

# ArCOV19-Rumors: Arabic COVID-19 Twitter Dataset for Misinformation Detection

Fatima Haouari, Maram Hasanain, Reem Suwaileh, Tamer Elsayed

Computer Science and Engineering Department, Qatar University

{200159617, maram.hasanain, rs081123, telsayed}@qu.edu.qa

## Abstract

In this paper we introduce ArCOV19-Rumors, an Arabic COVID-19 Twitter dataset for misinformation detection composed of tweets containing claims from 27<sup>th</sup> January till the end of April 2020. We collected 138 verified claims, mostly from popular fact-checking websites, and identified 9.4K relevant tweets to those claims. Tweets were manually-annotated by veracity to support research on misinformation detection, which is one of the major problems faced during a pandemic. ArCOV19-Rumors supports two levels of misinformation detection over Twitter: verifying free-text claims (called *claim-level verification*) and verifying claims expressed in tweets (called *tweet-level verification*). Our dataset covers, in addition to health, claims related to other topical categories that were influenced by COVID-19, namely, social, politics, sports, entertainment, and religious. Moreover, we present benchmarking results for tweet-level verification on the dataset. We experimented with SOTA models of versatile approaches that either exploit content, user profiles features, temporal features and propagation structure of the conversational threads for tweet verification.

## 1 Introduction

In addition to being a medium for the spread and consumption of news, Twitter has been shown to capture the dynamics of real-world events including the spread of diseases such as the seasonal influenza (Kagashe et al., 2017) or more severe epidemics like Ebola (Roy et al., 2020). Since the first reported case of the Novel Coronavirus (COVID-19) in China, in November 2019, the COVID-19 topic has drawn the interest of many Arab users over Twitter. Their interest, reflected in the Arabic content on the platform, has reached a peak after two months when the first case was reported in the United Arab Emirates late in Jan-

uary 2020. This ongoing pandemic has, unsurprisingly, spiked discussions on Twitter covering a wide range of topics such as general information about the disease, preventive measures, procedures and newly-enforced decisions by governments, up-to-date statistics of the spread in the world, and even the change in our daily habits and work styles. With the great importance and spread of COVID-19 information, misinformation and fake news have infected the Twitter stream. An early study quantifying COVID-19 medical misinformation on Twitter found that 25% of collected tweets contained misinformation (Kouzy et al., 2020). During COVID-19 pandemic, we observed that misinformation stretched beyond spreading fake and potentially-harmful medical information, to information that can have adverse negative political effects (example: “*In light of the unresponsiveness of Yemeni government to requests of evacuation from Yemeni students in Wuhan, Sultan of Oman orders their evacuation.*”) and economical effects too (example: “*Kuwaitis boycott AlMarai Saudi dairy company after reports on Coronavirus infected employees.*”). Combating the spread of such claims and verifying them becomes essential during this sensitive time.

In this work, we aim to facilitate research on misinformation detection on social media during this complex and historical period of our time by introducing a manually-annotated Arabic dataset, ArCOV19-Rumors, that covers tweets spreading COVID-19 related claims. To construct ArCOV19-Rumors, we start from an existing COVID-19 Arabic dataset (Haouari et al., 2021), ArCOV-19, that is the first Arabic COVID-19 Twitter dataset with propagation networks. Our proposed ArCOV19-Rumors includes a set of 138 COVID-19 verified claims and 9.4K corresponding relevant tweets that were manually-annotated to support both *claim-level* and *tweet-level* verification tasks. Claim-level verification is defined as follows: *given a short*

Dataset	# Tweets	# Annotated	Labels	Multi-task	Conversations	Manual Annot.
<a href="#">Elhadad et al. (2020a)</a>	220K	220K	misleading, real	-	-	-
<a href="#">Alsudias and Rayson (2020)</a>	1M	2K	false, true, unrelated	-	-	✓
<a href="#">Mubarak and Hassan (2020)</a>	30M	8K	rumor, info, advice, ...	-	-	✓
<a href="#">Alqurashi et al. (2021)</a>	4.5M	8.8K	misinformation, other	-	-	✓
<b>ArCOV19-Rumors</b>	1M	9.4K	false, true, other	✓	✓	✓

Table 1: Comparison between ArCOV19-Rumors and existing Arabic COVID-19 datasets for verification over tweets.

