

Detection of False Information Regarding Covid-19. A study

Grace Lawson¹

¹ University of Auckland - Nueva Zelanda, lawsongrace168@gmail.com,
<https://orcid.org/0009-0007-0232-1093>

ABSTRACT

The surge in the dissemination of false information on social media platforms, particularly regarding Covid-19, poses a significant threat to both the mental and physical well-being of individuals. Detecting and preventing the spread of such misinformation is a crucial undertaking. This article provides an overview of various approaches employed for the detection of fake news related to Covid-19, encompassing Classical Machine Learning models, Neural Network-based models, and those derived from alternative methodologies and preprocessing steps. The analysis includes insights from the "Constraint@AAAI2021 - COVID19 Fake News Detection" challenge, which aimed to binary classify news sourced from social media into fake and real categories. We examine the most effective approaches proposed by researchers during the challenge. Additionally, we detail datasets containing Covid-19-related fake news, offering valuable resources for the detection and classification of such misinformation.


Keywords: *Fake News Detection; Information Security; Social Media*



Copyright: © 2023 by the authors. This article is an open access article distributed under the terms and conditions of the Creative Commons

Received:
21 September, 2023

**Accepted for
publication:**
02 November, 2023



Detección de Información Falsa sobre Covid-19. Un estudio

RESUMEN

El aumento en la difusión de información falsa en plataformas de redes sociales, especialmente en relación con Covid-19, representa una amenaza significativa tanto para el bienestar mental como físico de las personas. Detectar y prevenir la propagación de desinformación es una tarea crucial. Este artículo proporciona una visión general de diversos enfoques empleados para la detección de noticias falsas relacionadas con Covid-19, abarcando modelos de Aprendizaje Automático Clásico, modelos basados en Redes Neuronales y aquellos derivados de metodologías alternativas y pasos de preprocesamiento. El análisis incluye aportes del desafío "Constraint@AAAI2021 - Detección de Noticias Falsas de COVID-19", que tenía como objetivo clasificar binariamente las noticias provenientes de las redes sociales en las categorías de falsas y reales. Examinamos los enfoques más efectivos propuestos por los investigadores durante el desafío. Además, detallamos conjuntos de datos que contienen noticias falsas relacionadas con Covid-19, ofreciendo recursos valiosos para la detección y clasificación de dicha desinformación.

Palabras Clave: *Detección de Noticias Falsas; Seguridad de la Información; Redes Sociales*

INTRODUCTION

As the role of social media continues to grow in our daily lives, the proliferation of fake news on these platforms has emerged as a significant problem. Fake news, deliberate misinformation crafted to mislead readers for financial or political gain, poses a threat that extends beyond mere deception, causing both mental and physical harm [1]. Controlling and preventing the spread of fake news is crucial to mitigating these negative effects.

Several notable works [2,3] aim to outline the primary approaches employed by researchers in the detection of fake news. In this context, "fake news" refers to misleading information across various domains such as politics, economics, goods and services, travel, and tourism. The spectrum of detection approaches is broad, encompassing classical machine learning models like Text Transformers, as well as neural network-based models. The swift detection and classification of fake news are of paramount importance, given the need for prompt identification.

Particularly critical is the identification and prevention of the dissemination of fake news related to Covid-19, an urgent and global issue that has intensified over the past two years with the onset of the pandemic. In [4], it was demonstrated that the proliferation of hoaxes and misinformation contributes to the increased spread of the virus and a decline in mental health among individuals. To avert such dire consequences, the authors propose fostering closer communication among mass media, healthcare organizations, and other key stakeholders, creating a unified platform for disseminating accurate health-related information. Additionally, leveraging AI, including natural language processing tools, becomes crucial in identifying and removing fake news from all social media platforms. Implementing specialized law enforcement measures is also deemed essential to control the spread of such misinformation.

This article aims to elucidate potential methodologies, recent research findings, and valuable datasets pertinent to the detection of fake news related to Covid-19. It is crafted to assist new researchers in gaining a comprehensive understanding of this subject.

