

Natural Language Inference with Self-Attention for Veracity Assessment of Pandemic Claims

Miguel Arana-Catania^{1,2}, Elena Kochkina^{3,2}, Arkaitz Zubiaga³, Maria Liakata^{3,2},
Rob Procter^{1,2}, Yulan He^{1,2}

¹Department of Computer Science, University of Warwick, UK

²The Alan Turing Institute, UK

³Queen-Mary University of London, UK

Abstract

We present a comprehensive work on automated veracity assessment from dataset creation to developing novel methods based on Natural Language Inference (NLI), focusing on misinformation related to the COVID-19 pandemic. We first describe the construction of the novel PANACEA dataset consisting of heterogeneous claims on COVID-19 and their respective information sources. The dataset construction includes work on retrieval techniques and similarity measurements to ensure a unique set of claims. We then propose novel techniques for automated veracity assessment based on Natural Language Inference including graph convolutional networks and attention based approaches. We have carried out experiments on evidence retrieval and veracity assessment on the dataset using the proposed techniques and found them competitive with SOTA methods, and provided a detailed discussion.

1 Introduction

In recent years, and particularly with the emergence of the COVID-19 pandemic, significant efforts have been made to detect misinformation online with the aim of mitigating its impact. With this objective, researchers have proposed numerous approaches and released datasets that can help with the advancement of research in this direction.

Most existing datasets (D’Ullizia et al., 2021) focus on a single medium (e.g., Twitter, Facebook, or specific websites), a unique information domain (e.g., health information, general news, or scholarly papers), a type of information (e.g., general claims or news), or a specific application (e.g., verifying claims, or retrieving useful information). This inevitably results in a limited focus on what is a complex, multi-faceted phenomenon. With the aim of furthering research in this direction, the contributions of our work are twofold: (1) creating a new comprehensive dataset of misinformation claims, and (2) introducing two novel approaches to veracity assessment.

In the first part of our work, we contribute to the global effort on addressing misinformation in the context of COVID-19 by creating a dataset for PANdemic Ai Claim vEracity Assessment, called the **PANACEA dataset**. It is a new dataset that combines different data sources with different foci, thus enabling a comprehensive approach that combines different media, domains and information types. To this effect our dataset brings together a heterogeneous set of *True* and *False* COVID claims and online sources of information for each claim. The collected claims have been obtained from online fact-checking sources, existing datasets and research challenges. We have identified a large overlap of claims between different sources and even within each source or dataset. Thus, given the challenges of aggregating multiple data sources, much of our efforts in dataset construction has focused on eliminating repeated claims. Distinguishing between different formulations of the same claim and nuanced variations that include additional information is a challenging task. Our dataset is presented in a large and a small version, accounting for different degrees of such similarity. Finally, the homogenisation of datasets and information media has presented an additional challenge, since fact-checkers use different criteria for labelling the claims, requiring a specific review of the different kinds of labels in order to combine them.

In the second part of our work, we propose NLI-SAN and NLI-graph, two novel veracity assessment approaches for automated fact-checking of the claims. Our proposed approaches are centred around the use of Natural Language Inference (NLI) and contextualised representations of the claims and evidence. NLI-SAN combines the inference relation between claims and evidence with attention techniques, while NLI-graph builds on graphs considering the relationship between all the different pieces of evidence and the claim.

Specifically we make the following contributions:

- We describe the development of a comprehensive COVID fact-checking dataset, PANACEA, as a result of aggregating and de-duplicating a set of heterogeneous data sources. The dataset is available in the project website¹, as well as a fully operational search platform to find and verify COVID-19 claims implementing the proposed approaches.
- We propose two novel approaches to claim verification, NLI-SAN and NLI-graph.
- We perform an evaluation of both evidence retrieval and the application of our proposed veracity assessment methods on our constructed dataset. Our experiments show that NLI-SAN and NLI-graph have state-of-the-art performance on our dataset, beating GEAR (Zhou et al., 2019) and matching KGAT (Liu et al., 2020). We discuss challenging cases and provide ideas for future research directions.

