
CoAID: COVID-19 HEALTHCARE MISINFORMATION DATASET

Limeng Cui Dongwon Lee

The Pennsylvania State University
University Park, PA 16802

{lzc334, dongwon}@psu.edu

ABSTRACT

As the COVID-19 virus quickly spreads around the world, unfortunately, misinformation related to COVID-19 also gets created and spreads like wild fire. Such misinformation has caused confusion among people, disruptions in society, and even deadly consequences in health problems. To be able to understand, detect, and mitigate such COVID-19 misinformation, therefore, has not only deep intellectual values but also huge societal impacts. To help researchers combat COVID-19 health misinformation, therefore, we present **CoAID** (Covid-19 healthcare mIsinformation Dataset), with diverse COVID-19 healthcare misinformation, including fake news on websites and social platforms, along with users' social engagement about such news. CoAID includes 4,251 news, 296,000 related user engagements, 926 social platform posts about COVID-19, and ground truth labels. The dataset is available at: <https://github.com/cuilimeng/CoAID>.

Keywords COVID-19, misinformation, fake news, benchmark, dataset, machine learning

1 Introduction

COVID-19 is believed to be caused by a novel coronavirus called SARS-CoV-2, which was initially discovered in Wuhan, China in December 2019, and has quickly spread throughout the world. WHO declared the outbreak a Public Health Emergency of International Concern on January 30, 2020 [16] and characterized COVID-19 as a pandemic on March 11, 2020 [18]. As of May 6, 2020, more than 3.5 million cases of COVID-19 and almost 250,000 deaths worldwide have been reported to WHO [17]. Common symptoms of COVID-19 includes cough, shortness of breath, fever, sore throat, and loss of taste or smell [7]. The incubation period is estimated to extend to 14 days, with a median time of 5.1 days [11]. Mortality for COVID-19 (3-4%) appears higher than that for seasonal influenza (<0.1%) [15]. Until now, there has been no known vaccine or specific cure for COVID-19 [14].

While COVID-19 situation has gotten severely worse, misinformation related to COVID-19 has rapidly risen and caused serious social disruptions as well. One the one hand, fake cures for COVID-19 have seriously threatened people's lives. For example, an Arizona man was dead and his wife was hospitalized after the couple ingested a form of Chloroquine to prevent COVID-19¹. On the other hand, the prevalent misinformation is disrupting social order. For example, 77 cell phone towers have been set on fire due to the conspiracy that 5G mobile networks can spread COVID-19².

In recent years, in mitigating these problems, many computational methods have been developed to auto-detect diverse types of misinformation in different genres [20, 26, 23, 3]. However, debunking COVID-19-related misinformation exhibits its own set of unique challenges. First, aided by the fear of the unknown, misinformation about emerging diseases can spread more quickly than ordinary misinformation does before being debunked³. Second, when one keeps

¹<https://www.npr.org/sections/coronavirus-live-updates/2020/03/24/820512107/man-dies-woman-hospitalized-after-taking-form-of-chloroquine-to-prevent-covid-19>

²<https://www.businessinsider.com/77-phone-masts-fire-coronavirus-5g-conspiracy-theory-2020-5>

³<https://www.jhunewsletter.com/article/2020/04/opposing-viewpoints-how-the-media-contributes-to-misinformation-in-crisis>

Table 1: Comparison with existing misinformation datasets

	Information Type			User Engagement	
	Claim	News Article	Post	Tweet	Reply
LIAR	-	✓	-	-	-
FA-KES	-	✓	-	-	-
FakeNewsNet	-	✓	-	✓	✓
FakeHealth	-	✓	-	✓	✓
CoAID	✓	✓	✓	✓	✓

seeing a piece of fake news in the news feed, she tends to think that it is true, even if she had doubts before⁴. Third, after people were convinced by misinformation, “myths correction message” from authority may be ineffective or even reduce people’s beliefs in other facts about an epidemic [2]. Therefore, it is not clear how easy or difficult for researchers to build satisfactory computational models to auto-detect COVID-19-related misinformation.

To aid these computational efforts, we have constructed a benchmark dataset, named as CoAID (Covid-19 healthcare Misinformation Dataset). CoAID includes confirmed fake and true news articles from the fact-checked or reliable websites and posts at social platforms. We also conducted quick data analysis to demonstrate the utility of CoAID. We contrasted the distinctive features between misinformation and facts about COVID-19, and ran several state-of-the-art misinformation detection methods on CoAID to show where solutions currently lie. Our aim is to call out for public attention to the problems of COVID-19-related misinformation and work together to help develop accurate detection and deterrence of such misinformation.