Article Structure

The article is organized as follows: The Introduction section furnishes a definition of fake news, outlines the hazards associated with such misinformation, and specifically addresses fake news concerning Covid-19. Section 2 delves into different approaches for detecting fake news related to Covid-19, encompassing classical machine learning models, neural network-based models, and models derived from alternative methodologies. Section 3 focuses on the "Constraint@AAAI2021 - COVID19 Fake News Detection" challenge, exploring the task, dataset, and models that yielded optimal results. Section 4 provides an overview of datasets dedicated to detecting fake news associated with Covid-19. The concluding Section 5 summarizes the key findings of the survey.

Approaches to Detecting Covid-19 Fake News

When discussing the most effective models for detecting Covid-19 fake news, attention is directed towards the structure of the review. It is essential to categorize these models into groups based on their nature, facilitating a more coherent and comprehensible review. While there is no strict delineation among model types, researchers typically work with a diverse array of models and model combinations, including ensembles, each involving different preprocessing steps to identify optimal combinations. Despite this lack of strict categorization, three primary groups of models can be conditionally distinguished: classical machine learning models, models based on neural networks, and models derived from alternative approaches. This section meticulously reviews each of these three model groups.

Classical Machine Learning Approach for Detecting Covid-19 Fake News

The classical machine learning approach is widely utilized for detecting fake news, especially those related to Covid-19. In scenarios requiring binary classification to distinguish between fake and real news, classical machine learning algorithms prove highly effective. Logistic Regression stands out among these algorithms, employing the fundamental concept of a linear classifier that divides feature space into two halves using a hyperplane, with each half-space corresponding to a class in binary classification.

Support Vector Machines are popular classifiers in the realm of classification problems, creating hyperplanes in multidimensional or infinite-dimensional spaces. Additionally, methods based on

classification trees, such as the Gradient Boosting Classifier and Random Forest Classifier, are employed for fake news detection. Notably, the Random Forest Classifier builds each tree independently and consolidates results at the conclusion of the process.

In the context of using classical machine learning algorithms for identifying fake news about Covid-19, researchers [5] conducted an analysis focused on fake news related to the COVID pandemic. They gathered a dataset from 150 users, extracting data from various social media accounts like Twitter, email, mobile, Whatsapp, and Facebook over a 4-month period from March 2020 to June 2020. The preprocessing stage involved removing non-Covid-19-related information and eliminating incomplete news. K-Nearest Neighbour was utilized for classification, yielding the best prediction results for June with a 0.91 F1-score and the least favorable outcomes for March with a 0.79 F1-score.

In [6,7], researchers sought to detect fake news related to Covid-19 using small datasets. Logistic Regression, Support Vector Machine, Gradient Boosting, and Random Forest were compared on a limited dataset of 1,000 fake and real messages. Support Vector Machine and Random Forest classifiers exhibited the most favorable results, achieving a 69% micro-F1 score. The authors highlighted that although these results may not match those obtained on larger datasets of fake news about Covid-19, this approach could prove valuable for researchers with limited datasets or time constraints where swift classification decisions are imperative.

In [8], the authors conducted a comparative analysis of four machine learning baselines (Decision Tree, Logistic Regression, Gradient Boost, and Support Vector Machine (SVM)) using a dataset focused on detecting fake news related to Covid-19. The dataset was personally collected by the authors from various social media platforms, including Facebook and Twitter. The SVM model demonstrated the best performance, achieving an F1-score of 0.93. This suggests that numerous machine learning algorithms exhibit promising results in binary classification tasks for distinguishing between fake and real news.

Neural Networks for Detecting Covid-19 Fake News

A diverse array of algorithms based on neural networks has proven effective in detecting fake news, particularly related to Covid-19. Researchers have employed various linguistic models, with BERT

[9]—Bidirectional Encoder Representations from Transformers—being a crucial one. Fine-tuning BERT with an additional output layer enables the creation of state-of-the-art models for diverse tasks. The DistilBERT [10] model reduces the size of a BERT model by 40%, retaining 97% of its language understanding capabilities while operating 60% faster. The COVID-Twitter-BERT (CT-BERT) model [11], pretrained on a substantial corpus of COVID-19-related Twitter messages, has demonstrated efficacy in detecting such messages.