*free-text claim (in 1 or 2 sentences) and its corresponding relevant tweets, predict whether the claim is true or false.* Tweet-level verification is defined as follows: *given a tweet containing a claim, detect whether it is true or false.*

To our knowledge, ArCOV19-Rumors is the only Arabic dataset made available to support both claim-level and tweet-level verification tasks in Twitter given the propagation networks of the tweets in general and on COVID-19 in particular. Some related Twitter datasets were recently released by [Alqurashi et al. \(2021\)](#), [Mubarak and Hassan \(2020\)](#), [Alsudias and Rayson \(2020\)](#), and [Elhadad et al. \(2020a\)](#). None of these datasets has the propagation networks of the tweets, and they either support the tweet verification task ([Elhadad et al., 2020a](#); [Alsudias and Rayson, 2020](#)), misinformation detection ([Alqurashi et al., 2021](#)) or multi-class categorization including rumors as one category ([Mubarak and Hassan, 2020](#)). Differently from ArCOV19-Rumors, [Elhadad et al. \(2020a\)](#) used an automatic approach to annotate tweets, while [Alsudias and Rayson \(2020\)](#) only cover COVID-19 health-oriented claims. Starting from some health misinformation reported by the Ministry of Health in Saudi Arabia and the World Health Organization (WHO), [Alqurashi et al. \(2021\)](#) annotated COVID-19 tweets as misinformation or not. Differently, our dataset covers in addition to health, other types of claims that were influenced by COVID-19, namely, social, politics, sports, entertainment, and religious. Table 1 demonstrates the differences between ArCOV19-Rumors and the aforementioned datasets.

The contribution of this paper is three-fold:

- We construct and release<sup>1</sup> the first Arabic dataset for misinformation detection over Twitter, covering both claim and tweet verification tasks. It contains 138 COVID-19 verified claims that scale to 9.4K labeled relevant

<sup>1</sup><https://gitlab.com/bigirqu/ArCOV-19/-/tree/master/ArCOV19-Rumors>

tweets along with their propagation networks.

- We suggest and motivate several research tasks that can be addressed using our *labeled* dataset for misinformation detection.
- We present benchmark results on tweet-level verification using SOTA models that either exploit content, user profiles, or the temporal features and the propagation structure of the conversational threads. Results offer baselines for future research on the problem.

The remainder of the paper is organized as follows. We present studies related to COVID-19 misinformation analysis and datasets in Section 2. The construction of ArCOV19-Rumors is presented in Section 3. Several use cases supported by our dataset are discussed in Section 4. Our benchmarks and their performances are presented in Section 5. We detail the released components of the dataset in Section 6, and conclude in Section 7.

## 2 Related Work

The negative effects and spread of misinformation about COVID-19, triggered efforts to understand this phenomenon. Many studies analyzed misinformation spreading on Twitter related to COVID-19 ([Shahi et al., 2020](#); [Cinelli et al., 2020](#); [Gallotti et al., 2020](#); [Kouzy et al., 2020](#); [Singh et al., 2020](#)). [Singh et al. \(2020\)](#) analyzed the spread of five health misinformation over time. [Shahi et al. \(2020\)](#) analyzed false and partially false tweets that have been previously fact-checked by a fact-checking platform. [Kouzy et al. \(2020\)](#) focused only on health misinformation and analyzed tweets negating health information by trusted sources (e.g., WHO). [Cui and Lee \(2020\)](#) and [Hossain et al. \(2020\)](#) published an English dataset for claim verification, and tweet verification respectively, while [Dharawat et al. \(2020\)](#) released a Twitter dataset for health risk assessment of COVID-19-related English tweets.

In this work, we extend an existing raw *Arabic* COVID-19 Twitter dataset, ArCOV-19 (Haouari et al., 2021), to include manually-annotated tweets to support both claim and tweet verification. To our knowledge, there is no work that released a large manually-annotated Arabic dataset for misinformation detection on COVID-19 at both claim and tweet levels. The closest work to ours is that done by Elhadad et al. (2020a) and Alsudias and Rayson (2020). However, neither have the propagation networks of the annotated tweets, and they both solely support the tweet verification task. Moreover, ArCOV19-*Rumors* has the largest number of *manually* annotated tweets.

Claim verification is a widely studied problem in general. Several non-COVID-19 datasets exist for claim verification (e.g., (Gorrell et al., 2019)), however, they either support English language (Zubiaga et al., 2016; Derczynski et al., 2017; Gorrell et al., 2019; Ma et al., 2017) or Chinese language (Ma et al., 2017). Among these studies, some only focused on tweet-level verification (Ma et al., 2017; Liu and Wu, 2018; Bian et al., 2020; Khoo et al., 2020) while others addressed claim-level verification (Vosoughi et al., 2017; Dang et al., 2019).

There are few initiatives targeting Arabic online content, but they support rumors detection (i.e., rumor or non-rumor) (Alzanin and Azmi, 2019; Alkhair et al., 2019), tweet verification without propagation networks (Elhadad et al., 2020a; Alsudias and Rayson, 2020), misinformation detection (Alqurashi et al., 2021), or tweet credibility (El Ballouli et al., 2017). Another notable work is a shared task by CLEF CheckThat! Lab for claim verification (Elsayed et al., 2019; Barrón-Cedeño et al., 2020; Barrón-Cedeno et al., 2020). To the best of our knowledge, this is the first Arabic dataset that supports both claim and tweet verification on social media using the propagation networks.