2 Preliminaries

According to the Oxford English Dictionary, both *misinformation*⁵ and *disinformation*⁶ are either wrong or misleading information, but disinformation is spread “deliberately” while misinformation is not necessarily. Literature provides more fine-grained definitions of various examples of misinformation on social media, including rumor, fake news, hoax, satire, propaganda, and even conspiracy [22, 13], such that: (1) *fake news* is the presentation of fake or misleading claims as news, where the claims are misleading deliberately [8], (2) *hoax* is deliberately fabricated falsehood, with the intention to deceive a certain group of the population [8], (3) *rumor* is unverified but relevant information that arise in contexts such as danger or potential threat, that helps people make sense and manage risk [6].

In this work, we do not differentiate between misinformation and disinformation as it is virtually impossible to computationally determine one’s intention. We also interchangeably use fake news and misinformation in this work. CoAID comprehensively collects various misinformation examples across different platforms.

3 Related Work

The construction of fake news dataset aim to extract useful features that can better distinguish fake news from true ones. Researchers have proposed a series of methods to extract news and verify the truthfulness of the news. There exists several benchmark datasets for fake news detection, that contain the linguistic features of news.

LIAR [25] is collected from political fact-checking website PolitiFact. It has 12.8K human labeled short statements from PolitiFact API, and each statement includes statement content, speaker and context. The truthfulness of each statement is evaluated by an editor from PolitiFact, by using detailed judgments and fine-grained labels: *pants-fire*, *false*, *barely true*, *half-true*, *mostly-true*, and *true*.

FA-KES [21] includes the news articles around the Syrian war. The authors first identified the major events in the Syrian war from VDC (Violations Documentation Center), and then retrieved 804 news articles related to the events from the following several news outlets. The truthfulness of each article is justified by comparing the articles against the VDC dataset.

Dhoju et al. [5] proposed a method to collect a healthcare news dataset. They first identified reliable (WHO, CDC/NIH/NCBI, mainstream media outlets) and unreliable media outlets (Dr. Melissa Zimdars’s list, fact-checking

⁴<https://www.bbc.com/future/article/20200406-why-smart-people-believe-coronavirus-myths>

⁵<https://www.oed.com/view/Entry/119699?redirectedFrom=misinformation>

⁶<https://www.oed.com/view/Entry/54579?redirectedFrom=disinformation>

websit snopes.com), and then used news article scraping tool Newspaper3k⁷ to crawl the news articles. Finally they used Google Cloud Natural Language API to filter out non-health news articles.

Besides the news content, other researchers also collected user engagement features of the news, such as user engagement on online social media.

Ghenai et al. [9] collected 126 unproven cancer treatments that are scrutinized by an oncologist from medical sources DC Science, Wikipedia and PubMed. Topic keywords were generated from the treatments and used to find related tweets.

FakeNewsNet [24] is collected from fact-checking websites PolitiFact and GossipCop. News contents and the related ground truth labels are crawled from the two websites. And the author collected social engagement through Twitter’s Advanced Search API. The dataset includes news ID, URL, article title, tweet IDs of tweets sharing the news and scripts to crawl news contents and social engagement. It contains 1,056 articles from PolitiFact and 22,864 articles from GossipCop. Each article is labeled as fake or true.

FakeHealth [4] is collected from healthcare information review website Health News Review, which reviews whether a news is satisfactory from 10 criteria. Daily publication of new content on this website was ceased at the end of 2018. The authors crawled users’ replies, retweets and profiles by using article URL and article title by using Twitter API.

In addition, existing datasets often focus only on the news coverage or articles, and little attention has been paid to social media posts. However, fake social media posts can easily go viral on multiple social platforms, and cause serious social disruptions. For example, a fake video on YouTube claiming that a US patent on the coronavirus was applied in 2006, quickly went viral on multiple social media platforms⁸, causing much confusion and discomfort.

