



HAL
open science

Analyzing COVID-Related Social Discourse on Twitter using Emotion, Sentiment, Political Bias, Stance, Veracity and Conspiracy Theories

Youri Peskine, Raphaël Troncy, Paolo Papotti

► To cite this version:

Youri Peskine, Raphaël Troncy, Paolo Papotti. Analyzing COVID-Related Social Discourse on Twitter using Emotion, Sentiment, Political Bias, Stance, Veracity and Conspiracy Theories. WWW 2023, Companion Proceedings of the ACM Web Conference 2023, 30 April-4 May 2023, Austin, USA, ACM, Apr 2023, Austin TX USA, United States. pp.688-693, 10.1145/3543873.3587622 . hal-04087004

HAL Id: hal-04087004

<https://hal.science/hal-04087004v1>

Submitted on 2 May 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Analyzing COVID-Related Social Discourse on Twitter using Emotion, Sentiment, Political Bias, Stance, Veracity and Conspiracy Theories

Youri Peskine*
your.i.peskine@eurecom.fr
EURECOM
Sophia Antipolis, France

Raphaël Troncy
raphael.troncy@eurecom.fr
EURECOM
Sophia Antipolis, France

Paolo Papotti
paolo.papotti@eurecom.fr
EURECOM
Sophia Antipolis, France

ABSTRACT

Online misinformation has become a major concern in recent years, and it has been further emphasized during the COVID-19 pandemic. Social media platforms, such as Twitter, can be serious vectors of misinformation online. In order to better understand the spread of these fake-news, lies, deceptions, and rumours, we analyze the correlations between the following textual features in tweets: emotion, sentiment, political bias, stance, veracity and conspiracy theories. We train several transformer-based classifiers from multiple datasets to detect these textual features and identify potential correlations using conditional distributions of the labels. Our results show that the online discourse regarding some topics, such as COVID-19 regulations or conspiracy theories, is highly controversial and reflects the actual U.S. political landscape.

CCS CONCEPTS

• **Computing methodologies** → **Natural language processing.**

KEYWORDS

Natural Language Processing, Transformers, BERT, COVID-datasets, Misinformation

ACM Reference Format:

Youri Peskine, Raphaël Troncy, and Paolo Papotti. 2023. Analyzing COVID-Related Social Discourse on Twitter using Emotion, Sentiment, Political Bias, Stance, Veracity and Conspiracy Theories. In *Companion Proceedings of the ACM Web Conference 2023 (WWW '23 Companion)*, April 30-May 4, 2023, Austin, TX, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3543873.3587622>

1 INTRODUCTION

As the amount of information shared online increase¹, we are prone to face more misinformation on the web. Events such as the 2016 U.S. Presidential Elections [1] or the Brexit [10] are prime examples of strongly discussed topics with large amount of false information

*Corresponding author

¹<https://www.domo.com/learn/infographic/data-never-sleeps-5>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

BeyondFacts'23, 30 April – 4 May 2023, Austin, Texas, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9419-2/23/04...\$15.00

<https://doi.org/10.1145/3543873.3587622>

shared online. Social media websites can have strong influence on shaping the beliefs of one individual, and can have consequences on real life topics such as politics [1, 10], science [26], economics² or health [15]. Considering that ‘fake-news’ tends to spread faster and wider than the truth [28], researchers have started to help fact-checkers scale-up their ability to verify information [18]. The need of such technology has been even more evident with the recent COVID-19 pandemic, with misinformation shared profusely online, and the World Health Organization (WHO) describing it as an *infodemic*³. According to [3], the number of fact-check reports rose by more than 900% between January and March 2020, reflecting the large amount of misinformation shared about COVID-19.

While some research is focused on *detecting misinformation*, in this work, we focus on better understanding the online discourse around COVID-19 on dimensions that go beyond misinformation classification. We explore the relationships among emotion, sentiment, political bias, stance, veracity and conspiracy theories, by leveraging a dataset for each textual feature. We use three datasets for training models that detect sentiment, emotion and political bias, and we use those models on the other datasets to study in detail their interactions. We then compute the conditional distribution of the labels between those features to analyze and share some insights about their relationships. Notable results show that political bias plays a role in the stance toward COVID-19 regulations and conspiracy theories or that emotion and sentiment are used by people who share potentially misleading content. Our results can be reproduced using the code available at <https://github.com/D2KLab/covid-twitter-discourse-analysis>.