### 3 Methodology

We extended ArCOV-19, an Arabic Twitter dataset about COVID-19 (Haouari et al., 2021), by annotating a subset of the tweets to support research on misinformation detection, which is one of the major problems faced during a pandemic. We aim to support two classes of the misinformation detection problem (with variants) over Twitter: verifying free-text claims (called *claim-level verification*) and verifying claims expressed in tweets (called *tweet-level verification*), covering two com-

mon use cases. To that end, we need to collect COVID-19 *verified* claims, then, for each claim, we need to find, in ArCOV-19, corresponding *relevant* tweets, and finally, among the relevant tweets, identify those that are either expressing the claim or negating it (to easily *propagate* the veracity of the claims to the tweets). This section details how that pipeline was implemented (Sections 3.1-3.6).

#### 3.1 Collecting COVID-19 Verified Claims

We manually collected a set of verified COVID-related claims from two popular Arabic fact-checking platforms: Fatabyyano<sup>2</sup> and Misbar.<sup>3</sup> We identified 113 false claims and 18 true claims during the period of ArCOV19-*Rumors*. To improve balance between false and true claims, we collect an additional 31 true claims from these sources:

1. **Authoritative Health Organizations:** We retrieved, from ArCOV-19, all tweets from Twitter accounts of the World Health Organization (WHO), ministries of public health in Arab countries, UNESCO, and UNICEF. Then, we manually extracted a set of true claims from the retrieved tweets.
2. **English Fact-checking Platforms:** We collected (and translated to Arabic) a few true claims from English Fact-checking platforms: PolitiFact,<sup>4</sup> Snopes,<sup>5</sup> and FullFact.<sup>6</sup>

That yielded a total of 162 claims, 113 false and 49 true, that are potentially suitable for our purpose.

#### 3.2 Finding Relevant Tweets

For each claim, we manually constructed and refined a set of search queries in order to retrieve potentially-relevant tweets from ArCOV-19 using Boolean search. To help identify the queries and maximize our coverage, we considered keywords that appeared in the social media posts, expressing the claim, provided as examples by the fact-checking platforms, and we also interactively search in Twitter to find the best possible keywords that can retrieve relevant tweets to the claim. We then used the queries to search ArCOV-19. An example is the true claim “Trump suggests injecting disinfectants as a treatment for COVID-19”.<sup>7</sup>

<sup>2</sup><https://fatabyyano.net>

<sup>3</sup><https://misbar.com/>

<sup>4</sup><https://www.politifact.com/>

<sup>5</sup><https://www.snopes.com/>

<sup>6</sup><https://fullfact.org/>

<sup>7</sup><https://tinyurl.com/y4btakuw>

We noticed that the claim and also the tweet given as an example by Misbar fact-checking platform are using two different Arabic words that reflect “disinfectants”, namely *mutaharat* (*disinfectants* in English) and *mueaqqamat* (*sterilizers* in English), so we used both in combination with either *Trump* or *US president* as search queries.

For each claim, we manually filtered the retrieved tweets to discard the non-relevant ones. This is conducted by one author of this paper. Examples of tweets that were considered relevant are tweets expressing or negating the claim, tweets having multiple claims including the target claim, tweets expressing advice, questions, or personal opinions about the claim, or even sarcastic tweets about it. Overall, more than half of the retrieved tweets were non-relevant, yielding 14,472 tweets relevant to the collected claims.

### 3.3 Filtering Claims

After identifying the relevant tweets, we excluded some claims for different reasons. We excluded claims with less than two relevant tweets due to extremely-insufficient content. We also excluded the ones (especially health-related claims) for which the veracity changed within the period of the dataset (e.g., “*Pets do not catch COVID-19*”), or those for which there are still no clear evidence with or against the claim (e.g., “*Bats are the source of COVID-19*”). Moreover, we discarded the claims that are not very specific, i.e., too general (e.g., “*Many Arab countries took security measures against those who refuse to quarantine*”). We eventually kept 95 false and 43 true claims, a total of 138 claims. Those claims have a total of 9,414 relevant tweets. The final set of claims is diverse and claims fall under different categories, as the misinformation propagating in Twitter during COVID-19 pandemic was not restricted to health. In fact, only 45 of them were health-related. The rest are distributed over social (38), political (22), religious (18), entertainment (9), and sports (6) topical categories.

### 3.4 Annotating Relevant Tweets

After collecting relevant tweets for each claim, it was time to identify tweets expressing or negating the claim, so that we can *propagate* the label of the claim to them. For each claim, one of the authors of this paper labelled all of its relevant tweets by

stance using the following three categories:<sup>8</sup>

- “Expressing same claim” if the main (focused) claim in the tweet is restating, expressing, or rephrasing the target claim. This tweet then receives the *same* veracity of the target claim; thus inheriting its label (whether *true* or *false*).
- “Negating the claim” if the main (focused) claim in the tweet is negating or denying the target claim. The veracity of this tweet is then the *opposite* of the veracity of the target claim, i.e., it is labeled as *true* if the target claim is *false*, and vice-versa.
- “Other” if the tweet cannot be labelled as one of the two earlier cases, e.g., expressing opinion or giving advice regarding the claim.