As shown in Table 1, CoAID dataset not only includes both true and fake news, but also has short claims and social media posts. Compared with news articles, a *claim* is much shorter—only one or two sentences—such as “Eating garlic prevents infection with the new coronavirus.” In this work, we term “myth” and “rumor” as “fake claim”, in order to contrast those with “true claim.” Our dataset collects all possible features of interests, including multi-modal information and user engagement data. In addition, CoAID can be updated automatically to fetch the latest news. The statistics of COVID-19 news dataset is shown in Table 2.

4 Dataset Construction

In this section, we introduce how we have collected healthcare misinformation related to COVID-19, their reliable ground truth labels, and associated user engagement features, and how we manage to update the dataset automatically.

4.1 Facts and Misinformation on Websites

We have collected both facts and misinformation related to COVID-19, including news articles and claims. The publishing dates of the collected information range from December 1, 2019 to September 1, 2020. The topics include: COVID-19, coronavirus, pneumonia, flu⁹, lock down, stay home, quarantine and ventilator.

- **News Article:** An example of news article is shown in Figure 1(a). We identified reliable media outlets and fact-checking websites and obtained URLs of misinformation (i.e., fake news) and facts (i.e., true news). To collect true news articles, we have crawled from 9 reliable media outlets that have been cross-checked as reliable, including: Healthline¹⁰, ScienceDaily¹¹, NIH¹² (National Institutes of Health), MNT¹³ (Medical News Today), Mayo Clinic¹⁴, Cleveland Clinic¹⁵, WebMD¹⁶, WHO¹⁷ (World Health Organization) and CDC¹⁸

⁷<https://newspaper.readthedocs.io/en/latest/>

⁸<https://factcheck.afp.com/false-claims-patents-fuel-novel-coronavirus-conspiracy-theories-online>

⁹We filtered out known flu types such as *Influenza A/B*, *bird flu* and *swine flu*.

¹⁰<https://www.healthline.com/health-news?ref=global>

¹¹https://www.sciencedaily.com/news/health_medicine/

¹²<https://www.nih.gov/news-events/news-releases>

¹³<https://www.medicalnewstoday.com/coronavirus>


¹⁴<https://www.mayoclinic.org/diseases-conditions/coronavirus/symptoms-causes/syc-20479963>

¹⁵<https://newsroom.clevelandclinic.org/>


¹⁶<https://www.webmd.com/news/articles>

¹⁷<https://www.who.int/news-room>

¹⁸<https://www.cdc.gov/coronavirus/2019-ncov/whats-new-all.html>



(a) A news article at the Healthline.com website



(b) Tweets



(c) Replies

Figure 1: Screenshots of (a) a news article at the Healthline.com website and its related user engagement, including (b) tweets and (c) replies of each tweet.




Figure 2: Screenshot of a social platform post.

(Centers for Disease Control and Prevention). We used the Selenium¹⁹ to parse the html structure of web pages to locate the URLs. To collect fake news, we retrieved the URLs from several fact-checking websites, including: LeadStories²⁰, PolitiFact²¹, FactCheck.org²², CheckYourFact²³, AFP Fact Check²⁴ and Health Feedback²⁵.

- **Claim:** To collect claims (of one or two sentences), we referred to the WHO official website²⁶, WHO official Twitter account²⁷ and MNT²⁸. We specifically separated true and fake claims. For example, “Only older adults

¹⁹<https://www.selenium.dev/>

²⁰<https://leadstories.com/hoax-alert/>

²¹<https://www.politifact.com/coronavirus/>

²²<https://www.factcheck.org/fake-news/>

²³<https://checkyourfact.com/>

²⁴<https://factcheck.afp.com/>

²⁵<https://healthfeedback.org/>

²⁶<https://www.who.int/>

²⁷<https://twitter.com/WHO>

²⁸<https://www.medicalnewstoday.com/articles/coronavirus-myths-explored>

Table 2: Statistics of CoAID Version 0.3

	Fake		True	
	Claim	News Article	Claim	News Article
# Information on Website	28	204	454	3,565
# Tweets	484	10,439	8,092	141,652
# Replies	626	7,436	12,451	114,820
# Social Platform Posts	650 (880)	-	42 (46)	-
# Total	1,788 (2,018)	18,079	21,039 (21,043)	260,037

and young people are at risk” is a fake claim, while “5G mobile networks DO NOT spread COVID-19” is a true claim.

After we obtained all URLs to true and fake news related to COVID-19, we used the Newspaper3k²⁹ to fetch their corresponding title, content, abstract, and keywords. In total, we have obtained 204 fake news articles, 3,565 true news articles, 28 fake claims and 454 true claims.

