

# A Second Pandemic? Analysis of Fake News About COVID-19 Vaccines in Qatar

Preslav Nakov,<sup>1</sup> Firoj Alam,<sup>1</sup> Shaden Shaar,<sup>1</sup>  
Giovanni Da San Martino<sup>2</sup> and Yifan Zhang<sup>1</sup>

<sup>1</sup> Qatar Computing Research Institute, HBKU, Qatar

<sup>2</sup> University of Padova, Italy

{pnakov, fialam, sshaar, yzhang}@hbku.edu.qa,  
dasan@math.unipd.it

## Abstract

While COVID-19 vaccines are finally becoming widely available, a second pandemic that revolves around the circulation of anti-vaxxer “fake news” may hinder efforts to recover from the first one. With this in mind, we performed an extensive analysis of Arabic and English tweets about COVID-19 vaccines, with focus on messages originating from Qatar. We found that Arabic tweets contain a lot of false information and rumors, while English tweets are mostly factual. However, English tweets are much more propagandistic than Arabic ones. In terms of propaganda techniques, about half of the Arabic tweets express doubt, and 1/5 use loaded language, while English tweets are abundant in loaded language, exaggeration, fear, name-calling, doubt, and flag-waving. Finally, in terms of framing, Arabic tweets adopt a health and safety perspective, while in English economic concerns dominate.

## 1 Introduction

During the COVID-19 pandemic, social media have become one of the main communication channels for information dissemination and consumption, and many people rely on them as their primary source of news (Perrin, 2015), attracted by the broader choice of information sources. Unfortunately, over time, social media have also become one of the main channels to spread disinformation. To tackle this issue, a number of (mostly manual) fact-checking initiatives have been launched, and there are over 200 fact-checking organizations currently active worldwide.<sup>1</sup> However, these efforts are insufficient, given the scale of disinformation, which, in the time of COVID-19, has grown into the *First Global Infodemic* (according to the World Health Organization).

<sup>1</sup><http://tiny.cc/zd1fnz>



Figure 1: Tweets about COVID-19 and vaccines.

Figure 1 shows examples of how online users discuss COVID-19 and vaccines. We can see that the problem goes beyond factuality: there are tweets spreading rumors, discussing action taken, instilling panic, making jokes, and promoting bad cure.

For the tweets in Figure 1, we might want to know whether they are factual, harmful, calling for action, etc. (see Section 4.1). It is also important to understand whether the content of the tweet is propagandistic (Section 4.2), what propaganda techniques are used, as well as the way the issue is framed (Section 4.3). Doing this in a timely manner is crucial to help organizations channel their efforts, and to counter the spread of disinformation, which may cause panic, mistrust, and other problems.

With this in mind, we performed an extensive analysis of Arabic and English tweets about COVID-19 and vaccines, with focus on messages originating from Qatar. Our analysis focuses on (i) COVID-19 disinformation, (ii) propaganda and its techniques, and (iii) framing.

Our contributions can be summarized as follows:

- We build and release a dataset of tweets related to COVID-19 and vaccines in Arabic and English.<sup>2</sup>
- We analyze the tweets from various perspectives (factuality, harmfulness, propaganda, and framing), and we discuss some interesting observations from our analysis.

## 2 Related Work

Below, we discuss relevant research directions.

### 2.1 Factuality

Work on fighting disinformation and misinformation online has focused on fact-checking and fake news detection (Li et al., 2016; Hardalov et al., 2016; Shu et al., 2017; Karadzhov et al., 2017; Lazer et al., 2018; Mihaylova et al., 2018; Vosoughi et al., 2018; Vo and Lee, 2018; Atanasova et al., 2019; Baly et al., 2019; Zlatkova et al., 2019; Baly et al., 2020; Nguyen et al., 2020; Shaar et al., 2020a). Research was further enabled by the emergence of datasets (Wang, 2017; Augenstein et al., 2019), often released as part of evaluation campaigns (Derczynski et al., 2017; Nakov et al., 2018; Da San Martino et al., 2019; Elsayed et al., 2019; Gorrell et al., 2019; Mihaylova et al., 2019; Barrón-Cedeño et al., 2020; Nakov et al., 2021c,d; Shaar et al., 2021b). As automated systems have credibility issues (Arnold, 2020), another research direction has emerged: building tools to facilitate human fact-checkers (Nakov et al., 2021b).

<sup>2</sup><http://gitlab.com/sshaar/a-second-pandemic.-analysis-of-fake-news-about-covid-19-vaccines-in-qatar>

### 2.2 Check-Worthiness Estimation

Given the volume of claims appearing in social media posts or in political statements, a problem that is crucial for fact-checkers is to identify which claims should be prioritized for fact-checking. The ClaimBuster system (Hassan et al., 2015) was a pioneering work in that direction. It categorized a political statement as *non-factual*, *unimportant factual*, or *check-worthy factual*. Gencheva et al. (2017) also focused on the 2016 US Presidential debates, for which they obtained binary (*check-worthy* vs. *non-check-worthy*) labels based on the fact-checking decisions of nine fact-checking organizations. An extension of this work was the ClaimRank system, which supports both English and Arabic (Jaradat et al., 2018). Note that political debates and speeches require modeling the context of the target sentence to classify. Indeed, context was a major focus for most research in the debates domain (Gencheva et al., 2017; Patwari et al., 2017; Vasileva et al., 2019; Shaar et al., 2021a). For example, Vasileva et al. (2019) modeled context in a multi-task learning neural network that predicts whether a sentence would be selected for fact-checking by each fact-checking organization (from a set of nine such organizations).

There has also been research on detecting check-worthy claims *in social media* (as opposed to the above research, which targeted political debates and speeches), featuring tweets about COVID-19 or general topics in Arabic and English (Hasanain et al., 2020; Shaar et al., 2020b, 2021c).

More directly related to our work here is the work of Alam et al. (2021c) and Alam et al. (2021a), who developed a multi-question annotation schema of tweets about COVID-19, organized around seven questions that aim to model the perspective of journalists, fact-checkers, social media platforms, policymakers, and the society. In our experiments, we use their schema and data to train classifiers for part of our analysis.