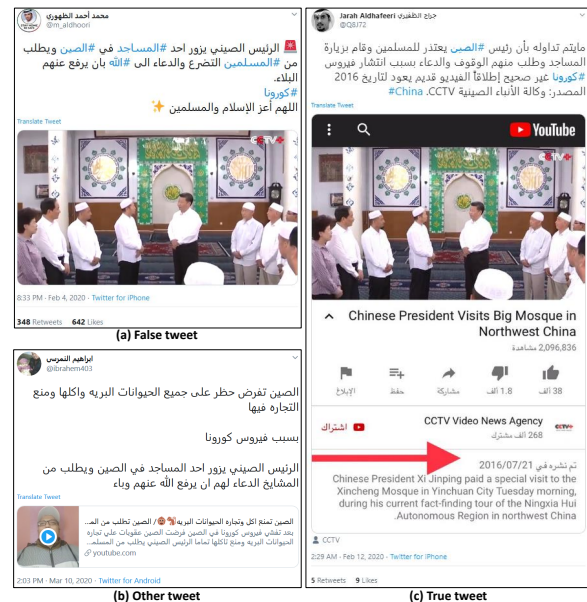


Figure 1: Example relevant tweets to a false claim: “*The Chinese President visits one of the mosques in China and asks Muslims to pray to God to alleviate the pandemic.*”

Figure 1 illustrates three example relevant tweets of a *false* claim translated as “*The Chinese President visits one of the mosques in China and asks Muslims to pray to God to alleviate the pandemic*”, which is literally stated in the tweet shown in Figure 1(a), thus it is *false*. On the contrary, the tweet illustrated in Figure 1(c) is denying the claim, and thus is *true*. It is translated as “*What is being circulated about the Chinese President, apologizing to*

<sup>8</sup>Full annotation guidelines: [https://gitlab.com/bigirqu/ArCOV-19/-/blob/master/ArCOV19-Rumors/annotations\\_guidelines.md](https://gitlab.com/bigirqu/ArCOV-19/-/blob/master/ArCOV19-Rumors/annotations_guidelines.md)



Muslims and visits the mosques and asking them to pray because of the spread of the #Corona virus, is totally incorrect. The video is old and dated back to 2016. Source: the Chinese news agency CCTV". Finally, the tweet illustrated in Figure 1(b) is labelled as "other" due to expressing multiple claims. It is translated as "China bans eating all wild animals and their trade due to Corona Virus. The Chinese President visits a mosque in China and asks the scholars to pray to alleviate the pandemic".

It is worth mentioning that expressing or negating the claim can be over an external link in the tweet, stated in an image, or said in a video, rather than in the text of the tweet. This was considered while annotating tweets. Providing such tweets in ArCOV19-Rumors allows the development of multi-modal systems, that can use signals in text, images, or videos, to make verification decisions.

### 3.5 Data Quality

As mentioned earlier, our dataset was annotated by a single annotator to reduce annotation time and because the annotator was well-experienced with the task. To measure annotation quality, we randomly selected 10% of the relevant tweets, and asked a second annotator to label them. We found that the agreement ratio between annotators is 0.87 and 0.80 for relevance and stance respectively. Due to subjectivity of the stance labelling task specifically (and since the annotator was also asked to consider images, videos, links, etc. in annotation which might lead to further subjectivity), we believe the agreement level is acceptable.

### 3.6 Collecting Propagation Networks

We also collected the propagation networks (i.e., retweets and conversational threads) for each relevant tweet. The propagation networks for tweets that contain misinformation are essential to study its spreading behaviour and can constitute evidential signals for verification. Figure 2 shows some replies to the false tweet presented in Figure 1(a). We notice that some replies have a clear stance against the claim. Moreover, one reply presents evidence that the claim in the original tweet is false. In addition to the replies, exploiting profiles of propagators can play a significant role in verifying the tweet (Liu and Wu, 2018).



Figure 2: Evidence from replies (with translation) against the false tweet in Fig. 1(a).

Table 2: Statistics of Claims in ArCOV19-Rumors (RPs=Replies, RTs=Retweets).

	Claims Subset	Tweet Vf. Subset
<b>Total tweets</b>	9,414	3,584
False Tweets	1,753 (18.6%)	1,753 (48.9%)
True Tweets	1,831 (19.4%)	1,831 (51.1%)
Other Tweets	5,830 (61.9%)	0
Average tweets / claim	68	26
Tweets with RPs	3,161 (33.6%)	1,222 (34.1%)
Tweets with RTs	3,810 (40.5%)	1,629 (45.5%)
Tweets with RPs & RTs	2,006 (21.3%)	772 (21.5%)
Average RPs / tweet	33	31
Average RTs per / tweet	16	19

## 4 Immediate Use Cases

Table 2 presents a statistical summary of the labeled tweets in ArCOV19-Rumors. We define two subsets of tweets. The first is denoted as **Claims Subset**, which includes all relevant tweets of the claims (labeled as *true*, *false*, or *other*). This is of high interest since all tweets relevant to a specific claim can be used in verifying it. The other is denoted as **Tweet Verification Subset**, which only includes the relevant tweets that are either *expressing* or *denying* the original claims, excluding the ones labeled as *other*. This is also of high interest since each of those tweets is subject to verification.