2 Related Work

COVID-19 and misinformation datasets. Comprehensive information on COVID-19 datasets is provided in Appendix A. Such datasets include the CoronaVirusFacts/DatosCoronaVirus Alliance Database, the largest existing collection of COVID claims and the largest existing network of journalists working together on COVID misinformation, an essential reference for our work; COVID-19-TweetIDs (Chen et al., 2020) the widest dataset of COVID tweets with more than 1 billion tweets; Cord-19: The COVID-19 open research dataset (Wang et al., 2020a), the largest downloadable set of scholarly articles on the pandemic with nearly 200,000 articles. General misinformation datasets linked to our verification work include: Emergent (Ferreira and Vlachos, 2016) collection of 300 labeled claims by journalists; LIAR (Wang, 2017) with 12,836 statements from PolitiFact with detailed justifications; FakeNewsNet (Shu et al., 2020) collecting not only claims from news content, but also social context and spatio-temporal information; NELA-GT-2018 (Nørregaard et al., 2019) with 713,534 articles from 194 news outlets; FakeHealth (Dai et al., 2020) collecting information from HealthNewsReview, a project critically analysing claims about health care interventions; PUBHEALTH (Kotonya and Toni, 2020) with 11,832 claims related to health topics; FEVER (Thorne et al., 2018a) as well as its later versions

¹<https://panacea2020.github.io/>
<https://doi.org/10.5281/zenodo.6493847>

FEVER 2.0 (Thorne et al., 2018b) and FEVEROUS (Aly et al., 2021), containing claims based on Wikipedia and therefore constituting a well-defined, informative and non-duplicated information corpus; SciFact (Wadden et al., 2020) also from a very different domain, containing 1,409 scientific claims. Our dataset is a real-world dataset bringing together heterogeneous sources, domains and information types.

Approaches to claim veracity assessment. We employ our dataset for automated fact-checking and veracity assessment (Zeng et al., 2021). Researchers such as Hanselowski et al. (2018); Yoneda et al. (2018); Luken et al. (2018); Soleimani et al. (2020); Pradeep et al. (2021) analysed the veracity relation between the claim and each piece of evidence independently, combining this information later. Other authors considered multiple pieces of evidence together (Thorne et al., 2018a; Nie et al., 2019; Stambach and Neumann, 2019). Different pieces of evidence have been previously combined using graph neural networks (Zhou et al., 2019; Liu et al., 2020; Zhong et al., 2020). Many of these authors have centred their techniques on the use of NLI (Chen et al., 2017; Ghaeini et al., 2018; Parikh et al., 2016; Li et al., 2019) to verify the claim. In our work we also make use of NLI results of claim-evidence pairs, but propose alternative approaches built on a self-attention network and a graph convolutional network for veracity assessment.

3 Dataset Construction

This section describes our dataset construction by selecting COVID-19 related data sources (§3.1), and applying information retrieval and re-ranking techniques to remove duplicate claims (§3.2).

3.1 Data Sources

We first identified a set of COVID-19 related data sources to build our dataset. Our aim is to have the largest compilation of non-overlapping, labelled and verified claims from different media and information domains (Twitter, Facebook, general websites, academia), and used for different applications (media reporting, veracity evaluation, information retrieval challenges, etc.). We have included any large dataset or media, to our knowledge, related to that objective that includes claims together with their information sources. The data sources identified are shown in Table 1. More details and pre-processing steps are presented in Appendix A. By