### 2.3 Propaganda

Propaganda is a communication tool that is deliberately designed to influence the opinions and the actions of other people in order to achieve a pre-determined objective. When automatic means are being used to spread such influencing messages on social media platforms, this is referred to as *computational propaganda* (Woolley and Howard, 2018).

Research on propaganda detection has focused on textual content (Barrón-Cedeno et al., 2019; Rashkin et al., 2017; Da San Martino et al., 2019; Da San Martino et al., 2020a). Suitable datasets were made available by Rashkin et al. (2017) and Barrón-Cedeno et al. (2019), where the documents (news articles) were annotated using distant supervision, according to the reputation of their source, as judged by journalists. Rashkin et al. (2017) focused on analyzing the language of propaganda (vs. trusted, satire, and hoaxes) based on LIWC lexicons, while Barrón-Cedeno et al. (2019) studied a variety of stylistic features.

Habernal et al. (2017, 2018) developed a corpus annotated with five fallacies, including *ad hominem*, *red herring*, and *irrelevant authority*. Fine-grained propaganda analysis was done by Da San Martino et al. (2019), who developed a corpus of news articles annotated with 18 propaganda techniques. Subsequently, the Prta system was released (Da San Martino et al., 2020b), and improved models were proposed, focusing on interpretability (Yu et al., 2021) or addressing the limitations of transformers (Chernyavskiy et al., 2021). Finally, multimodal content was explored in memes using 22 propaganda techniques (Dimitrov et al., 2021a,b).

## 2.4 Framing

Framing refers to representing different salient aspects and perspectives for the purpose of conveying the latent meaning about an issue (Entman, 1993). Recent work on automatically identifying media frames includes developing coding schemes and semi-automated methods (Boydston et al., 2013), datasets such as the Media Frames Corpus (Card et al., 2015), and systems to automatically detect media frames (Liu et al., 2019; Zhang et al., 2019), large-scale automatic analysis of news articles (Kwak et al., 2020), and semi-supervised approaches (Cheeks et al., 2020).

## 2.5 Fighting the COVID-19 Infodemic

Related work on fighting the COVID-19 infodemic includes developing multi-question annotation schemes of tweets about COVID-19 (Alam et al., 2021c,a), studying credibility (Cinelli et al., 2020; Pulido et al., 2020; Zhou et al., 2020), racial prejudices and fear (Medford et al., 2020; Vidgen et al., 2020), situational information, e.g., caution and advice (Li et al., 2020), as well as detecting mentions and stance with respect to known misconceptions (Hossain et al., 2020).

## 3 Dataset

We collected Arabic tweets from February 2020 till March 2021. For the English tweets, we had two separate time periods (before and after COVID-19 vaccines became available): (i) from February till August 2020 (644 tweets), and (ii) from November 2020 till January 2021 (1,945 tweets). We used the following keywords to collect the tweets:

**English:** #covid19, #CoronavirusOutbreak, #Coronavirus, #Corona, #CoronaAlert, #CoronaOutbreak, Corona, covid-19, COVID vaccine, Covid-19 vaccine, #covidvaccine, corona vaccine, #vaccinate, #vaccine, vaccine  
**Arabic:** #كورونا, كورونا (Corona), #فيروس\_كورونا\_المستجد, فيروس\_كورونا\_الجديد, #فيروس\_كورونا, (novel Coronavirus), #كورونا, #كورونا\_الجديد (new Corona), مطعوم, تطعيم, لقاح, لقاءات,

We collected original tweets (no retweets or replies), we removed the duplicates using the similarity-based approach in Alam et al. (2021b), and we filtered out tweets with less than five words. Finally, we kept the most frequently liked and retweeted tweets for annotation. Our final corpus consists of 606 Arabic and 2,589 English tweets.

## 4 Method

Figure 2 shows the architecture of our system. Below, we discuss each analysis step in detail.

### 4.1 Disinformation Analysis

For disinformation analysis, we used the dataset from (Alam et al., 2021c,a), which is organized around seven questions: asking whether the tweet (Q1) contains a verifiable factual claim, (Q2) is likely to contain false information, (Q3) is of interest to the general public, (Q4) is potentially harmful to a person, a company, a product, or the society, (Q5) requires verification by a fact-checker, (Q6) poses harm to society, or (Q7) requires the attention of policy makers. The dataset consist of 504 English and 218 Arabic tweets, and we used it to train an SVM classifier, whose hyper-parameters we optimized using 10-fold cross-validation. Table 1 shows the performance of the classifier for English and Arabic for all questions. Note the multiclass nature of the tasks and the skewed class distribution for Q2 to Q6 (Alam et al., 2021a).

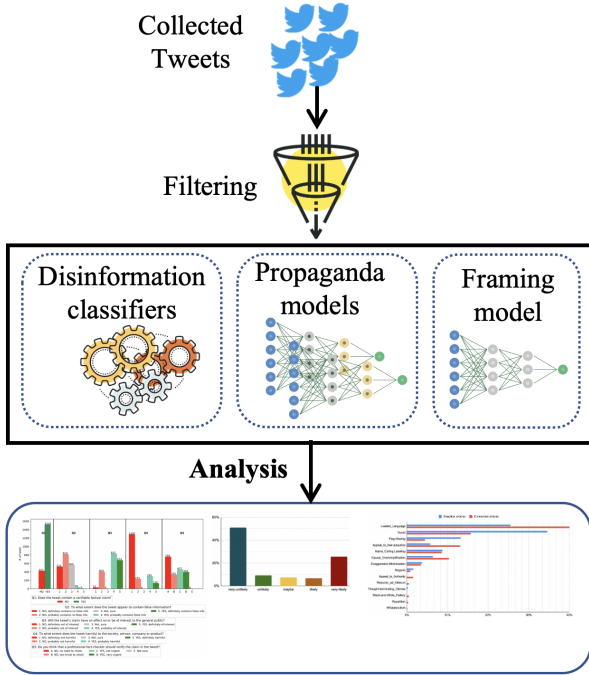


Figure 2: The architecture of our system for tweet analysis. The arrows show the information flow.