The table indicates that, out of 9.4k labelled relevant tweets, about 3.6k of them are either true or false (constituting the tweet verification subset); each is considered a separate tweet-level verification query. We notice that the distribution of *true* vs. *false* tweets is balanced, making it a good resource for training verification systems. We also notice that a good portion of both subsets have retweets and replies, indicating potentially-useful propagation networks. Accordingly and based on the two subsets, our labeled data can support three different misinformation detection tasks.

## 4.1 Claim-level Verification

This task is defined as follows: *given a claim and all corresponding relevant tweets (with their propagation networks), detect the veracity of the claim, i.e., whether the claim is true or false.* There are some initiatives to support Arabic claim verification. However, in one of the most prominent ones (CheckThat! lab at CLEF-2019), a pre-defined set of Web page was used in verification (Elsayed et al., 2019) while in ArCOV19-Rumors we focus on Twitter.

## 4.2 Tweet-level Verification

This task is defined as follows: *given a tweet (with its propagation networks), detect its veracity, i.e., whether the tweet is true or false.* Addressing this task in Arabic has never been studied. Existing studies for Arabic tweet verification mainly rely on the source tweet content only (Elhadad et al., 2020b; Alsudias and Rayson, 2020), or, additionally, the potentially-relevant Web pages (Barrón-Cedeño et al., 2020; Barrón-Cedeño et al., 2020).

ArCOV19-Rumors supports another variant of this task. This variant makes also available (to the verification system) the tweets that are relevant to its target claim but were posted *earlier* (with their propagation networks) allowing for early claim verification. To our knowledge, there is no study that addressed this problem. This is an interesting problem for several reasons. First, with the lack of any propagation networks for a target tweet, a verification system can still exploit networks of earlier relevant tweets to verify the tweet. Second, as time is critical to debunk fake claims, exploiting relevant tweets posted earlier allows verifying the tweet as soon as it is posted, without waiting for its propagation networks. Third, even if the target tweet has propagation networks, relevant tweets might provide more evidence, hence improving verification accuracy.

## 4.3 Claim Retrieval

This task is defined as follows: *given a tweet that expresses a claim (i.e., a tweet in the tweet verification subset), find all tweets that are expressing the same claim.* To our knowledge, this task is under-studied; an exception is the work done by Shaar et al. (2020) and the task proposed by Barrón-Cedeño et al. (2020) in CLEF CheckThat! 2020

lab;<sup>9</sup> however, they focus more on claims than tweets and they release English-only datasets to support this task. Solving this problem helps in applications like finding previously-verified tweets, or clustering tweets expressing the same claim, to avoid re-verification.

## 5 Experiments and Evaluation

In this section, we present *benchmarking results* using SOTA models on the *tweet-level* verification task to facilitate future research. We experimented with variant models that either exploit content only, user profile features, temporal features, or propagation structure of the conversational threads for tweet verification. We present our preprocessing approach in Section 5.1, the SOTA models used for benchmarking in Section 5.2, and finally a discussion of the results in Section 5.3.

### 5.1 Preprocessing

In all experiments, we used our **tweets with RPs** subset, since some of the SOTA methods we experiment with depend on replies of tweets to be verified (*target tweets* hereafter). In ArCOV19-Rumors, some of the target tweets were replies to previous tweets. In our experiments, we considered the direct and indirect replies that were posted *only after* the target tweet. We also eliminated target tweets that have no textual replies, i.e., containing only emojis, non Arabic content, images, or videos. We ended up with 1,108 tweets, 597 True and 511 False, with 16 replies on average. We split our data by target tweets over 5 folds, ensuring there is no overlap of claims across folds (i.e., no overlap between tweets of the same claim).<sup>10</sup>

We processed the tweets and the replies by removing non-Arabic letters, URLs, handles, special characters, and diacritics.

### 5.2 Verification Models

We experimented with two SOTA models that exploit the propagation networks for tweet verification. Moreover, due to proven effectiveness of BERT-based classifiers in versatile text classification tasks, we elect to develop a simple and effective BERT-based classifier that we fine-tune for the

<sup>9</sup><https://sites.google.com/view/clef2020-checkthat/>

<sup>10</sup>Folds used in our experiments: [https://gitlab.com/bigirqu/ArCOV-19/-/tree/master/ArCOV19-Rumors/baselines\\_folds](https://gitlab.com/bigirqu/ArCOV-19/-/tree/master/ArCOV19-Rumors/baselines_folds)

task. In this BERT-based classification architecture, we experimented with two Arabic pre-trained BERT models, namely, AraBERT and MARBERT. Further details on the models are presented below.