Data Source	Description	Domain	No. of claims (False / True)
CoronaVirusFacts Database	Published by Poynter, this online source combines fact-checking articles from more than 100 fact-checkers from all over the world, being the largest journalist fact-checking collaboration on the topic worldwide.	Heterogeneous	11,647 (11,647 / 0)
CoAID dataset (Cui and Lee, 2020)	This contains fake news from fact-checking websites and real news from health information websites, health clinics, and public institutions.	News	5,485 (953 / 4,532)
MM-COVID (Li et al., 2020)	This multilingual dataset contains fake and true news collected from Poynter and Snopes.	News	3,409 (2,035 / 1,374)
CovidLies (Hossain et al., 2020)	This contains a curated list of common misconceptions about COVID appearing in social media, carefully reviewed to contain very relevant and unique claims.	Social media	62 (62 / 0)
TREC Health Misinformation track	Research challenge using claims on the health domain focused on information retrieval from general websites through the Common Crawl corpus (commoncrawl.org).	General websites	46 (39 / 7)
TREC COVID challenge (Voorhees et al., 2021; Roberts et al., 2020)	Research challenge using claims on the health domain focused on information retrieval from scholar peer-reviewed journals through the CORD19 dataset (Wang et al., 2020a), the largest existing compilation of COVID-related articles.	Scholar papers	40 (3 / 37)

Table 1: Data sources used for the construction of our dataset. The last column shows the number of claims before de-duplication.

processing and combining these sources we obtained 20,689 initial claims.

3.2 Claim De-duplication

We processed claims and removed: exact duplicates; claims making only a direct reference to existing content in other media (audio, video, photos); automatically obtained content not representing claims; entries with claims or fact-checking sources in languages other than English.

The similarity of claims was then analysed using: BM25 (Robertson et al., 1995; Crestani et al., 1998; Robertson and Zaragoza, 2009) and BM25 with MonoT5 re-ranking (Nogueira et al., 2020). BM25 is a commonly-used ranking function that estimates the relevance of documents to a given query. MonoT5 uses a T5 model trained using as input the template ‘Query: [query] Document: [doc] Relevant:’, fine-tuned to produce as output the token ‘True’ or ‘False’. A softmax layer applied to those tokens gives the respective relevance probabilities. These methods are used to identify not only claims similar in content, but also distinct claims that are sufficiently relevant when searching for information about them. This ensures that the claims presented are unique, and avoids overlap between training and testing cases when using the data to train veracity assessment models. These methods were carried out using

Pyserini² and PyGaggle³. The set of claims was indexed and a search was performed for each of the claims to detect similar claims. We created two versions of the dataset by varying the similarity threshold between claims. The LARGE dataset excludes claims with a 90% probability of being similar, while in the SMALL dataset the probability is increased to 99%, as obtained through the MonoT5 model. These thresholds were chosen empirically by manual inspection of the results with simultaneous consideration of the efficiency of the method.

As a further assessment of the uniqueness of the claims, we evaluated the de-duplication process using BERTScore⁴ (Zhang et al., 2019) on the resulting datasets. We used the linked code with a RoBERTa-large model with baseline rescaling. We compared each claim with all the other claims in the dataset and kept the score of the most similar match. The mean and standard deviation, and the 90th percentile of claim similarity values are shown in the upper part of Table 3. The average claim similarity has been drastically reduced in the LARGE dataset compared to the original dataset and further reduced in the SMALL dataset.

To illustrate the difference between the two ver-

²<https://github.com/castorini/pyserini>

³<https://github.com/castorini/pygaggle>

⁴https://github.com/Tiiiger/bert_score

<p>Claim 1: Losing your sense of smell may be an early symptom of COVID-19.</p> <hr/> <p>Exclude from LARGE and SMALL: Loss of smell may suggest milder COVID-19.</p> <p>Exclude from SMALL only: Loss of smell and taste validated as COVID-19 symptoms in patients with high recovery rate.</p> <hr/> <p>Claim 2: COVID-19 hitting some African American communities harder.</p> <hr/> <p>Exclude from LARGE and SMALL: The African American community is being hit hard by COVID-19.</p> <p>Exclude from SMALL only: COVID-19 impacts in African-Americans are different from the rest of the U.S. population.</p>

Table 2: Claim de-duplication examples.