Q (# Cl.)	English		Arabic	
	Acc	W-F1	Acc	W-F1
Q1 (2)	64.5	64.8	72.5	72.9
Q2 (5)	40.0	41.1	44.3	43.3
Q3 (5)	42.0	41.7	51.4	49.1
Q4 (5)	41.6	41.5	57.1	56.4
Q5 (5)	37.0	37.6	57.1	57.4
Q6 (8)	50.0	50.4	69.7	68.6
Q7 (10)	60.7	58.6	68.8	69.1

Table 1: Performance of our models for disinformation analysis. Here, # Cl. shows the number of classes for the corresponding question.

## 4.2 Propaganda Analysis

For the propaganda analysis, we used two systems: Proppy (Barrón-Cedeño et al., 2019) and Prta (Da San Martino et al., 2020b).

**Proppy** uses a maximum entropy classifier trained on 51k articles, represented with various style-related features, such as character  $n$ -grams and a number of vocabulary richness and readability measures. The performance of the model is 82.89 in terms of F1 score, as evaluated on a separate test set of 10k articles. It outputs the following propaganda labels based on the output score  $p \in [0, 1]$ : *very unlikely* ( $0.0 \leq p < 0.2$ ), *unlikely* ( $0.2 \leq p < 0.4$ ), *somehow* ( $0.4 \leq p < 0.6$ ), *likely* ( $0.6 \leq p < 0.8$ ), and *very likely* ( $0.8 \leq p \leq 1.0$ ).

The **Prta system** offers a fragment-level and a sentence-level classifiers. They were trained on a corpus of 350K tokens. The performance of the sentence-level classifier is 60.71 in terms of F1 score. The fragment-level classifier identifies the text fragments and the propaganda techniques that occur in them. They consider the following 18 techniques: (i) Loaded language, (ii) Name calling or labeling, (iii) Repetition, (iv) Exaggeration or minimization, (v) Doubt, (vi) Appeal to fear/prejudice, (vii) Flag-waving, (viii) Causal oversimplification, (ix) Slogans, (x) Appeal to authority, (xi) Black-and-white fallacy, dictatorship, (xii) Thought-terminating cliché, (xiii) Whataboutism, (xiv) Reductio ad Hitlerum, (xv) Red herring, (xvi) Bandwagon, (xvii) Obfuscation, intentional vagueness, confusion, and (xviii) Straw man.

Note that both Proppy and Prta were developed for English. Thus, for the classification of Arabic content, we first translated it to English using the Google translation API, and then we ran the tools.

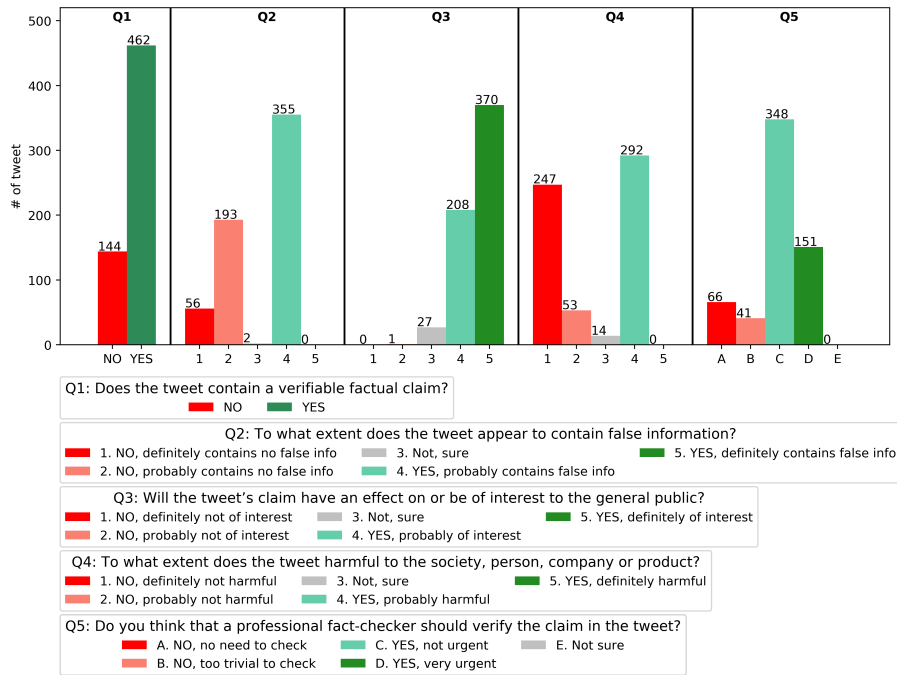
## 4.3 Framing

We used the *Tanbih Framing Bias Detection system* (Zhang et al., 2019), trained on the Media Frames Corpus (11k training news articles) by fine-tuning BERT to detect topic-agnostic media frames, achieving accuracy of 66.7% on the test set (1,138 news articles). It can predict the following 15 frames: (i) Economy, (ii) Capacity and resources, (iii) Morality, (iv) Fairness and equality, (v) Legality, constitutionality and jurisprudence, (vi) Policy prescription and evaluation, (vii) Crime and punishment, (viii) Security and defense, (ix) Health and safety, (x) Quality of life, (xi) Cultural identity, (xii) Public opinion, (xiii) Politics, (xiv) External regulation and reputation, and (xv) Other.