RoBERTa [12], an essential model often employed, is a robust BERT variant trained on a larger dataset, for more iterations, and with a larger batch size. Other noteworthy models include ELECTRA [13], which employs a unique approach of pre-training text encoders as discriminators, and ALBERT [14], a Lite BERT designed for self-supervised learning of language representations.

The XLNet model [15], akin to BERT, learns bidirectional context with an autoregressive formulation. Hierarchical Attention Networks (HAN), based on LSTM, consist of four sequential levels—word encoder, word-level attention, sentence encoder, and sentence-level attention [16].

ELMO [17], a deep contextualized word representation, captures complex characteristics of word usage, including syntax and semantics, while considering variations across linguistic contexts (polysemy modeling).

The paper [18] showcases a successful application of neural networks in detecting Covid-19 fake news. The authors gathered 4.8K expert-annotated social media posts related to COVID-19 and 86 common misconceptions about the disease. This dataset was utilized to assess the performance of misinformation detection during the Coronavirus pandemic. Each message was categorized within the context of selected misconceptions as "misinformative," "informative," or "irrelevant."

For modeling, the authors employed TF-IDF and GloVe parameterizations to generate vectorized representations. The RoBERTabase implementation, incorporating two textual similarity models (cosine similarity computed by averaging token vectors and BERTScore involving cosine similarities between RoBERTa token embeddings), was used to derive contextual word embeddings. Additionally, the RoBERTa-base model was fine-tuned using the COVID-19 tweet dataset. The domain-adapted BERTScore emerged as the top-performing similarity model.

In [19], the authors devised an ensemble of linguistic models, including XLNet, RoBERTa, XLMRoBERTa, DeBERTa, ERNIE 2.0, and ELECTRA, for Covid-19 fake news detection. They employed a tweet-preprocessor library in Python [20] to filter out noise such as usernames, URLs, and emojis. The authors implemented Heuristic Post-Processing, considering Soft-voting prediction vectors. Through preprocessing, ensemble construction, and adopting soft-voting, the researchers achieved an impressive F1-score of 0.9831.

Addressing the challenge of Covid-19 fake news detection, [21] implemented a combination of topical distributions from Latent Dirichlet Allocation (LDA) with contextualized representations from XLNet. The implementation utilized the Transformers library [22] by Hugging Face, providing the PyTorch interface for XLNet. This model yielded a notable 0.967 F1-score on the test dataset. To benchmark the results, the authors compared various models, including SVM with contextualized representations using Universal Sentence Encoder (USE), BERT with document-topic distributions from LDA, fine-tuned pretrained XLNet model, and a combination of BERT and BERT+topic models. Among these, the XLNet model with topic distributions demonstrated the most favorable outcomes.

Additional Approaches for Detecting Covid-19 Fake News

In addressing the challenge of identifying fake news related to Covid-19, classical machine learning models and neural networks are pivotal. However, numerous other features and methodologies, widely embraced in natural language processing, offer valuable contributions to fake news detection.

Various algorithms leverage n-grams of words or characters [23], while GloVe [24], an unsupervised learning algorithm, yields vector representations for words. Fasttext [25], a classifier known for matching deep learning classifiers in accuracy but excelling in terms of training and evaluation speed, is also noteworthy. Label Smoothing [26], a regularization technique introducing noise to labels, and adversarial training [27], a novel regularization method enhancing model robustness against perturbations, are crucial considerations. Additionally, tax2vec [28], a semantic space vectorization algorithm, adds to the repertoire of available approaches.

Research investigating the impact of Covid-19 fake news, such as [29], delves into the social consequences of misinformation within the context of health information on social media. Scholars aim to discern between false information and evidence of social impact shared on social platforms.

The study's focal point is data gathered from three major social media platforms: Twitter, Facebook, and Reddit. The research questions revolve around understanding how social media messages addressing fake health information interact and how discussions based on health evidence with social impact counteract misinformation.