1. Bi-GCN [Bian et al. \(2020\)](#): a bidirectional Graph Convolutional Networks model that leverages Graph Convolutional Networks to verify tweets given the target tweet and replies content, in addition to the replies tree structure. The model was called bidirectional for its ability to capture the bottom-up and top-down replies tree for each tweet. We used the authors implementation and setup<sup>11</sup> in our experiments. Each tweet/reply was represented using a vector of 5K content features, which are the most frequent 5K words in the dataset excluding stop words.
2. PPC-RNN+CNN [Liu and Wu \(2018\)](#): a multivariate time series model that exploits user profiles to verify tweets. Each time step involves a single user represented by a vector of user features. A user can be the author of the target tweet or a reply to it. The user vectors were sorted by posting time starting from the target tweet to capture the temporal features. In the original paper, the model uses a fixed number of time steps. In our experiments, we find the average number of replies per target tweet is 16, thus, we set the time steps to 17 taking into account the target tweet. We implemented this model using Keras<sup>12</sup> following the same setup presented in the original paper.
3. AraBERT [Baly et al. \(2020\)](#): a pretrained BERT model on a large-scale Arabic corpus of news articles. We fine-tuned the model to classify the tweets as True or False *given the target tweet content only*. We used the fine-tuning implementation with the same hyper-parameters setup shared by the model developers,<sup>13</sup> using AraBERTv1, which is based on pre-segmentation using Farasa segmenter,<sup>14</sup> and we set the max sequence length to 128. The model was fine-tuned with a batch size of 16 for 8 epochs, with a learning rate of  $10^{-5}$ .

<sup>11</sup><https://github.com/TianBian95/BiGCN>

<sup>12</sup><https://keras.io/>

<sup>13</sup><https://colab.research.google.com/drive/1P9iQHtUH5KUBTVtp8B4-AopZzEEPE0lw?usp=sharing>

<sup>14</sup><https://farasa.qcri.org/>

Table 3: Tweet-level verification results.

Model	Class	Acc	P	R	F1	macro-F1
Majority	-	0.533	0.266	0.500	-	0.343
Bi-GCN	F	0.669	0.731	0.526	0.580	0.649
	T		0.662	0.830	0.718	
RNN+CNN	F	0.682	0.631	0.783	0.688	0.673
	T		0.749	0.600	0.657	
AraBERT	F	0.730	0.733	0.671	0.691	0.713
	T		0.719	0.763	0.735	
MARBERT	F	0.757	0.763	0.692	0.717	0.740
	T		0.743	0.788	0.762	

4. MARBERT [Abdul-Mageed et al. \(2020\)](#): a pretrained BERT model on a large Twitter dataset. We selected this model since in our task, we are solely working with Twitter data and thus, a pretrained model from the same domain might be more suitable. We fine-tuned the model for a classification task given the same input as AraBERT. We used the authors' code<sup>15</sup> and setup to fine-tune the model. The model was fine-tuned with a batch size of 32 for 5 epochs, with a learning rate of  $2 * 10^{-6}$ . Similar to AraBERT, the max sequence length was set to 128.

### 5.3 Results and Discussion

Table 3 presents the performance of the four models, in addition to a simple majority baseline. For each model, we report the overall accuracy, and also F1, precision, and recall for each class. Starting from the same set of 5 folds, we trained each model 5 times, each with a random seed, and we report the average over those runs then over folds.<sup>16</sup>

Results show that all models significantly outperform the majority baseline. Moreover, the pretrained BERT models are superior to the other models, despite the fact that they leverage the reply tree in addition to the target tweet. More specifically, MARBERT exhibits better performance than AraBERT, indicating the effectiveness of matching the pre-training domain with the testing one, and achieving 0.74 macro-averaged F1.

The results also demonstrate that the models are better in detecting True claims than False claims, with an exception of the PPC-RNN+CNN model.

<sup>15</sup><https://colab.research.google.com/drive/1M01s7EPUi1dwqIDh6HNfJ5y826XvcgGX?usp=sharing>

<sup>16</sup>It is worth noting that since we set the hyper-parameters values as reported by original papers, better performance might be achieved if hyper-parameters are tuned on our dataset.

This can be attributed to the fact that the average number of replies for False tweets is 12 vs. 19 for True tweets. We note that, in case the time steps are actually less than 17 (the average indicated earlier), the model *randomly* fills the missing time steps with user features from other repliers to the target tweet. Such limitation in the model might lead to added noise and thus a poor performance with False tweets.

Moreover, the Bi-GCN model is language-dependent, and since we are working with Arabic data, we may need to test preprocessing the data with different techniques, such as considering stemming the content or replacing the top 5000 words features with content embeddings instead.