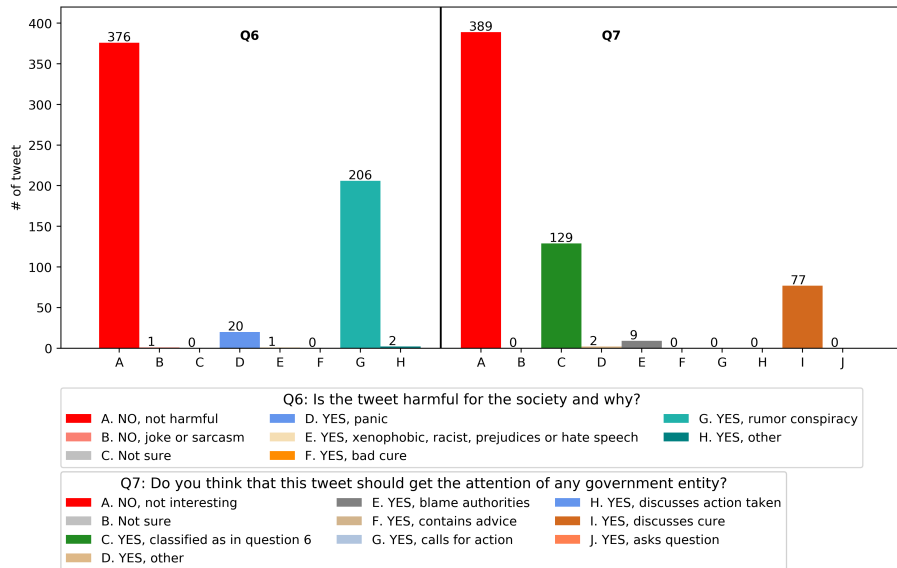
## 5 Results and Discussion

### 5.1 Disinformation Analysis

**Arabic:** Figure 3 shows the distribution for the questions for Arabic. We can see that (i) most tweets contain a verifiable factual claim, (ii) about half of the tweets contain false information, (iii) most tweets are of general interest to the public, (iv) about half of the tweets are harmful to the society, a person, a company, or a product (Question 6), (v) many tweets are worth fact-checking, (vi) most tweets are not harmful to the society, and many spread rumors, and (vii) some tweets discuss possible cure, and very few spread panic.



(a) Questions Q1-Q5.



(b) Questions Q6-Q7.

Figure 3: Statistics about the distribution of the Arabic tweets from February 2020 till March 2021.

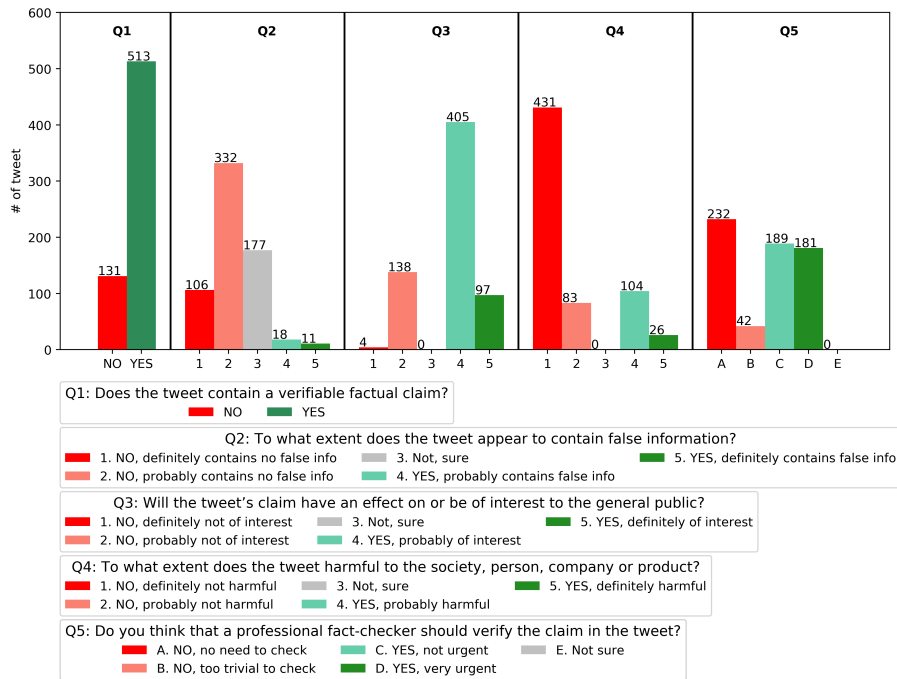
**English:** Figure 4 shows the distribution for the English tweets from February till August 2020. We can see that most tweets contain a verifiable factual claim, contain no false information, are of general interest to the public, are not harmful, and are worth fact-checking. Moreover, many tweets contain jokes, some contain rumors, and some blame the authorities.

We also analyzed the English tweets from November 2020 till January 2021. The results are shown in Figure 5, and follow a very similar trend.


**Summary:** Arabic tweets contain relatively more false information and rumors, some discuss possible cure, and very rarely spread panic. English tweets contain mostly factual statements, many make jokes, and rarely spread rumors.

## 5.2 Propaganda Analysis

Figure 6 shows the propaganda analysis in Arabic vs. English tweets. We can see that Arabic propagandistic tweets are extremely rare, while for English they are about 33% of all tweets.



(a) Questions Q1-Q5.



(b) Questions Q6-Q7.

Figure 4: Statistics about the distribution of the **English tweets** collected from February 2020 till August 2020.

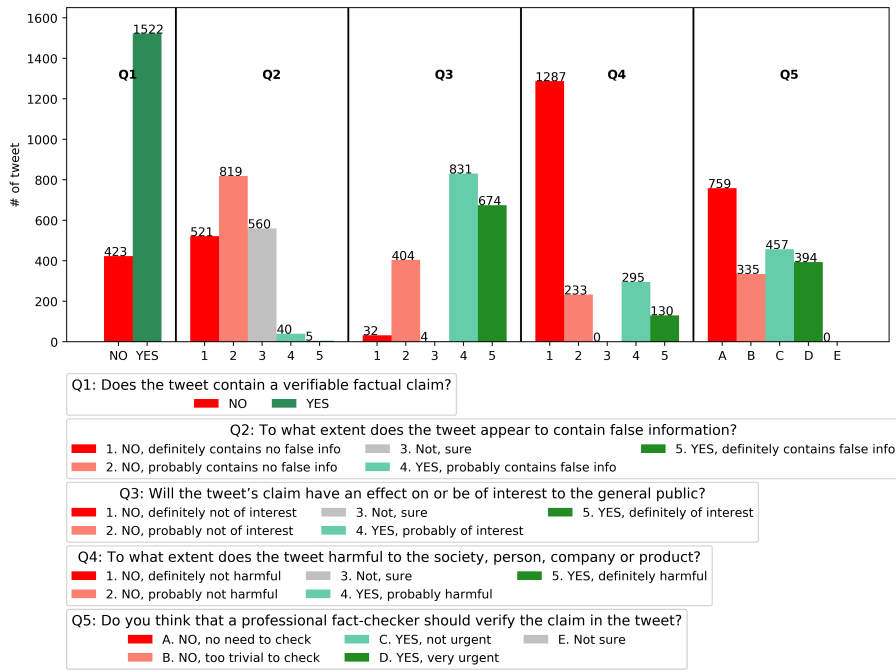
We also analyzed English tweets collected from November 2020 till January 2021, to cover tweets about COVID-19 vaccines, and we found that there were fewer propagandistic tweets: about 25%.