2 RELATED WORK

Online misinformation during the COVID-19 pandemic has urged researchers to study its prevalence in social media websites, such as Twitter. Many datasets have been built around annotating textual features in tweets during the pandemic. In this work, we selected three different core datasets, each one allowing the training of a model for the detection of one textual feature: COVID LTSE Attributes (Emotion) [9], COVIDSenti (Sentiment) [19] and Russian Troll (Political Bias)⁴ [13]. We also selected three additional external datasets, which will only be used for the evaluation of the correlation: COVID 19 Stance (Stance) [8], Birdwatch (Veracity)

²<https://www.forbes.com/sites/kenrapoza/2017/02/26/can-fake-news-impact-the-stock-market/>

³<https://www.who.int/health-topics/infodemic>

⁴<https://fivethirtyeight.com/features/why-were-sharing-3-million-russian-troll-tweets/>

[23] and MediaEval-FND (Conspiracy theories) [21]. The three core datasets are also used for the evaluation of the correlation.

2.1 Datasets

2.1.1 [Core] COVID LTSE Attributes. The COVID LTSE Attributes dataset [9] contains 252 million of tweets from January 2020 to June 2021. The data was searched using keywords, such as ‘wuhan’ or ‘corona’. This dataset is labeled with 17 attributes, such as topics or emotions. In this work, we focus on the emotion attributes, which has been labeled using the *CrystalFeel*⁵ pre-trained machine learning algorithm.

2.1.2 [Core] COVIDSenti. COVIDSenti [19] is a dataset with labels for sentiment in COVID-related tweets. The tweets have been crawled from February 2020 to March 2020, using keywords such as ‘coronavirus’ or ‘Corona Outbreak’. The data is annotated with *TextBlob*⁶ using the methodology described in [2] resulting in 90,000 annotations.

2.1.3 [Core] Russian Troll. The Russian Troll dataset [13] contains 2.9 million of tweets from February 2012 to May 2018. It is the only dataset with no COVID-related data. The tweets are from accounts associated with the Internet Research Agency, which interfered during the U.S. 2016 presidential elections. A detailed analysis of the disinformation tactics used by this group of people is available in [12]. The data is labeled at the account level using five main categories (‘Right troll’, ‘Left troll’, ‘Fearmonger’, ‘HashtagGamer’ and ‘NewsFeed’). In this work, we focus on ‘Right Troll’ and ‘Left Troll’ while the remaining labels are labeled as ‘Other’.

2.1.4 [External] COVID 19 Stance. In the COVID 19 Stance dataset [8], tweets are labeled with a stance towards a topic related to the pandemic. The data was crawled from February 2020 to August 2020 using keywords (‘coronavirus’, ‘covid-19’, etc.) or hashtags (‘#lockdown’, ‘#washhands’, ‘#socialdistancing’, etc.). The topics are ‘Anthony S. Fauci’, ‘keeping schools closed’, ‘stay at home orders’ and ‘wearing a face mask’, and the annotation was done with Amazon Mechanical Turk. Since the release of this dataset, numerous tweets have been deleted or removed, and we were only able to retrieve 3,616 tweets.

2.1.5 [External] Birdwatch. Birdwatch⁷ is a crowdsourced fact-checking program that allow Twitter users to identify potentially misleading tweets. The dataset used in [23] contains a total of 9,851 tweets that have been labeled by Birdwatch users with ‘Misleading’ or ‘Not Misleading’, from January 2021 to September 2021. As tweets can have multiple notes from multiple Birdwatch users, with different labels, we performed majority-voting to assign a single label to the tweet. If majority cannot be reached, the tweet is discarded.