To analyze this, the authors employed a methodology combining quantitative and qualitative content analysis of collected messages. The selection criteria for messages were based on relevance to the number of active users, availability of public messages, and suitability for online discussion. Hashtags like “health,” “vaccines,” “nutrition,” and “Ebola” were chosen, aligning with the researchers' interests. The categorized data fell into four groups: ESISM (messages exemplifying evidence of social impact shared on social media), MISFA (messages containing fake news), OPINION (messages reflecting a user's opinion), and INFO (messages presenting facts or news).

Consequently, the study unveiled that misinformation and fake news are more prevalent on Twitter (19%) compared to Facebook (4%) or Reddit (7%). The analysis also highlighted that messages centered around false health information tended to exhibit aggression in the examined data, whereas those conveying evidence of social impact were characterized by a more peaceful and respectful tone. Furthermore, messages founded on evidence of social impact effectively countered false information, especially when the authors of the misinformation messages demonstrated politeness and openness to discussion. Conversely, challenges arose when misinformation authors exhibited disrespect and adopted an aggressive stance against scientific principles.

In their work [30], the authors leveraged natural language processing to scrutinize messages related to COVID-19 and identify predominant topics discussed on social media during the pandemic. The dataset was constructed by extracting messages and comments from platforms such as Twitter, Youtube, Facebook, and three online discussion forums.

During the preprocessing phase, the authors eliminated URLs, HTML tags, number words, unnecessary special characters for sentence boundary detection, hashtags, and mentions. They also expanded contractions, condensed words with repeated characters, and converted slangs to standard English words.

Subsequently, the Keyphrase Extractor algorithm was implemented using Python, encompassing seven steps:

- (1) grammar determination,
- (2) sentence breaking and tokenization,
- (3) POS tagging,
- (4) lemmatization,
- (5) chunking,
- (6) transformation and filtering, and
- (7) sentiment scoring and filtering.

Through this process, the authors, utilizing keyphrases and expert assessments, identified 34 negative and 20 positive categories. The top-5 negative themes included concerns about social distancing and isolation policies, misinformation, political influence, financial issues, and poor governance. On the positive side, the top-5 themes associated with the pandemic comprised public awareness, spiritual support, encouragement, charity, and entertainment.

Constraint@AAAI2021 - COVID-19 Fake News Detection

In the contemporary landscape, while some challenges incorporate Covid-19-related fake news in their datasets, such as the IberLEF-2021 shared task [31], encompassing 237 instances of Covid-19 fake news, there is a noticeable scarcity of challenges exclusively dedicated to the detection of fake news about Covid-19. This section delves into the significance of the "Constraint@AAAI2021 Fake News Detection" challenge [32], particularly in the context of Covid-19-related fake news detection. The challenge presented participants with a dataset comprising such fake news, and the diverse models developed by participants aimed to detect fake news about Covid-19. Here, we provide an overview of the task, the dataset, and the noteworthy models created within the challenge.

Challenge and Dataset

The primary objective of the mentioned challenge was to devise a system for binary classification, distinguishing between fake news and real news. The challenge encompassed the same task for two languages: English and Hindi. However, our focus in this survey is solely on the English task and dataset.

The dataset designed for detecting Covid-19 fake news comprises 10,700 messages, with 5,100 categorized as real and 5,600 as fake. Real news sources were drawn from reputable outlets like the World Health Organization (WHO), Centers for Disease Control and Prevention (CDC), and others. In contrast, fake news was sourced from social media platforms such as Facebook, Twitter, Instagram, and more. The dataset encompasses 37,503 unique words, and key statistics are presented in Table 1.

Table 1. Quantitative Characteristics of the Dataset

Attribute	Fake	Real	Combined
Unique words	19728	22916	37503
Avg. words per post	22	32	27
Avg. chars per post	143	218	183

A notable observation from the dataset reveals that, on average, real news tends to be lengthier than fake news by approximately 10 words, and the character count in a fake message is nearly 75 characters lower than in a real one. Table 2 provides examples of both real and fake news from the dataset, sourced from Twitter.