4.2 User Engagement

We used Twitter API to fetch user engagement data of both facts and misinformation. Figure 1(b) shows user tweets related to the news in Figure 1(a), and Figure 1(c) shows user replies under a tweet.

- **Tweets:** We used the titles of news articles as the search queries, and specified start and end dates to get the tweets discussing about the news in question in a certain period. The retrieved user engagement features include: user ID, tweets, replies, favorites, retweets, and location. In total, we obtained 10,439 tweets about fake articles, 141,652 tweets about true news articles, 484 tweets about fake claims and 8,092 tweets about true claims, which is summarized in Table 2.
- **Replies:** We further obtained users’ replies of each tweet, by using tweet IDs. In total, we obtained 7,436 replies under fake articles, 114,820 replies under true news articles, 626 tweets about fake claims and 12,451 tweets under true claims, which is summarized in Table 2.

4.3 Social Platform Posts

In addition to traditional mass media, on social media, individual users act as the producers of user generated contents, forming the so-called self media. An example of such self-media is shown in Figure 2. As both true and fake news spread through self-media as well, we have also collected both fake and true posts originated from five well-known social media platforms—Facebook, Twitter, Instagram, Youtube, and TikTok—also fact-checked by: Lead Stories, PolitiFact, FactCheck.org, CheckYourFact, AFP Fact Check, and Health Feedback. We only count once for each distinctive fake post. For example, if a post was originally posted by a Facebook user and then posted again by another user on Twitter, it is counted as 1 (i.e., no duplicate fake posts). The numbers in parentheses in Table 2 represent the total number of posts, before duplicates were removed.

Per each post, we only crawled its title without news content, and listed the posts under “Claim” in Table 2.

4.4 Automatic Updates

CoAID can be updated automatically with the latest news and tweets. We record the timestamp of the most recently added data each time, and use it as the start date for the next search.

In summary, we list the description of the extracted features in Table 3. Note that “Title” is the title of the related fact-checking article, while “Article title” is the title of the article crawled. Due to the limitation under Twitter’s Terms & Conditions³⁰, only tweet IDs are included in the released dataset of CoAID.

²⁹<https://newspaper.readthedocs.io/en/latest/>

³⁰<https://developer.twitter.com/en/developer-terms/agreement-and-policy>

Table 3: Feature descriptions of CoAID

Type	Features
Information on Websites	ID, Fact-checking URL, Information URLs, Title, Article title, Content, Abstract, Publish date, Keywords
User engagement: Tweets	ID, Tweet ID
User engagement: Replies	ID, Tweet ID, Reply ID
Social Platform Posts	ID, Fact-checking URL, Post URLs, Title




Figure 3: Sentiments of tweets related to fake and true news articles

5 Data Analysis

CoAID includes multi-modal news, related ground truth labels, and user engagement. The detailed statistics of the dataset are shown in Table 2. In this section, we perform preliminary data analysis in order to illustrate the features of CoAID. Then, we perform several fake news detection methods to evaluate the challenges of COVID-19 healthcare misinformation detection. The following data analysis is based on CoAID Version 0.1 in Table 5.

5.1 Prescriptive Analysis

We use VADER [10] to analyze users’ sentiments in tweets related to fake and true news articles. We exclude the tweets that are completely neutral, and plot the sentiment scores of tweets related to fake articles in Figure 3(a) and true news articles in Figure 3(b), respectively. The bars show the numbers of tweets with different negative and positive degrees, while the scatter plot shows the negative and positive scores. The intensity of the hive represents the density of data points. We can see from Figure 3 that tweets related to fake news are more negative, and have stronger sentiment polarities than those related to true news articles.

Next, we analyze the top 30 frequent hashtags in tweets related to fake and true news articles, respectively. For a more intuitive view, we delete the most frequent hashtags “#coronavirus”, “#covid_19”, “#covid19” and “#covid” as they appear almost in every tweet. We show the frequency of hashtags in tweets related to fake and true news articles in Figure 4(a) and Figure 4(b), respectively. We find that the hashtag distributions of tweets about fake and true news articles are quite different. While the hashtags in tweets about true news articles are mainly related to healthcare, those in tweets about fake news cover more diverse topics, including conspiracy (#bioweapon) and fake cure (#vitaminc).