2.1.6 [External] MediaEval - FND. The MediaEval - FND dataset [21] contains 1,912 tweets about COVID-19-related conspiracy theories. The tweets are labeled as ‘Discussing’, ‘Promoting’ or not related to a given conspiracy theory. In the original dataset, each

tweet is annotated for nine different named conspiracy⁸. In this work, we merged the labels into three broad categories: ‘Not related to any conspiracy’, ‘Discussing at least one conspiracy’, ‘Promoting at least one conspiracy’.

2.2 Classification Models

Transformer-based models [27] have largely contributed to progress in many Natural Language Processing (NLP) tasks, including machine translation [4, 24], question answering [14, 30] and text classification [6, 16]. Most notably, BERT [5] has outperformed other methods, such as TF-IDF or Recurrent Neural Networks, while providing a pre-trained model that can be fine-tuned for specific tasks [25]. For example, Covid-Twitter-BERT (CT-BERT) [17] has been trained on textual data from Twitter during the COVID-19 pandemic, which improves results on domain-specific datasets.

3 METHODOLOGY

In order to detect correlations, we build three text-classification models on the core datasets mentioned in Section 2.1, for sentiment, emotion and political bias. In this section, we will discuss the methodology to train the different models on their respective data, and explain how we analyze potential correlations between the textual features.

3.1 Model Training

Models are trained to perform text classification using supervised learning. As some datasets are very large, we sample 25,000 tweets from the COVID LTSE Attributes dataset and 42,000 from the Russian Troll dataset. We first apply some basic pre-processing on the text, by removing links and special characters. We then split datasets into a train set and a validation set with a 80/20 stratified split ratio. Models have pre-trained CT-BERT weights and a classification layer depending on the number of output classes. They are trained using Adam [11] optimizer, with a weight decay of 1.10^{-2} and a learning rate of 1.10^{-5} for 25 epochs. We use a Cross Entropy loss with weights proportional to the inverse of the class distribution. We monitor the performance of the model with the F1 score on the validation set, and save the best performing model. We use the pytorch [20] and huggingface-transformers [29] libraries to implement our code. Results for the models during training are available in Section 4.1.

3.2 Correlation

In order to find possible correlations between the textual features, we used the different models to predict features on the other datasets. Core datasets are used for the training of the models and the prediction of other features, while external datasets are only used for prediction. For example, the model trained on COVIDSenti is used on the COVID LTSE Attributes dataset, the Russian Troll dataset, the COVID-19-Stance dataset, the Birdwatch dataset and the MediaEval-FND dataset. This way, we can analyze the conditional distribution of the predicted labels given the ground truth labels. The results and analysis are presented in Section 4.2.

⁵<https://socialanalyticsplus.net/crystalfeel/>

⁶<https://textblob.readthedocs.io/en/dev/>

⁷Recently renamed *Community Notes* <https://help.twitter.com/en/using-twitter/community-notes>

⁸List of all the conspiracy theories: Suppressed Cures, Behaviour and Mind Control, Antivax, Fake Virus, Intentional Pandemic, Harmful Radiation/Influence, Population Reduction/Control, New World Order, Satanism

Table 1: F1-Score of the models on a validation set

Core Dataset	F1-score
COVID LTSE Attributes	0.622
COVIDSenti	0.769
Russian Troll	0.636

4 RESULTS

In this section, we discuss the results of the models during the training (4.1), and the analysis of the possible correlations between the features (4.2). Some examples of tweets in the dataset are available in Table 2, highlighting a particular predicted label.

4.1 Model performance

As discussed in Section 3.1, we split all core datasets into training and validation sets with a 80/20 stratified split ratio. Table 1 shows the performance of the models on the validation set.

The model based on the COVIDSenti dataset obtains the best score on its own evaluation set. This might be expected as sentiment detection is arguably the easiest task of the three. The overall performance of the models are fair, given the noise in the datasets, as COVID LTSE Attributes and COVIDSenti have been automatically annotated, and the Russian Troll dataset has been labelled at the user level, resulting in some generic tweets having annotations towards political bias.

4.2 Correlation Analysis

In order to detect some correlations between the studied textual features, we compute the matrix of frequency of the labels of two textual features in the data. The y-axis label⁹ is the ground truth of the corresponding dataset, while the x-axis label represents the prediction of the model. Rows have been normalized to represent the conditional distribution of the predicted label given a ground-truth label.

