpus with associated credibility annotations', in *Proceedings of the 9th International Conference on Web and Social Media, ICWSM 2015*.
7. Rapoza, K. (2017) *Can 'Fake News' Impact The Stock Market?*, *Website*. Available at: https://www.forbes.com/sites/kenrapoza/2017/02/26/can-fake-news-impact-the-stock-market/#1cba88cb2fac.
8. Robb, A. (2017) *Anatomy of a Fake News Scandal*, *Website*.
9. Ruchansky, N., Seo, S. and Liu, Y. (2017) 'CSI: A hybrid deep model for fake news detection', in *International Conference on Information and Knowledge Management, Proceedings*. doi: 10.1145/3132847.3132877.
10. Santia, G. C. and Williams, J. R. (2018) 'Buzz face: A news veracity dataset with facebook user commentary and egos', in *12th International AAAI Conference on Web and Social Media, ICWSM 2018*.
11. Shu, K. *et al.* (2018) 'FakeNewsNet: A Data Repository with News Content, Social Context and Spatialtemporal Information for Studying Fake News on Social Media'. Available at: http://arxiv.org/abs/1809.01286.
12. Shu, K., Mahudeswaran, D. and Liu, H. (2019) 'FakeNewsTracker: a tool for fake news collection, detection, and visualization', *Computational and Mathematical Organization Theory*. doi: 10.1007/s10588-018-09280-3.
13. Taboada, M. *et al.* (2011) 'Lexicon-basedmethods for sentiment analysis', *Computational Linguistics*. doi: 10.1162/COLI_a_00049.
14. Tacchini, E. *et al.* (2017) 'Some like it Hoax: Automated fake news detection in social networks', in *CEUR Workshop Proceedings*.
15. Vosoughi, S., Roy, D. and Aral, S. (2018) 'News On-line', *Science*, 1151(March), pp. 1146–1151.
16. Wang, W. Y. (2017) '"Liar, liar pants on fire": A new benchmark dataset for fake news detection', in *ACL 2017 - 55th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference (Long Papers)*. doi: 10.18653/v1/P17-2067.
17. Zhang, X. and Ghorbani, A. A. (2019) 'An overview of online fake news: Characterization, detection, and discussion', *Information Processing and Management*. doi: 10.1016/j.ipm.2019.03.004.

Appendix

Link to Google Doc which carefully documents each step of the work leading to this paper:
https://docs.google.com/document/d/1qIWrO0DgFlqohDWPyBitfM78L9bFpz_hUU5UZnmTuLw/edit

GitHub repository containing all relevant code: https://github.com/rlleshi/StickyFakeNews