**Fine-Grained Propaganda Analysis** Next, we aimed to detect the specific propaganda techniques used in the tweets. Figure 7 shows the top propaganda techniques for Arabic and English.


We can see that, for Arabic, 50% of the tweets express *doubt*, and 20% use *loaded language*.

For English, we see a different distribution: about 33% of the tweets use *loaded language*, while each of the following techniques appears in about 10% of the tweets: *exaggeration*, *fear*, *name-calling*, *doubt*, and *flag-waving*.

Yet another trend is observed for English tweets collected from November 2020 till January 2021 (discussing vaccines): 50% of the tweets use *loaded language*, and each of the following four techniques appears in about 10% of the tweets: *flag-waving*, *name-calling*, and *exaggeration*.



(a) Questions Q1-Q5.



(b) Questions Q6-Q7.

Figure 5: Statistics about the distribution of the **English tweets** collected from November 2020 till January 2021.

### 5.3 Framing

Finally, we performed analysis in terms of framing, which reflects the perspective taken in the COVID-19 related tweets we analyzed. The results are shown on Figure 8.

We can see that in the Arabic tweets *health and safety* is the dominant frame, with *economy* coming second, and *cultural identity* being third.

For English, in both studied time periods, *economy* is the primary frame, and *health and safety* comes second.

We speculate that the difference in framing between Arabic and English tweets reflects the perspective of Qatari locals (who tweet primarily in Arabic) vs. that of expats (who tweet primarily in English). Thus, it is to be expected that the former are concerned primarily with health aspects (e.g., COVID-19 vaccination, social distancing, and other measures to keep one safe during the pandemic), while the latter worry more about the economic consequences of the pandemic (and respectively, about the security of their jobs).

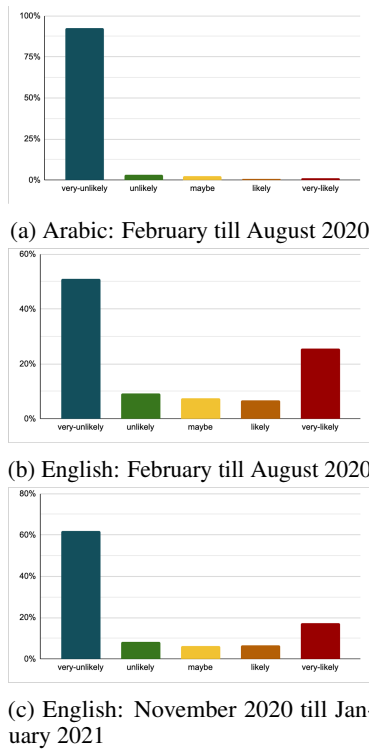


Figure 6: Analysis: use of propaganda.

## 6 Conclusion and Future Work

We have presented our analysis of COVID-19 tweets in Arabic and English aiming to help in the fight against the global infodemic, which emerged as a result of the COVID-19 pandemic. In particular, we collected tweets in different time frames starting from February 2020 till January 2021, and we analyzed them using different aspects of disinformation, propaganda, and framing. We believe that such analysis should help in better understanding the trends over time and across languages.

Many interesting directions could be pursued in future work. For example, the analysis could be applied to other languages; in fact, we already did a related study for Bulgarian (Nakov et al., 2021a). Moreover, while here we focused on tweets, the approach is applicable to other social media platforms such as Facebook and WhatsApp.

## Acknowledgments

This research is part of the Tanbih mega-project,<sup>3</sup> developed at the Qatar Computing Research Institute, HBKU, which aims to limit the impact of “fake news”, propaganda, and media bias by making users aware of what they are reading, thus promoting media literacy and critical thinking.

<sup>3</sup><http://tanbih.qcri.org>




Figure 7: Analysis: propaganda techniques.




Figure 8: Analysis: framing.