4.2.1 Sentiment feature. Figure 1 shows the correlation between the sentiment feature and the other features. We can see in Figure 1a 1b 1c 1d that people against the mentioned topics tend to share a more negative sentiment. This is especially true for the topic 'Face masks', and less apparent for the topic 'School closures'. People in favor of the 'Face masks' and 'School closures' topics tend to use more negative sentiment as well. However, people use more positive sentiment when supporting 'Stay at home' orders.

Figure 1e shows that all emotions except happiness tend to be more negative than positive, which is expected. It also shows that anger is the emotion where negative sentiment is the most prevalent. Figure 1f shows that tweets that have a political bias use more sentiment (positive and negative) than other tweets. However the distribution of sentiment is the same for both Left and Right bias.

Figure 1g shows that tweets that share potentially misleading content tend to use slightly more negative sentiment than the non-misleading tweets. Very few tweets share positive sentiment in this dataset overall, suggesting that people on Birdwatch are more

⁹The labels 'N', 'H', 'A', 'S', 'F' represent the following emotions: 'None', 'Happiness', 'Anger', 'Sadness', 'Fear'.

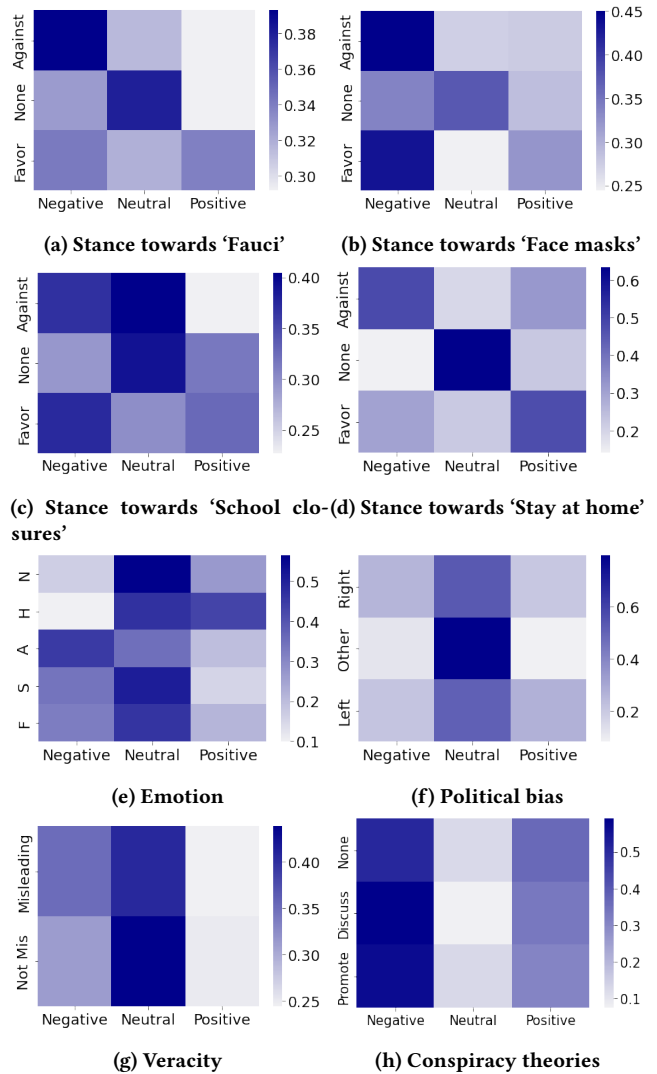


Figure 1: Distribution of the labels for the sentiment feature

interested in labeling negative tweets. However, conspiracy theories do not seem to be particularly correlated with sentiment, as shown in Figure 1h.

4.2.2 Emotion feature. Figure 2 shows how the emotion feature is correlated to the other features. First, it seems clear in Figure 2a 2b 2c 2d that the four topics 'Fauci', 'Face masks', 'School closures' and 'Stay at home' are quite controversial on Twitter, with the anger emotion dominating almost all stances. The sadness emotion is the most used when discussing the topic of 'School closures', showing empathy for the teachers and the children. The 'Stay at home' topic sees more happiness in the tweets, with people enjoying working from home.