Figure 4: Frequency of hashtags in tweets about fake and true news articles





Figure 5: Frequency of common claims


In Figure 5, next, we list ten most common claims of COVID-19 on Twitter, and their corresponding frequencies. We merge similar claims into one, such as “Eating garlic prevents infection with the new coronavirus” and “Garlic protects against coronaviruses.” We can see that the most frequently discussed claim is “COVID-19 is just like the flu”, followed by “COVID-19 is airborne” and “The new coronavirus can be transmitted through mosquito bites”. We then plot the daily counts of tweets related to two claims, ‘COVID-19 is just like the flu’ and “5G mobile networks spread COVID-19” in Figure 6(a) and 6(b), respectively. We can see that there were most tweets about “COVID-19 is just like the flu” at the beginning of the global outbreak, which is around March 12, 2020. Note that this topic is mentioned almost everyday. The tweets about “5G mobile networks spread COVID-19” were mostly on the April 10, 2020, when 5G towers were being set on fire, but afterward the number declined as the claim was debunked.

5.2 COVID-19 Misinformation Detection

To demonstrate the main utility of m , here, we conduct comparative experiments on the misinformation detection task. We include both simple methods and state-of-the-art methods as baselines. We consider the following baseline methods for COVID-19 misinformation detection:



(a) Daily counts of tweets related to “COVID-19 is just like the flu”



(b) Daily counts of tweets related to “5G mobile networks spread COVID-19”

Figure 6: Claim trends

- **SVM (Support Vector Machine)**: We use bag-of-words to represent texts, and feed the representations to a linear kernel SVM.
- **LR (Logistic Regression)**: We concatenate all the word embeddings together, and feed it to the model.
- **RF (Random Forest)**: We use bag-of-words to represent texts, and feed the representations to the model.
- **CNN**: We feed the word embeddings into a Convolution1D, which will learn filters. Then we add a max pooling layer.
- **BiGRU [1]**: BiGRU gets annotations of words by summarizing information from both directions for words.
- **CSI [20]**: CSI is a hybrid deep learning-based misinformation detection model that utilizes information from article content and user response. The article representation is modeled via an LSTM model with the article embedding via Doc2Vec [12] and user response.
- **SAME_w [3]**: SAME uses news image, content, metadata and users’ sentiments for fake news detection. As the majority of news does not have a cover image, the visual part of the model is ignored, and the baseline is termed as SAME_v.
- **HAN [27]**: HAN has two levels of attention mechanisms applied at the word and sentence-level, enabling it to attend differently to more and less important content for document classification.
- **dEFEND [23]**: dEFEND utilizes the hierarchical attention network as HAN on article content, and a co-attention mechanism between article content and user comment for misinformation detection.

To evaluate the performance of misinformation detection algorithms, we use the following metrics, which are commonly used to evaluate classifiers in related areas: PR-AUC, Precision, Recall, and F1 score.

We implement all models with Keras. We randomly use the labels of 75% news articles for training and predict the remaining 25%. We set the hidden dimension of all models to 100. The word embeddings are initialized by GloVe [19] and the dimension of pre-trained word embeddings is set to 100. For deep learning methods, we use Adam with a minibatch of 50 articles on the dataset, and the training epoch is set as 10. For a fair comparison, we use the cross-entropy loss for all methods. Finally, we run each method five times and report the average score in Table 4.

From Table 4, we can see that state-of-the-art methods perform better than simple methods, as they incorporate signals from user engagement, better capturing contextual information. However, as the dataset is quite imbalanced, the models tend to generate many fake positive cases. Thus, the recall and F1 values are far from being satisfactory. As the proportion of true and fake information is likely to be even more skewer, practical detection solutions need to handle this type of imbalance setting more.

In addition, with regard to healthcare news, it is often more scarce and rare for lay persons to give discriminating comments due to lack of professional knowledge. For example, WHO tweeted “There is no evidence that regularly rinsing the nose with saline solution has protected people from infection with the new coronavirus.” The majority of comments for the tweet are unrelated, useless, and even includes hate speech and other misinformation about

Table 4: Misinformation detection performance

	PR AUC	Precision	Recall	F1
SVM	0.3365	0.4036	0.1322	0.1986
LR	0.2871	0.4287	0.0690	0.1143
RF	0.3937	0.6056	0.0581	0.1045
CNN	0.8126	0.9653	0.1238	0.1983
BiGRU	0.2241	0.7476	0.0524	0.0930
CSI	0.3576	0.6814	0.2109	0.2283
SAME\w	0.7901	0.8922	0.2991	0.3400
HAN	0.6824	0.6965	0.4659	0.5471
dEFEND	0.7229	0.8965	0.4847	0.5814

COVID-19. Therefore, the state-of-the-art methods [20, 3, 23] that can exploit the signals from user engagement have not achieved satisfactory results.