## References

- Firoj Alam, Fahim Dalvi, Shaden Shaar, Nadir Durrani, Hamdy Mubarak, Alex Nikolov, Giovanni Da San Martino, Ahmed Abdelali, Hassan Sajjad, Kareem Darwish, and Preslav Nakov. 2021a. Fighting the COVID-19 infodemic in social media: A holistic perspective and a call to arms. In *Proceedings of the International AAAI Conference on Web and Social Media, ICWSM '21*, pages 913–922.
- Firoj Alam, Hassan Sajjad, Muhammad Imran, and Ferda Ofli. 2021b. CrisisBench: Benchmarking crisis-related social media datasets for humanitarian information processing. In *Proceedings of the International AAAI Conference on Web and Social Media, ICWSM '21*, pages 923–932.
- Firoj Alam, Shaden Shaar, Fahim Dalvi, Hassan Sajjad, Alex Nikolov, Hamdy Mubarak, Giovanni Da San Martino, Ahmed Abdelali, Nadir Durrani, Kareem Darwish, Abdulaziz Al-Homaid, Wajdi Zaghouni, Tommaso Caselli, Gijs Danoe, Friso Stolk, Britt Bruntink, and Preslav Nakov. 2021c. Fighting the COVID-19 infodemic: Modeling the perspective of journalists, fact-checkers, social media platforms, policy makers, and the society. In *Findings of EMNLP 2021*.
- Phoebe Arnold. 2020. The challenges of online fact checking. Technical report, Full Fact.
- Pepa Atanasova, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, Georgi Karadzhov, Tsvetomila Mihaylova, Mitra Mohtarami, and James Glass. 2019. Automatic fact-checking using context and discourse information. *Journal of Data and Information Quality (JDIQ)*, 11(3):1–27.
- Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. MultiFC: A real-world multi-domain dataset for evidence-based fact checking of claims. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP '19*, pages 4685–4697, Hong Kong, China.
- Ramy Baly, Georgi Karadzhov, Jisun An, Haewoon Kwak, Yoan Dinkov, Ahmed Ali, James Glass, and Preslav Nakov. 2020. What was written vs. who read it: News media profiling using text analysis and social media context. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL '20*, pages 3364–3374.
- Ramy Baly, Georgi Karadzhov, Abdelrhman Saleh, James Glass, and Preslav Nakov. 2019. Multi-task ordinal regression for jointly predicting the trustworthiness and the leading political ideology of news media. In *Proceedings of the 17th Annual Conference of the North American Chapter of the Association for Computational Linguistics, NAACL-HLT '19*, pages 2109–2116, Minneapolis, MN, USA.
- Alberto Barrón-Cedeño, Giovanni Da San Martino, Israa Jaradat, and Preslav Nakov. 2019. Proppy: A system to unmask propaganda in online news. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence, AAAI'19*, pages 9847–9848, Honolulu, HI, USA.
- Alberto Barrón-Cedeño, Tamer Elsayed, Preslav Nakov, Giovanni Da San Martino, Maram Hasanain, Reem Suwaileh, and Fatima Haouari. 2020. Check-That! at CLEF 2020: Enabling the automatic identification and verification of claims in social media. In *Proceedings of the 42nd European Conference on Information Retrieval, ECIR '19*, pages 499–507, Lisbon, Portugal.
- Alberto Barrón-Cedeno, Israa Jaradat, Giovanni Da San Martino, and Preslav Nakov. 2019. Proppy: Organizing the news based on their propagandistic content. *Information Processing & Management*, 56(5):1849–1864.
- Amber E Boydston, Justin H Gross, Philip Resnik, and Noah A Smith. 2013. Identifying media frames and frame dynamics within and across policy issues. In *New Directions in Analyzing Text as Data Workshop*.
- Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. The Media Frames Corpus: Annotations of frames across issues. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing, ACL-IJCNLP '15*, pages 438–444, Beijing, China.
- Loretta H Cheeks, Tracy L Stepien, Dara M Wald, and Ashraf Gaffar. 2020. Discovering news frames: An approach for exploring text, content, and concepts in online news sources. In *Cognitive Analytics: Concepts, Methodologies, Tools, and Applications*, pages 702–721. IGI Global.
- Anton Chernyavskiy, Dmitry Ilvovsky, and Preslav Nakov. 2021. Transformers: “The end of history” for NLP? In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases, ECML-PKDD'21*, Bilbao, Spain (virtual).
- Matteo Cinelli, Walter Quattrociocchi, Alessandro Galeazzi, Carlo Michele Valensise, Emanuele Brugnoli, Ana Lucia Schmidt, Paola Zola, Fabiana Zollo, and Antonio Scala. 2020. The COVID-19 social media infodemic. *Sci Rep.*, 10:16598.
- Giovanni Da San Martino, Alberto Barron-Cedeno, and Preslav Nakov. 2019. Findings of the NLP4IF-2019 shared task on fine-grained propaganda detection. In *Proceedings of the 2nd Workshop on NLP for Internet Freedom (NLP4IF): Censorship, Disinformation, and Propaganda, NLP4IF '19*, pages 162–170, Hong Kong, China.

- Giovanni Da San Martino, Stefano Cresci, Alberto Barrón-Cedeño, Seunghak Yu, Roberto Di Pietro, and Preslav Nakov. 2020a. A survey on computational propaganda detection. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, IJCAI '20, pages 4826–4832.
- Giovanni Da San Martino, Shaden Shaar, Yifan Zhang, Seunghak Yu, Alberto Barrón-Cedeño, and Preslav Nakov. 2020b. Prta: A system to support the analysis of propaganda techniques in the news. In *Proceedings of the Annual Meeting of Association for Computational Linguistics*, ACL '20, pages 287–293.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, Rostislav Petrov, and Preslav Nakov. 2019. Fine-grained analysis of propaganda in news articles. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, EMNLP-IJCNLP '19, pages 5636–5646, Hong Kong, China.
- Leon Derczynski, Kalina Bontcheva, Maria Liakata, Rob Procter, Geraldine Wong Sak Hoi, and Arkaitz Zubiaga. 2017. SemEval-2017 Task 8: RumourEval: Determining rumour veracity and support for rumours. In *Proceedings of the 11th International Workshop on Semantic Evaluation*, SemEval '17, pages 60–67, Vancouver, Canada.
- Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, and Giovanni Da San Martino. 2021a. Detecting propaganda techniques in memes. In *Proceedings of the Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, ACL-IJCNLP '21, pages 6603–6617.
- Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, and Giovanni Da San Martino. 2021b. Task 6 at SemEval-2021: Detection of persuasion techniques in texts and images. In *Proceedings of the 15th International Workshop on Semantic Evaluation*, SemEval '21, pages 70–98.
- Tamer Elsayed, Preslav Nakov, Alberto Barrón-Cedeño, Maram Hasanain, Reem Suwaileh, Pepa Atanasova, and Giovanni Da San Martino. 2019. CheckThat! at CLEF 2019: Automatic identification and verification of claims. In *Proceedings of the 41st European Conference on Information Retrieval*, ECIR '19, pages 309–315, Cologne, Germany.
- Robert M Entman. 1993. Framing: Towards clarification of a fractured paradigm. *McQuail's reader in mass communication theory*, pages 390–397.
- Pepa Gencheva, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, and Ivan Koychev. 2017. A context-aware approach for detecting worth-checking claims in political debates. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*, RANLP '17, pages 267–276, Varna, Bulgaria.
- Genevieve Gorrell, Elena Kochkina, Maria Liakata, Ahmet Aker, Arkaitz Zubiaga, Kalina Bontcheva, and Leon Derczynski. 2019. SemEval-2019 task 7: RumourEval, determining rumour veracity and support for rumours. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, SemEval '19, pages 845–854, Minneapolis, MN, USA.
- Ivan Habernal, Raffael Hannemann, Christian Pollak, Christopher Klamm, Patrick Pauli, and Iryna Gurevych. 2017. Argotario: Computational argumentation meets serious games. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, EMNLP '17, pages 7–12, Copenhagen, Denmark.
- Ivan Habernal, Patrick Pauli, and Iryna Gurevych. 2018. Adapting serious game for fallacious argumentation to German: Pitfalls, insights, and best practices. In *Proceedings of the 11th International Conference on Language Resources and Evaluation*, LREC '18, pages 3329–3335, Miyazaki, Japan.
- Momchil Hardalov, Ivan Koychev, and Preslav Nakov. 2016. In search of credible news. In *Proceedings of the 17th International Conference on Artificial Intelligence: Methodology, Systems, and Applications*, AIMSA '16, pages 172–180, Varna, Bulgaria.
- Maram Hasanain, Fatima Haouari, Reem Suwaileh, Zien Sheikh Ali, Bayan Hamdan, Tamer Elsayed, Alberto Barrón-Cedeño, Giovanni Da San Martino, and Preslav Nakov. 2020. Overview of CheckThat! 2020 Arabic: Automatic identification and verification of claims in social media. In *Working Notes of Conference of the Cross-Language Evaluation Forum for European Languages*, CLEF '2020.
- Naemul Hassan, Chengkai Li, and Mark Tremayne. 2015. Detecting check-worthy factual claims in presidential debates. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*, CIKM '15, pages 1835–1838, Melbourne, Australia.
- Tamanna Hossain, Robert L. Logan IV, Arjuna Ugarte, Yoshitomo Matsubara, Sean Young, and Sameer Singh. 2020. COVIDLies: Detecting COVID-19 misinformation on social media. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP*.
- Israa Jaradat, Pepa Gencheva, Alberto Barrón-Cedeño, Lluís Màrquez, and Preslav Nakov. 2018. Claim-Rank: Detecting check-worthy claims in Arabic and English. In *Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics*, NAACL-HLT '18, pages 26–30, New Orleans, LA, USA.