Figure 2e shows that a majority of tweets from the COVIDSenti dataset use the fear emotion, even in positive tweets. This seems counter-intuitive and may be due to having numerous tweets about

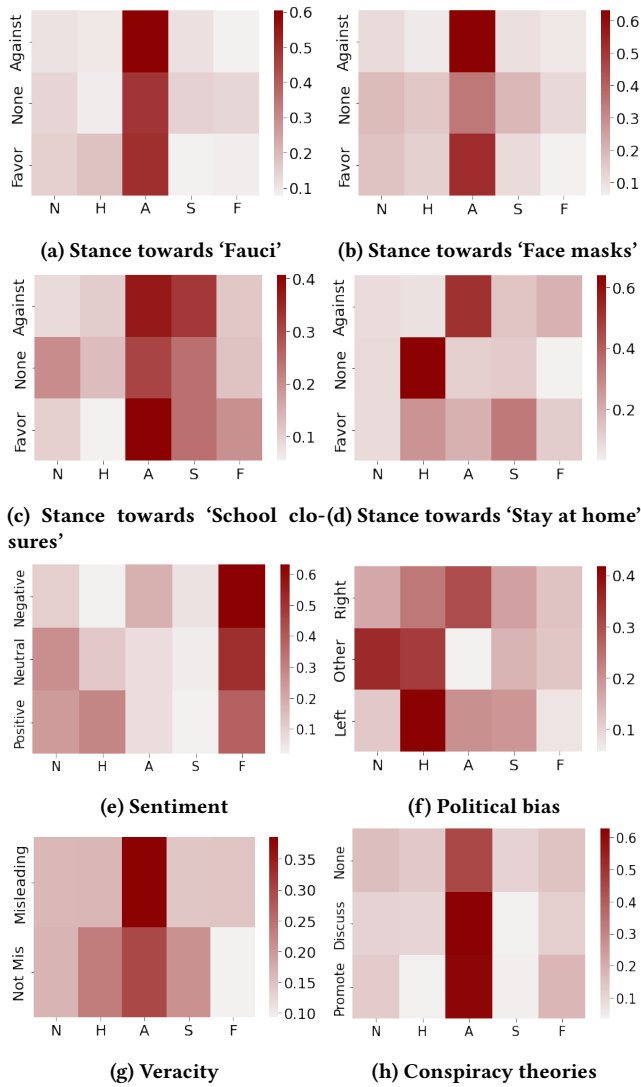


Figure 2: Distribution of the labels for the emotion feature

wishing people to stay safe, in fear of covid. However, positive tweets also use happiness a lot, which is to be expected.

Political biased tweets are more likely to have an emotion than not, as shown in Figure 2f. Tweets from users tagged as having left bias tend to contain more happiness, while tweets from users having right bias tend to contain more anger.

Regarding veracity, in Figure 2g, anger is dominating the tweets sharing potentially misleading information, while emotion are slightly more even on not misleading tweets. In Figure 2h, we can see that emotion and conspiracy theories are not heavily correlated. We notice a slight decrease in anger in non-conspiracist tweets.

4.2.3 Political bias feature. Lastly, we analyze correlation between political bias and other textual features, highlighted in Figure 3. We again notice that some topics are controversial, for example, 'Face masks' and 'Stay at home'. In those topics, we see that people against



Figure 3: Distribution of the labels for the political bias feature

(face masks) have more right political bias and people in favor have more left political bias. This reflects the U.S. political landscape during the pandemic, as Democrats governors had generally more strict mandates towards wearing face masks than their Republican counterparts [7]. The topic of school closures and re-opening was also highly controversial, with Republicans leaning toward having more in-person classes and Democrats toward having more online-classes¹⁰.