## 6 Data Release

In summary, we release the following resources as ArCOV19-*Rumors* dataset, taking into consideration Twitter content redistribution policy:<sup>17</sup>

- **Verified Claims:** 138 verified claims, each labeled as true or false.
- **Claims Subset:** IDs of the tweets relevant to the verified claims, each labeled as true, false, or other.
- **Propagation Networks of the Claims Subsets:** which includes for each tweet in the claim subset:
  - **Retweets:** IDs of the full retweet set.
  - **Conversational Threads:** tweet IDs of the full reply thread (including direct and indirect replies).
- **Annotation guidelines:** guidelines used to annotate the relevant tweets.
- **Baselines folds:** tweet IDs for folds used to train our baselines.

Along with the dataset, we plan to provide some pointers to publicly-available crawlers that users can easily use to crawl the tweets given their IDs.

## 7 Conclusion

In this paper, we presented ArCOV19-*Rumors*, which is an Arabic Twitter dataset that supports both claim-level and tweet-level verification tasks given the propagation networks. We released 138 verified claims associated with 9.4K relevant tweets. Our dataset covers, in addition to health, other types of claims that were influenced by

COVID-19, namely, social, politics, sports, entertainment, and religious. To facilitate future research, we presented benchmarking performance results of some SOTA models on the tweet verification task.

## Acknowledgments

The work of Tamer Elsayed and Maram Hasanain was made possible by NPRP grant# NPRP 11S-1204-170060 from the Qatar National Research Fund (a member of Qatar Foundation). The work of Reem Suwaileh was supported by GSRA grant# GSRA5-1-0527-18082 from the Qatar National Research Fund and the work of Fatima Haouari was supported by GSRA grant# GSRA6-1-0611-19074 from the Qatar National Research Fund. The statements made herein are solely the responsibility of the authors.

## References

- Muhammad Abdul-Mageed, AbdelRahim Elmadany, and El Moatez Billah Nagoudi. 2020. Arbert & marbert: Deep bidirectional transformers for arabic. *arXiv preprint arXiv:2101.01785*.
- Maysoon Alkhair, Karima Meftouh, Kamel Smaïli, and Nouha Othman. 2019. An arabic corpus of fake news: Collection, analysis and classification. In *International Conference on Arabic Language Processing*, pages 292–302. Springer.
- Sarah Alqurashi, Btool Hamoui, Abdulaziz Alashaikh, Ahmad Alhindi, and Eisa Alanazi. 2021. Eating garlic prevents covid-19 infection: Detecting misinformation on the arabic content of twitter. *arXiv preprint arXiv:2101.05626*.
- Lama Alsudias and Paul Rayson. 2020. COVID-19 and Arabic Twitter: How can Arab world governments and public health organizations learn from social media? In *Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020*. Association for Computational Linguistics.
- Samah M Alzanin and Aqil M Azmi. 2019. Rumor detection in arabic tweets using semi-supervised and unsupervised expectation–maximization. *Knowledge-Based Systems*, 185:104945.
- Fady Baly, Hazem Hajj, et al. 2020. Arabert: Transformer-based model for arabic language understanding. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 9–15.
- Alberto Barrón-Cedeño, Tamer Elsayed, Preslav Nakov, Giovanni Da San Martino, Maram Hasanain,

<sup>17</sup><https://developer.twitter.com/en/developer-terms/agreement-and-policy>