Table 2. Exemplars of Genuine and Counterfeit News

Label	Text
Fake	1972 #Watch Italian Billionaire commits suicide by throwing himself from 20th Floor of his tower after his entire family was wiped out by #Coronavirus #Suicide has never been the way, may soul rest in peace May God deliver us all from this time
Fake	2 Trump announced that Roche Medical Company will launch the vaccine next Sunday and millions of doses are ready from it !!! The end of the play
Fake	China Muslims hidden at Bihari mosque has been taken to corona virus test by Bihari police. Erode police has caught Thailand Muslim mullahs infected with corona virus. Today Salem Police has caught 11 Indonesian Muslim mullahs at Salem mosque. This video shows that they are applying and putting saliva on spoons plates and utensils and also they are in the intention of spreading corona virus disease. Nobody knows what's happening in the Nation
Real	Almost 200 vaccines for #COVID19 are currently in clinical and pre-clinical testing. The history of vaccine development tells us that some will fail and some will succeed-@DrTedros #UNGA #UN75
Real	14 new cases of #COVID19 have been confirmed in Nigeria: 2 in FCT 12 in Lagos Of the 14 6 were detected on a vessel 3 are returning travellers into Nigeria: 1 is close contact of a confirmed case As at 7:35 pm 26th March there are 65 confirmed cases 3 discharged 1 death
Real	Currently most cases of #COVID19 in the US are in California and Washington State. However many other communities are also dealing with cases of COVID-19. See CDC recommendations for preventing spread of COVID-19 in communities.

Submissions from participants in the shared task were evaluated based on their weighted average F1-score. The F1-score was computed separately for the fake news and real news classes, and the average was weighted by the number of true instances for each class.

Optimal Approaches Identified in the Challenge

A total of 166 teams participated in the challenge with the English dataset, and notably, 114 of them surpassed the baseline F1-score of 93%. The top-performing results were closely clustered, surpassing an impressive 98% F1-score. The three leading outcomes are detailed in Table 3.

Table 3. Leading Three Outcomes in the Constraint@AAAI2021 - COVID-19 Fake News Detection Shared Task (percentage)

Rank	Team	Precision/Recall/ F1-score
1	g2tmn	98.69 / 98.69 / 98.69
2	saradhix	98.65 / 98.64 / 98.65
3	xiangyangli	98.60 / 98.60 / 98.60

The winning team of the challenge, g2tmn [33], attained a remarkable 98.69% F1-score by employing an ensemble of three pretrained CT-BERT models with varying random seed values and distinct data splitting methods for training and validation samples. In the preprocessing phase, researchers utilized the Python emoji library for emoji replacement with concise textual descriptions [34], tokenization of URLs (substituted with the \$URL\$ token), and conversion of message texts to lowercase.

Securing the second position with a 98.65% F1-score, the saradhix team [35] employed a diverse set of classical machine learning methods, including Naive Bayes, Logistic Regression, Random Forest, XGBoost, and Support Vector Machine, alongside Transformer models such as BERT, DB-BERT, RoBERTa, Electra, and XLNet. The optimal results were achieved with a RoBERTa-based model featuring 12 layers, 768 hidden units, 12 heads, and 125 million parameters. The researchers highlighted RoBERTa as a robust BERT model, trained on a significantly larger dataset and for an extended period with a larger batch size.

Claiming the third spot with a 98.6% F1-score, the xiangyangli team [36] utilized Text Transformers in their research. Additionally, the authors implemented a Pseudo Label Algorithm for data augmentation, considering the relatively small size of the dataset. Any test data predicted with a probability exceeding 0.95 was deemed correct with high confidence and was incorporated into the training set. The ensemble, comprising BERT, Ernie, XL-Net, RoBERTa, and Electra models, was created using Text Transformers, and cross-validation with the Pseudo Label Algorithm contributed to their outstanding performance.

To synthesize the key findings from the Constraint@2021 Fake News Detection open shared task:

1. The most effective models emerged as ensembles of Text Transformers.
2. Fine-tuning these Transformers proved to be the pivotal step, with preprocessing steps, typically crucial for classical machine learning models, playing a relatively insignificant role in these cases.