Figures 3e 3f show that specific sentiment and emotion are not strongly correlated to one political bias or the other. However, Figure 3g shows that potentially misleading tweets are not tied to one political bias, while non-misleading tweets are more likely to be shared with left political bias. Moreover, Figure 3h shows

¹⁰<https://www.pewresearch.org/fact-tank/2020/08/05/republicans-democrats-differ-over-factors-k-12-schools-should-consider-in-deciding-whether-to-reopen/>

that tweets discussing conspiracy theories are more likely to have left political bias, while it is the opposite for tweets promoting conspiracy theories. This supports the findings in [22], which states that conservatives tend to share more anti-science information than pro-science, thus being more inclined towards conspiracy theories.

5 CONCLUSION

In this paper, we analyzed the correlations between five textual features: emotion, sentiment, political bias, stance and veracity. We leveraged relevant datasets to train three models to predict emotion, sentiment and political bias on COVID-19 related tweets. These models allowed us to analyze the conditional distribution of the different labels to better understand the online discourse. Main findings include that COVID-19-related regulations topics, such as 'Face masks', 'School closures' or 'Stay at home orders' are highly controversial, generating a lot of negative sentiment and anger emotion in the Twitter discourse. The users' political bias on those topics also outlined the stance of US politicians in the debate. Similarly, conspiracy theories are usually promoted with negative sentiment and right political bias, which might reflect the inclination of conservatives towards anti-science information.

ACKNOWLEDGMENTS

This work has been partially supported by CHIST-ERA within the CIMPLE project (CHIST-ERA-19-XAI-003).