- Reem Suwaileh, and Fatima Haouari. 2020. Checkthat! at clef 2020: Enabling the automatic identification and verification of claims in social media. In *European Conference on Information Retrieval*, pages 499–507. Springer.
- Alberto Barrón-Cedeno, Tamer Elsayed, Preslav Nakov, Giovanni Da San Martino, Maram Hasanain, Reem Suwaileh, Fatima Haouari, Nikolay Babulakov, Bayan Hamdan, Alex Nikolov, et al. 2020. Overview of checkthat! 2020: Automatic identification and verification of claims in social media. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 215–236. Springer.
- Tian Bian, Xi Xiao, Tingyang Xu, Peilin Zhao, Wenbing Huang, Yu Rong, and Junzhou Huang. 2020. Rumor detection on social media with bi-directional graph convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 549–556.
- Matteo Cinelli, Walter Quattrociochi, Alessandro Galeazzi, Carlo Michele Valensise, Emanuele Brugnoli, Ana Lucia Schmidt, Paola Zola, Fabiana Zollo, and Antonio Scala. 2020. The covid-19 social media infodemic. *arXiv preprint arXiv:2003.05004*.
- Limeng Cui and Dongwon Lee. 2020. Coaid: Covid-19 healthcare misinformation dataset. *arXiv preprint arXiv:2006.00885*.
- Anh Dang, Abidalrahman Moh’d, Aminul Islam, and Evangelos Milios. 2019. Early detection of rumor veracity in social media. In *Proceedings of the 52nd Hawaii International Conference on System Sciences*.
- 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-2017)*. Association for Computational Linguistics.
- Arkin Dharawat, Ismini Lourentzou, Alex Morales, and ChengXiang Zhai. 2020. Drink bleach or do what now? covid-hera: A dataset for risk-informed health decision making in the presence of covid19 misinformation. *arXiv preprint arXiv:2010.08743*.
- Rim El Ballouli, Wassim El-Hajj, Ahmad Ghandour, Shady Elbassuoni, Hazem Hajj, and Khaled Shaban. 2017. Cat: Credibility analysis of arabic content on twitter. In *Proceedings of the Third Arabic Natural Language Processing Workshop*, pages 62–71.
- Mohamed K Elhadad, Kin Fun Li, and Fayez Gebali. 2020a. Covid-19-fakes: a twitter (arabic/english) dataset for detecting misleading information on covid-19. In *International Conference on Intelligent Networking and Collaborative Systems*, pages 256–268. Springer.
- Mohamed K Elhadad, Kin Fun Li, and Fayez Gebali. 2020b. An ensemble deep learning technique to detect covid-19 misleading information. In *International Conference on Network-Based Information Systems*, pages 163–175.
- Tamer Elsayed, Preslav Nakov, Alberto Barrón-Cedeno, Maram Hasanain, Reem Suwaileh, Giovanni Da San Martino, and Pepa Atanasova. 2019. Overview of the clef-2019 checkthat! lab: Automatic identification and verification of claims. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 301–321. Springer.
- Riccardo Gallotti, Francesco Valle, Nicola Castaldo, Pierluigi Sacco, and Manlio De Domenico. 2020. Assessing the risks of “infodemics” in response to covid-19 epidemics. *arXiv preprint arXiv:2004.03997*.
- 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*, pages 845–854.
- Fatima Haouari, Maram Hasanain, Reem Suwaileh, and Tamer Elsayed. 2021. Arcov-19: The first arabic covid-19 twitter dataset with propagation networks. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop, WANLP 2021*.
- 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 2020*.
- Ireneus Kagashe, Zhijun Yan, and Imran Suheryani. 2017. Enhancing seasonal influenza surveillance: topic analysis of widely used medicinal drugs using twitter data. *Journal of medical Internet research*, 19(9):e315.
- Ling Min Serena Khoo, Hai Leong Chieu, Zhong Qian, and Jing Jiang. 2020. Interpretable rumor detection in microblogs by attending to user interactions. *arXiv preprint arXiv:2001.10667*.
- Ramez Kouzy, Joseph Abi Jaoude, Afif Kraitem, Molly B El Alam, Basil Karam, Elio Adib, Jabra Zarka, Cindy Traboulsi, Elie W AkI, and Khalil Baddour. 2020. Coronavirus goes viral: quantifying the covid-19 misinformation epidemic on twitter. *Cureus*, 12(3).
- Yang Liu and Yi-Fang Brook Wu. 2018. Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

- Jing Ma, Wei Gao, and Kam-Fai Wong. 2017. Detect rumors in microblog posts using propagation structure via kernel learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 708–717.
- Hamdy Mubarak and Sabit Hassan. 2020. Ar-corona: Analyzing arabic tweets in the early days of coronavirus (covid-19) pandemic. *arXiv preprint arXiv:2012.01462*.
- Melissa Roy, Nicolas Moreau, Cécile Rousseau, Arnaud Mercier, Andrew Wilson, and Laëtitia Atlani-Duault. 2020. Ebola and localized blame on social media: analysis of twitter and facebook conversations during the 2014–2015 ebola epidemic. *Culture, Medicine, and Psychiatry*, 44(1):56–79.
- Shaden Shaar, Giovanni Da San Martino, Nikolay Babulikov, and Preslav Nakov. 2020. That is a known lie: Detecting previously fact-checked claims. *arXiv preprint arXiv:2005.06058*.
- Gautam Kishore Shahi, Anne Dirkson, and Tim A Majchrzak. 2020. An exploratory study of covid-19 misinformation on twitter. *arXiv preprint arXiv:2005.05710*.
- Lisa Singh, Shweta Bansal, Leticia Bode, Ceren Budak, Guangqing Chi, Kornraphop Kawintiranon, Colton Padden, Rebecca Vanarsdall, Emily Vraga, and Yanchen Wang. 2020. A first look at covid-19 information and misinformation sharing on twitter. *arXiv preprint arXiv:2003.13907*.
- Soroush Vosoughi, Mostafa ‘Neo’ Mohsenvand, and Deb Roy. 2017. Rumor gauge: Predicting the veracity of rumors on twitter. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 11(4):1–36.
- Arkaitz Zubiaga, Maria Liakata, Rob Procter, Geraldine Wong Sak Hoi, and Peter Tolmie. 2016. Analysing how people orient to and spread rumours in social media by looking at conversational threads. *PloS one*, 11(3).