6 Conclusion

In this paper, we present a comprehensive COVID-19 misinformation dataset CoAID, which contains news articles, user engagement, and social platform posts. We describe how we collected the dataset. In addition, we conduct brief data analysis to show the distinctive features between misinformation and facts, and demonstrate the future research directions through a fake news detection task over several state-of-the-art methods. We hope researchers to find CoAID useful for their research and together contribute to flatten the curve of COVID-19.

Acknowledgments

We thank all healthcare providers who risk their lives for others, people staying home for others, and the kindness that connects us all.

This work was in part supported by NSF awards #1742702, #1820609, #1915801, and #1934782.

References

- [1] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*, 2015.
- [2] John M Carey, Victoria Chi, DJ Flynn, Brendan Nyhan, and Thomas Zeitzoff. The effects of corrective information about disease epidemics and outbreaks: Evidence from zika and yellow fever in brazil. *Science advances*, 6(5):eaaw7449, 2020.
- [3] Limeng Cui, Suhan Wang, and Dongwon Lee. Same: sentiment-aware multi-modal embedding for detecting fake news. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 41–48, 2019.
- [4] Enyan Dai, Yiwei Sun, and Suhan Wang. Ginger cannot cure cancer: Battling fake health news with a comprehensive data repository. *arXiv preprint arXiv:2002.00837*, 2020.
- [5] Sameer Dhoju, Md Main Uddin Rony, Muhammad Ashad Kabir, and Naeemul Hassan. Differences in health news from reliable and unreliable media. In *Companion Proceedings of The 2019 World Wide Web Conference*, pages 981–987, 2019.
- [6] Nicholas DiFonzo and Prashant Bordia. *Rumor psychology: Social and organizational approaches*. American Psychological Association, 2007.
- [7] Centers for Disease Control and Prevention. Symptoms of coronavirus. <https://www.cdc.gov/coronavirus/2019-ncov/symptoms-testing/symptoms.html>. Accessed: 2020-05-08.
- [8] Axel Gelfert. Fake news: A definition. *Informal Logic*, 38(1):84–117, 2018.
- [9] Amira Ghenai and Yelena Mejova. Fake cures: user-centric modeling of health misinformation in social media. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):1–20, 2018.