REFERENCES

- [1] Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives* 31, 2 (2017), 211–236.
- [2] Ajay Bandi and Aziz Fellah. 2019. Socio-Analyzer: A Sentiment Analysis Using Social Media Data. In *Proceedings of 28th International Conference on Software Engineering and Data Engineering (EPIC Series in Computing, Vol. 64)*, Frederick Harris, Sergiu Dascalu, Sharad Sharma, and Rui Wu (Eds.). EasyChair, 61–67. <https://doi.org/10.29007/kzkl>
- [3] J S Brennen, F M Simon, P N Howard, and R K Nielsen. 2020. Types, sources, and claims of COVID-19 misinformation. *Reuters Institute for the Study of Journalism* (2020).
- [4] Raj Dabre, Chenhui Chu, and Anoop Kunchukuttan. 2020. A Survey of Multilingual Neural Machine Translation. *ACM Comput. Surv.* 53, 5, Article 99 (sep 2020), 38 pages. <https://doi.org/10.1145/3406095>
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [6] Ashraf Elnagar, Ridhwan Al-Debsi, and Omar Einea. 2020. Arabic text classification using deep learning models. *Information Processing & Management* 57, 1 (2020), 102121. <https://doi.org/10.1016/j.ipm.2019.102121>
- [7] Michele Gelfand, Ren Li, Eftychia Stamkou, Dylan Pieper, Emmy Denison, Jessica Fernandez, Virginia Choi, Jennifer Chatman, Joshua Jackson, and Eugen Dimant. 2022. Persuading republicans and democrats to comply with mask wearing: An intervention tournament. *Journal of Experimental Social Psychology* 101 (2022), 104299.
- [8] Kyle Glandt, Sarthak Khanal, Yingjie Li, Doina Caragea, and Cornelia Caragea. 2021. Stance Detection in COVID-19 Tweets. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Online, 1596–1611. <https://doi.org/10.18653/v1/2021.acl-long.127>
- [9] Raj Gupta, Ajay Vishwanath, and Yiping Yang. 2022. COVID-19 Twitter Dataset with Latent Topics, Sentiments and Emotions Attributes. <https://doi.org/10.3886/E120321V12>
- [10] Max Hänska and Stefan Bauchowitz. 2017. Tweeting for Brexit: how social media influenced the referendum. (2017).
- [11] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1412.6980>
- [12] Darren L. Linvill, Brandon C. Boatwright, Will J. Grant, and Patrick L. Warren. 2019. "THE RUSSIANS ARE HACKING MY BRAIN!" investigating Russia's internet research agency twitter tactics during the 2016 United States presidential campaign. *Computers in Human Behavior* 99 (2019), 292–300. <https://doi.org/10.1016/j.chb.2019.05.027>
- [13] Darren L. Linvill and Patrick L. Warren. 2020. Troll Factories: Manufacturing Specialized Disinformation on Twitter. *Political Communication* 37, 4 (2020), 447–467. <https://doi.org/10.1080/10584609.2020.1718257> arXiv:<https://doi.org/10.1080/10584609.2020.1718257>
- [14] Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. 2018. The Natural Language Decathlon: Multitask Learning as Question Answering. <https://doi.org/10.48550/ARXIV.1806.08730>
- [15] Areeb Mian and Shujhat Khan. 2020. Coronavirus: the spread of misinformation. *BMC Medicine* 18, 1 (18 Mar 2020), 89. <https://doi.org/10.1186/s12916-020-01556-3>
- [16] Shervin Minaee, Nal Kalchbrenner, Erik Cambria, Narjes Nikzad, Meysam Chenaghlu, and Jianfeng Gao. 2021. Deep Learning–Based Text Classification: A Comprehensive Review. *ACM Comput. Surv.* 54, 3, Article 62 (apr 2021), 40 pages. <https://doi.org/10.1145/3439726>
- [17] Martin Müller, Marcel Salathé, and Per E Kummervold. 2020. COVID-TwitterBERT: A Natural Language Processing Model to Analyse COVID-19 Content on Twitter. <https://arxiv.org/abs/2005.07503>
- [18] Preslav Nakov, David Corney, Maram Hasanain, Firoj Alam, Tamer Elsayed, Alberto Barrón-Cedeño, Paolo Papotti, Shaden Shaar, and Giovanni Da San Martino. 2021. Automated Fact-Checking for Assisting Human Fact-Checkers. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Zhi-Hua Zhou (Ed.). International Joint Conferences on Artificial Intelligence Organization, 4551–4558. <https://doi.org/10.24963/ijcai.2021/619> Survey Track.
- [19] Usman Naseem, Imran Razzak, Matloob Khushi, Peter W Eklund, and Jinman Kim. 2021. COVIDSenti: A large-scale benchmark Twitter data set for COVID-19 sentiment analysis. *IEEE Transactions on Computational Social Systems* 8, 4 (2021), 1003–1015.
- [20] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. Curran Associates Inc., Red Hook, NY, USA.
- [21] Konstantin Pogorelov, Daniel Thilo Schroeder, Stefan Brenner, Asep Maulana, and Johannes Langguth. 2022. Combining Tweets and Connections Graph for FakeNews Detection at MediaEval 2022. In *Multimedia Benchmark Workshop*.
- [22] Ashwin Rao, Fred Morstatter, Minda Hu, Emily Chen, Keith Burghardt, Emilio Ferrara, and Kristina Lerman. 2021. Political partisanship and antiscience attitudes in online discussions about COVID-19: Twitter content analysis. *J. Med. Internet Res.* 23, 6 (June 2021), e26692.
- [23] Mohammed Saeed, Nicolas Traub, Maëlle Nicolas, Gianluca Demartini, and Paolo Papotti. 2022. Crowdsourced Fact-Checking at Twitter: How Does the Crowd Compare With Experts?. In *31st ACM International Conference on Information & Knowledge Management (Atlanta, GA, USA) (CIKM '22)*. Association for Computing Machinery, New York, NY, USA, 1736–1746. <https://doi.org/10.1145/3511808.3557279>
- [24] Felix Stahlberg. 2020. Neural machine translation: A review. *Journal of Artificial Intelligence Research* 69 (2020), 343–418.
- [25] Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to Fine-Tune BERT for Text Classification?. In *Chinese Computational Linguistics*, Maosong Sun, Xuanjing Huang, Heng Ji, Zhiyuan Liu, and Yang Liu (Eds.). Springer International Publishing, Cham, 194–206.
- [26] Kathie M d'I Treen, Hywel TP Williams, and Saffron J O'Neill. 2020. Online misinformation about climate change. *Wiley Interdisciplinary Reviews: Climate Change* 11, 5 (2020), e665.
- [27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems (NeurIPS)*. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>
- [28] Sorous Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *science* 359, 6380 (2018), 1146–1151.
- [29] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 38–45. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>