Datasets for the Detection of Covid-19 Fake News

Previously, we detailed the dataset from the Constraint@AAAI2021 - COVID-19 Fake News Detection challenge, suitable for binary classification of fake news and real news. In this section, we introduce additional datasets for the detection of fake news related to Covid-19.

The CoAID dataset (COVID-19 Healthcare Misinformation Dataset) [37] encompasses diverse COVID-19 healthcare misinformation, including fake news on websites and social platforms, along with user engagement related to such news. CoAID consists of 4,251 news articles, 296,000 user engagements, 926 social platform posts about COVID-19, and associated ground truth labels.

CONCLUSION

The identification of fake news, especially misinformation related to Covid-19, has become a matter of unprecedented importance. This article delved into various methodologies for detecting such news, covering classical machine learning models, neural network-based models, and alternative approaches. Our exploration involved a comprehensive analysis of the Constraint@AAAI2021 - COVID19 Fake News Detection challenge and its corresponding dataset. The results from the challenge emphasized the effectiveness of ensemble models, specifically Transformers like BERT and CT-BERT, for the binary classification of fake news and real news.

Additionally, we spotlighted datasets exclusively crafted for the detection of Covid-19-related fake news, providing valuable assets for addressing the complexities inherent in this detection and classification challenge.

The surge in Covid-19 fake news is seen as a manifestation of information warfare, prompting consideration of information war models by experts in mathematical modeling [42, 43]. Future endeavors may involve utilizing such models to forecast the dynamics of Covid-19 fake news and real news.

It remains crucial to acknowledge established models and methods for fake news detection, particularly those tailored to account for the stylistic nuances of texts [44]. These approaches continue to be pertinent and valuable in navigating the intricacies of Covid-19 fake news, assuming that such misinformation is an integral aspect of the broader landscape of misleading information.

The first multilingual dataset for Covid-19 fake news, FakeCovid, spans news from 150 countries in 40 languages [38]. Comprising 5,182 fact-checked news articles for COVID-19 collected between 04/01/2020 and 15/05/2020, the dataset draws from 92 different fact-checking websites referenced by Poynter and Snopes. The authors categorized the fact-checked news into 11 different content-related categories.

An alternative approach to obtaining Covid-19 fake news involves using Elasticsearch3 to retrieve validated fake news from FakeHealth4 [39]. Additionally, the COVID-19 Infodemic Twitter dataset [40] provides annotated tweets with fine-grained labels related to disinformation about COVID-19, along with an annotation schema and detailed instructions for dataset creation.


Lastly, in [41], the author describes a dataset where tweets with the hashtag #covid19 are collected using the Twitter API and a Python script. This ongoing collection, initiated on 25/7/2020, allows for trends analysis and the creation of other Covid-19 fake news datasets.

BIBLIOGRAPHICAL REFERENCES

1. Hunt, E. (2016). What is fake news? How to spot it and what you can do to stop it. The Guardian. Retrieved from <https://www.theguardian.com/media/2016/dec/18/what-is-fake-news-pizzagate>
2. Choraś, M., Demestichas, K., Giełczyk, A., Herrero, Á., Ksieniewicz, P., Remoundou, K., Urda, D., Wozniak, M.(2020). Advanced Machine Learning techniques for fake news (online disinformation) detection: A systematic mapping study. Applied Soft Computing. 101. 107050. DOI: 10.1016/j.asoc.2020.107050.
3. Zhou, X., Zafarani, R. (2020). A survey of fake news: Fundamental theories, detection methods, and opportunities. ACM Computing Surveys. 53. DOI: 10.1145/3395046.
4. Tasnim, S., Hossain, M., Mazumder, H. (2020). Impact of rumors and misinformation on COVID-19 in social media. J Prev Med Public Health. 53(3):171–174. DOI: 10.3961/jpmph.20.094.