- Georgi Karadzhov, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, and Ivan Koychev. 2017. Fully automated fact checking using external sources. In *Proceedings of the Conference on Recent Advances in Natural Language Processing, RANLP '17*, pages 344–353, Varna, Bulgaria.
- Haewoon Kwak, Jisun An, and Yong-Yeol Ahn. 2020. A systematic media frame analysis of 1.5 million New York Times articles from 2000 to 2017. In *Proceedings of the 12th ACM Conference on Web Science, WebSci '20*, pages 305–314, Southampton, United Kingdom.
- David M.J. Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. The science of fake news. *Science*, 359(6380):1094–1096.
- Lifang Li, Qingpeng Zhang, Xiao Wang, Jun Zhang, Tao Wang, Tian-Lu Gao, Wei Duan, Kelvin Kam-fai Tsoi, and Fei-Yue Wang. 2020. Characterizing the propagation of situational information in social media during COVID-19 epidemic: A case study on Weibo. *IEEE Transactions on Computational Social Systems*, 7(2):556–562.
- Yaliang Li, Jing Gao, Chuishi Meng, Qi Li, Lu Su, Bo Zhao, Wei Fan, and Jiawei Han. 2016. A survey on truth discovery. *SIGKDD Explor. Newsl.*, 17(2):1–16.
- Siyi Liu, Lei Guo, Kate Mays, Margrit Betke, and Derry Tanti Wijaya. 2019. Detecting frames in news headlines and its application to analyzing news framing trends surrounding US gun violence. In *Proceedings of the 23rd Conference on Computational Natural Language Learning, CoNLL '19*, pages 504–514, Hong Kong, China.
- Richard J Medford, Sameh N Saleh, Andrew Sumarsono, Trish M Perl, and Christoph U Lehmann. 2020. An “Infodemic”: Leveraging high-volume Twitter data to understand early public sentiment for the coronavirus disease 2019 outbreak. *OFID*, 7(7).
- Tsvetomila Mihaylova, Georgi Karadzhov, Pepa Atanasova, Ramy Baly, Mitra Mohtarami, and Preslav Nakov. 2019. SemEval-2019 task 8: Fact checking in community question answering forums. In *Proceedings of the 13th International Workshop on Semantic Evaluation, SemEval '19*, pages 860–869, Minneapolis, MN, USA.
- Tsvetomila Mihaylova, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, Mitra Mohtarami, Georgi Karadjov, and James Glass. 2018. Fact checking in community forums. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, AAAI '18*, pages 879–886, New Orleans, LA, USA.
- Preslav Nakov, Firoj Alam, Shaden Shaar, Giovanni Da San Martino, and Yifan Zhang. 2021a. COVID-19 in Bulgarian social media: Factuality, harmfulness, propaganda, and framing. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing, RANLP '21*.
- Preslav Nakov, Alberto Barrón-Cedeño, Tamer Elsayed, Reem Suwaileh, Lluís Màrquez, Wajdi Zaghoulani, Pepa Atanasova, Spas Kyuchukov, and Giovanni Da San Martino. 2018. Overview of the CLEF-2018 CheckThat! Lab on automatic identification and verification of political claims. In *Proceedings of the Ninth International Conference of the CLEF Association: Experimental IR Meets Multilinguality, Multimodality, and Interaction*, Lecture Notes in Computer Science, pages 372–387, Avignon, France.
- Preslav Nakov, David Corney, Maram Hasanain, Firoj Alam, Tamer Elsayed, Alberto Barrón-Cedeño, Paolo Papotti, Shaden Shaar, and Giovanni Da San Martino. 2021b. Automated fact-checking for assisting human fact-checkers. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence, IJCAI '21*, pages 4551–4558.
- Preslav Nakov, Giovanni Da San Martino, Tamer Elsayed, Alberto Barrón-Cedeño, Rubén Míguez, Shaden Shaar, Firoj Alam, Fatima Haouari, Maram Hasanain, Nikolay Babulkov, Alex Nikolov, Gautam Kishore Shahi, Julia Maria Struß, and Thomas Mandl. 2021c. The CLEF-2021 CheckThat! Lab on detecting check-worthy claims, previously fact-checked claims, and fake news. In *Proceedings of the 43rd European Conference on Information Retrieval, ECIR '21*, pages 639–649.
- Preslav Nakov, Giovanni Da San Martino, Tamer Elsayed, Alberto Barrón-Cedeño, Rubén Míguez, Shaden Shaar, Firoj Alam, Fatima Haouari, Maram Hasanain, Nikolay Babulkov, Alex Nikolov, Gautam Kishore Shahi, Julia Maria Struß, and Thomas Mandl. 2021d. The CLEF-2021 CheckThat! lab on detecting check-worthy claims, previously fact-checked claims, and fake news. In *Proceedings of the 43rd European Conference on Information Retrieval, ECIR '21*, pages 639–649.
- Van-Hoang Nguyen, Kazunari Sugiyama, Preslav Nakov, and Min-Yen Kan. 2020. FANG: Leveraging social context for fake news detection using graph representation. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management, CIKM '20*, pages 1165–1174.
- Ayush Patwari, Dan Goldwasser, and Saurabh Bagchi. 2017. TATHYA: a multi-classifier system for detecting check-worthy statements in political debates. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17*, pages 2259–2262, Singapore.
- Andrew Perrin. 2015. Social media usage. *Pew research center*, pages 52–68.