Table 2: Examples of tweets from all the datasets. Ground Truth indicates the label of the tweet in its original dataset while Predicted Label is the output of one of the trained models

	Tweets	Ground Truth	Predicted Label
(a)	<i>Idc what you say, you're selfish if you refuse to wear a mask. This shouldn't be political. #MaskUp #MaskMoaners</i>	Favor 'Face masks'	Negative sentiment
(b)	<i>@kylegriffin1 Close the damn schools until there is a vaccine. #NotMyChild</i>	Favor 'School closures'	Negative sentiment
(c)	<i>If grocery stores can be open and people can risk their lives working there, then so can the schools and teachers. #OpenSchools</i>	Against 'School closures'	Neutral sentiment
(d)	<i>Corona virus Day 4 diary entry: I have now been social distancing for the past 26 years.</i>	Sadness	Neutral sentiment
(e)	<i>Italy Declares State of Emergency Over Wuhan Coronavirus</i>	Fear	Neutral sentiment
(f)	<i>Biden blames rise of COVID-19 cases on the unvaccinated: "This is a pandemic of the unvaccinated."</i>	Not Misleading	Neutral sentiment
(g)	<i>@saraecook Fauci is such a hypocrite! He knew back during the SARS outbreak most people who died was due largely to cytokine storm. Much the same with Coronavirus. He had no problem with Hydroxychloroquine being used then. #FauciFraud</i>	Against 'Fauci'	Anger emotion
(h)	<i>Trump and the White House are straight up publicly attacking the country's leading infectious disease expert during a #pandemic that has already killed nearly 140,000 Americans. Yup, that tracks. #COVID19 #DrFauci</i>	Favor 'Fauci'	Anger emotion
(i)	<i>The policymakers need to consider the fact that schools can't run without fees and teachers can't survive without salary. #SaveOurSchools</i>	Against 'School closures'	Sadness emotion
(j)	<i>@GrandadJohn5 Good news that County cricket is starting up but no news on recreational cricket the cut off point for our league is August 8th after that only friendlies or right off the season #StaySafeStayHome</i>	No stance towards 'Stay at home'	Happiness emotion
(k)	<i>Coronavirus: is it safe to travel and should children be kept home?</i>	Positive sentiment	Fear emotion
(l)	<i>I had such a good day with the students at @UNCG and @ncatsuaggies discussing activism, social justice, & organizing. They were incredible.</i>	Left political bias	Happiness emotion
(m)	<i>THE ATTACK ON FREEDOM OF SPEECH CONTINUES! #CrookedHillary will destroy the 1st Amendment Right of her Opposition!</i>	Right political bias	Anger emotion
(n)	<i>Take sports away and Social Interaction in schools. Your kids will have a great immune system! Way to teach your kids your saving them from the coronavirus. Bill gates and all Ted Talk technocrats have wanted online learning for years. Wake up! #OpenSchools</i>	Against 'School closures'	Right political bias
(o)	<i>Wearing a mask and social distancing doesn't mean you are "living in fear." It's like wearing a seat belt or using your headlights in the rain, it's for your safety and the safety of others. #WearADamnMask</i>	Favor 'Face masks'	Left political bias
(p)	<i>@Athens108 @realDonaldTrump It may have worked for an old Coronavirus that is a different virus from COVID19. All the studies coming out says chloroquine does not work for this virus. Dr. Fauci is always on the frontlines for all viruses. #FauciIsAHero</i>	Favor 'Fauci'	Right political bias
(q)	<i>87% of the deaths were caused by democrat leadership. Things like forcing nfected patents into nursing homes by executive order and banning HCQ. We now know that HCQ could have easily saved over 100000 lives over 20000 in NY alone. Trump was right. Democrats own the pandemic</i>	Promote conspiracy	Right political bias

[30] Zhou Yu, Jun Yu, Yuhao Cui, Dacheng Tao, and Qi Tian. 2019. Deep Modular Co-Attention Networks for Visual Question Answering. In *Proceedings of the*