5. Bandyopadhyay, S., Dutta, S. (2020). Analysis of fake news in social medias for four months during lockdown in COVID-19. DOI: 10.20944/preprints 202006.0243.v1.
6. Shushkevich, E. & Cardiff, J. (2021). Detecting fake news about Covid-19 on small datasets with machine learning algorithms. Proceedings of the 30th Conference of Open Innovations Association FRUCT, pp. 253–258.
7. Shushkevich, E., M. Alexandrov, M., Cardiff, J. (2021). Detecting fake news about Covid-19 using classifiers from Scikit-learn. International Workshop on Inductive Modeling IWIM'2021, 5 pp.
8. Patwa, P., Sharma, S., Pykl, S., Guptha, V., Kumari, G., Akhtar, M., Ekbal, A., Das, A., Chakraborty, T. (2021). Fighting an infodemic: COVID-19 fake news dataset. arXiv:2011.03327.
9. Devlin J., Chang, M., Lee, K., Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
10. Sanh, V., Debut, L., Chaumond, J., Wolf, T. (2019). Distilbert, a distilled version of Bert: Smaller, faster, cheaper and lighter. CoRR 1910.01108.
11. Muller, M., Salathe, M., Kummervold, P. E.: (2020). COVID-Twitter-BERT: A natural language processing model to analyse COVID-19 content on Twitter. arXiv preprint arXiv:2005.07503.
12. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V. (2019). Roberta: A robustly optimized Bert pretraining approach. arXiv preprint arXiv:1907.11692.
13. Clark, K., Luong, M.T., Le, Q., Manning, C. (2020). Electra: Pre-training text encoders as discriminators rather than generators. arXiv preprint arXiv: 2003.10555.
14. Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., Soricut, R., Carvalho, M. (2019). Albert: A lite Bert for self-supervised learning of language representations. arXiv preprint arXiv:1909.11942.
15. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R. & Le, Q. (2019). Xlnet: Generalized autoregressive pretraining for language understanding. In Advances in neural information processing systems, pp. 5753–5763.

16. Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., Hovy, E. (2016). Hierarchical attention networks for document classification. Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies. pp. 1480–1489.
17. Peters, M.E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., Zettlemoyer, L. (2018). Deep contextualized word representations. arXiv preprint arXiv:1802.05365.
18. Hossain, T., Logan, R., Ugarte, A., Matsubara, Y., Young, S., & Singh, S. (2020). Detecting COVID19 misinformation on social media. Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020. DOI:10.18653/v1/2020.nlpccovid19- 2.1.
19. Dipta, S., Basak, A., Dutta, S. (2021). A heuristicdriven ensemble framework for COVID-19 fake news detection. In Combating Online Hostile Posts Computación y Sistemas, Vol. 25, No. 4, 2021, pp. 783–792 doi: 10.13053/CyS-25-4-4089
20. Hancock, J., Markowitz, D. (2014). Linguistic traces of a scientific fraud: The case of Diederik Stapel. PLoS One 9, no. 8.
21. Gautam, A., Venkatesh, V., Masud, S. (2021). Fake news detection system using XLNet model with topic distributions: CONSTRAINT@AAAI2021 Shared Task, 2101.11425, arXiv, cs.CL.
22. Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T., Gugger, S., Rush, A. (2020). Huggingface's transformers: State-of-the-art natural language processing. ArXiv.
23. Martinc, M., Skrlj, B., Pollak, S. (2018). Multilingual gender classification with multiview deep learning: Notebook for PAN at CLEF 2018. In: Cappellato, L., Ferro, N., Nie, J., Soulier, L. editors Working Notes of CLEF 2018 - Conference and Labs of the Evaluation Forum, Avignon, France, September 10-14, 2018. CEUR Workshop Proceedings, vol. 2125.
24. Pennington, J., Socher, R., Manning, C.D. (2014). Glove: Global vectors for word representation. Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP). pp. 1532–1543.



25. Joulin, A., Grave, E., Bojanowski, P., Mikolov, T. (2016). Bag of tricks for efficient text classification. arXiv preprint arXiv:1607.01759.
26. Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A. (2016). Inception-v4, inception-Resnet and the impact of residual connections on learning. arXiv preprint arXiv:1602.07261