- [10] Clayton J Hutto and Eric Gilbert. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth international AAAI conference on weblogs and social media*, 2014.
- [11] Stephen A Lauer, Kyra H Grantz, Qifang Bi, Forrest K Jones, Qulu Zheng, Hannah R Meredith, Andrew S Azman, Nicholas G Reich, and Justin Lessler. The incubation period of coronavirus disease 2019 (covid-19) from publicly reported confirmed cases: estimation and application. *Annals of internal medicine*, 2020.
- [12] Quoc Le and Tomas Mikolov. Distributed representations of sentences and documents. In *International conference on machine learning*, pages 1188–1196, 2014.
- [13] Maria Molina, Shyam Soundar, Thai Le, and Dongwon Lee. “fake news” is not simply false information: A concept explication and taxonomy of online content. *American Behavioral Scientist*, pages 1–33, 2019.
- [14] World Health Organization. Q&a on coronaviruses (covid-19). <https://www.who.int/emergencies/diseases/novel-coronavirus-2019/question-and-answers-hub/q-a-detail/q-a-coronaviruses>. Accessed: 2020-05-08.
- [15] World Health Organization. Q&a: Similarities and differences – covid-19 and influenza. <https://www.who.int/emergencies/diseases/novel-coronavirus-2019/question-and-answers-hub/q-a-detail/q-a-similarities-and-differences-covid-19-and-influenza>. Accessed: 2020-05-08.
- [16] World Health Organization. Statement on the second meeting of the international health regulations (2005) emergency committee regarding the outbreak of novel coronavirus (2019-ncov). [https://www.who.int/news-room/detail/30-01-2020-statement-on-the-second-meeting-of-the-international-health-regulations-\(2005\)-emergency-committee-regarding-the-outbreak-of-novel-coronavirus-\(2019-ncov\)](https://www.who.int/news-room/detail/30-01-2020-statement-on-the-second-meeting-of-the-international-health-regulations-(2005)-emergency-committee-regarding-the-outbreak-of-novel-coronavirus-(2019-ncov)). Accessed: 2020-05-08.
- [17] World Health Organization. Who director-general’s opening remarks at the media briefing on covid-19 - 6 may 2020. <https://www.who.int/dg/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19—6-may-2020>. Accessed: 2020-05-08.
- [18] World Health Organization. Who director-general’s opening remarks at the media briefing on covid-19-11 march 2020. <https://www.who.int/dg/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19—11-march-2020>. Accessed: 2020-05-08.
- [19] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.
- [20] Natali Ruchansky, Sungyong Seo, and Yan Liu. Csi: A hybrid deep model for fake news detection. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 797–806, 2017.
- [21] Fatima K Abu Salem, Roaa Al Feel, Shady Elbassuoni, Mohamad Jaber, and May Farah. Fa-kes: a fake news dataset around the syrian war. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 13, pages 573–582, 2019.
- [22] Chengcheng Shao, Giovanni Luca Ciampaglia, Alessandro Flammini, and Filippo Menczer. Hoaxy: A platform for tracking online misinformation. In *Proceedings of the 25th international conference companion on world wide web*, pages 745–750, 2016.
- [23] Kai Shu, Limeng Cui, Suhang Wang, Dongwon Lee, and Huan Liu. defend: Explainable fake news detection. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 395–405, 2019.
- [24] Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 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.
- [25] William Yang Wang. “liar, liar pants on fire”: A new benchmark dataset for fake news detection. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 422–426, 2017.
- [26] Yaqing Wang, Fenglong Ma, Zhiwei Jin, Ye Yuan, Guangxu Xun, Kishlay Jha, Lu Su, and Jing Gao. Eann: Event adversarial neural networks for multi-modal fake news detection. In *Proceedings of the 24th acm sigkdd international conference on knowledge discovery & data mining*, pages 849–857, 2018.
- [27] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. Hierarchical attention networks for document classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*, pages 1480–1489, 2016.

A Releases

A.1 Version 0.1

This initial dataset corresponding to the paper includes COVID-19 information collected from December 1, 2019 through May 1, 2020. It includes 1,896 news, 183,564 related user engagements, 516 social platform posts about COVID-19, and ground truth labels.

Table 5: Statistics of CoAID Version 0.1

	Fake		True	
	Claim	News Article	Claim	News Article
# Information on Website	27	135	166	1,568
# Tweets	457	9,218	6,342	87,324
# Replies	623	5,721	9,764	64,115
# Social Platform Posts	414 (492)	-	21 (24)	-
# Total	1,521 (1,599)	15,074	16,293 (16,296)	153,007

A.2 Version 0.2

Besides the original data in Version 0.1, the additional data includes COVID-19 information collected from May 1, 2020 through July 1, 2020. It includes 1,339 news, 111,128 related user engagements, 335 social platform posts about COVID-19, and ground truth labels.

The statistics of the additional data in CoAID Version 0.2 is shown in Table 6.

Table 6: Statistics of the additional data in CoAID Version 0.2

	Fake		True	
	Claim	News Article	Claim	News Article
# Information on Website	1	55	172	1,111
# Tweets	27	1,198	1,453	54,224
# Replies	3	1,672	1,909	50,642
# Social Platform Posts	197 (320)	-	14 (15)	-
# Total	228 (351)	2,925	3,548 (3,549)	105,977

A.3 Version 0.3

Besides the original data in Version 0.2, the additional data includes COVID-19 information collected from July 1, 2020 through September 1, 2020. It includes 1,016 news, 1,308 related user engagements, 75 social platform posts about COVID-19, and ground truth labels.

The statistics of the additional data in CoAID Version 0.3 is shown in Table 7.

Table 7: Statistics of the additional data in CoAID Version 0.3

	Fake		True	
	Claim	News Article	Claim	News Article
# Information on Website	0	14	116	886
# Tweets	0	23	297	104
# Replies	0	43	778	63
# Social Platform Posts	39 (68)	-	7	-
# Total	39 (68)	80	1,198	1,053