- Cristina M Pulido, Beatriz Villarejo-Carballido, Gisela Redondo-Sama, and Aitor Gómez. 2020. COVID-19 infodemic: More retweets for science-based information on coronavirus than for false information. *International Sociology*, 35(4):377–392.
- Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, EMNLP '17, pages 2931–2937, Copenhagen, Denmark.
- Shaden Shaar, Firoj Alam, Giovanni Da San Martino, and Preslav Nakov. 2021a. The role of context in detecting previously fact-checked claims. *Arxiv/2104.07423*.
- Shaden Shaar, Firoj Alam, Giovanni Da San Martino, Alex Nikolov, Wajdi Zaghouani, Preslav Nakov, and Anna Feldman. 2021b. Findings of the NLP4IF-2021 shared tasks on fighting the COVID-19 infodemic and censorship detection. In *Proceedings of the Fourth Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda*, pages 82–92.
- Shaden Shaar, Nikolay Babulkov, Giovanni Da San Martino, and Preslav Nakov. 2020a. That is a known lie: Detecting previously fact-checked claims. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, ACL '20, pages 3607–3618.
- Shaden Shaar, Maram Hasanain, Bayan Hamdan, Zien Sheikh Ali, Fatima Haouari, Alex Nikolov, Mucahid Kutlu, Yavuz Selim Kartal, Firoj Alam, Giovanni Da San Martino, Alberto Barrón-Cedeño, Rubén Míguez, Javier Beltrán, Tamer Elsayed, and Preslav Nakov. 2021c. Overview of the CLEF-2021 CheckThat! lab task 1 on check-worthiness estimation in tweets and political debates. In *Working Notes of CLEF 2021—Conference and Labs of the Evaluation Forum*, CLEF '2021. CEUR-WS.org.
- Shaden Shaar, Alex Nikolov, Nikolay Babulkov, Firoj Alam, Alberto Barrón-Cedeño, Tamer Elsayed, Maram Hasanain, Reem Suwaileh, Fatima Haouari, Giovanni Da San Martino, and Preslav Nakov. 2020b. Overview of CheckThat! 2020 English: Automatic identification and verification of claims in social media. In *Working Notes of CLEF 2021—Conference and Labs of the Evaluation Forum*, CEUR Workshop Proceedings. CEUR-WS.org.
- Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake news detection on social media: A data mining perspective. *SIGKDD Explor. Newsl.*, 19(1):22–36.
- Slavena Vasileva, Pepa Atanasova, Lluís Màrquez, Alberto Barrón-Cedeño, and Preslav Nakov. 2019. It takes nine to smell a rat: Neural multi-task learning for check-worthiness prediction. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*, RANLP '19, pages 1229–1239, Varna, Bulgaria.
- Bertie Vidgen, Scott Hale, Ella Guest, Helen Margetts, David Broniatowski, Zeerak Waseem, Austin Botelho, Matthew Hall, and Rebekah Tromble. 2020. Detecting East Asian prejudice on social media. In *Proceedings of the Workshop on Online Abuse and Harms*, WOAHA '20, pages 162–172.
- Nguyen Vo and Kyumin Lee. 2018. The rise of guardians: Fact-checking URL recommendation to combat fake news. In *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, SIGIR '18, pages 275–284, Ann Arbor, MI, USA.
- Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science*, 359(6380):1146–1151.
- William Yang Wang. 2017. “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*, ACL '17, pages 422–426, Vancouver, Canada.
- Samuel C Woolley and Philip N Howard. 2018. *Computational propaganda: political parties, politicians, and political manipulation on social media*. Oxford University Press.
- Seunghak Yu, Giovanni Da San Martino, Mitra Moughtarami, James Glass, and Preslav Nakov. 2021. Interpretable propaganda detection in news articles. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*, RANLP '21.
- Yifan Zhang, Giovanni Da San Martino, Alberto Barrón-Cedeño, Salvatore Romeo, Jisun An, Hae-woon Kwak, Todor Staykovski, Israa Jaradat, Georgi Karadzhov, Ramy Baly, Kareem Darwish, James Glass, and Preslav Nakov. 2019. Tanbih: Get to know what you are reading. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing: System Demonstrations*, EMNLP-IJCNLP '19, pages 223–228, Hong Kong, China.
- Xinyi Zhou, Apurva Mulay, Emilio Ferrara, and Reza Zafarani. 2020. ReCOVeRY: A multimodal repository for COVID-19 news credibility research. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*, CIKM '20, pages 3205–3212.
- Dimitrina Zlatkova, Preslav Nakov, and Ivan Koychev. 2019. Fact-checking meets fauxtography: Verifying claims about images. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, EMNLP '19, pages 2099–2108, Hong Kong, China.