sions of the dataset, we present some examples of claims in Table 2. For Claim 1, the semantically similar claim ‘*Loss of smell may suggest milder COVID-19*’ is identified and excluded from both LARGE and SMALL datasets. But the claim ‘*Loss of smell and taste validated as COVID-19 symptoms in patients with high recovery rate*’, which includes mentions of another symptom and the recovery rate, is only excluded from the SMALL dataset. For Claim 2, the rephrased claim ‘*The African American community is being hit hard by COVID-19*’ is excluded from both datasets. But the claim ‘*COVID-19 impacts in African-Americans are different from the rest of the U.S. population*’, which refers specifically to the U.S. population, is only excluded from the SMALL dataset.

3.3 Dataset Statistics

Our final dataset statistics are shown in the lower part of Table 3, where the original and the two reduced versions are presented. After the steps described in Section 3.2 the LARGE dataset contains 5,143 claims, and the SMALL version 1,709 claims.

Category	Orig.	LARGE	SMALL
Similarity	0.67 ± 0.23	0.43 ± 0.13	0.37 ± 0.14
$\eta_{.90}$	0.99	0.60	0.56
False	14,739	1,810	477
True	5,950	3,333	1,232
Total	20,689	5,143	1,709

Table 3: The average claim similarity values and the PANACEA LARGE and SMALL dataset statistics. $\eta_{.90}$ denotes the 90th percentile value.

Example claims contained in the dataset are shown in Table 4. Each of the entries in the dataset

contains the following information:

- *Claim*. Text of the claim.
- *Claim label*. The labels are: *False*, and *True*.
- *Claim source*. The sources include mostly fact-checking websites, health information websites, health clinics, public institutions sites, and peer-reviewed scientific journals.
- *Original information source*. Information about which general information source was used to obtain the claim.
- *Claim type*. The different types, explained in Section A.2, are: *Multimodal*, *Social Media*, *Questions*, *Numerical*, and *Named Entities*.

4 Claim Veracity Assessment

We develop a pipeline approach consisting of three steps: document retrieval, sentence retrieval and veracity assessment for claim veracity evaluation. Given a claim, we first retrieve the most relevant documents from COVID-19 related sources and then further retrieve the top N most relevant sentences. Considering each retrieved sentence as evidence, we train a veracity assessment model to assign a *True* or *False* label to the claim.

4.1 Document Retrieval

Document Dataset. In order to retrieve documents relevant to the claims, we first construct an additional dataset containing documents obtained from reliable COVID-19 related websites. These information sources represent a real-world comprehensive database about COVID-19 that can be used as a primary source of information on the pandemic. We have selected four organisations from which to collect the information: (1) *Centers for Disease Control and Prevention (CDC)*, national public health agency of the United States; (2) *European Centre for Disease Prevention and Control (ECDC)*, EU agency aimed at strengthening Europe’s defenses against infectious diseases; (3) *WebMD*, online publisher of news and information on health; and (4) *World Health Organization (WHO)*, agency of the United Nations responsible for international public health.

All pages corresponding to the COVID-19 sub-domains of each site have been downloaded. The

Claim	Category	Source	Orig. data src.	Type
Stroke Scans Could Reveal COVID-19 Infection.	True	ScienceDaily	CoAID	
Whiskey and honey cure coronavirus.	False	Independent news site	CovidLies	
COVID-19 is more deadly than Ebola or HIV.	False	Australian Associated Press	Poynter	
Dextromethorphan worsens COVID-19.	True	Nature	TREC Health Misinformation track	
ACE inhibitors increase risk for coronavirus.	False	Infectious Disorders - Drug Targets journal	TREC COVID challenge	
Nancy Pelosi visited Wuhan, China, in November 2019, just a month before the COVID-19 outbreak there.	False	Snopes	MM-COVID	Named Entity, Numerical content

Table 4: Example entries in the constructed PANACEA dataset.

web content was downloaded using the BeautifulSoup⁵ and Scrapy⁶ packages. Social networking sites and non-textual content were discarded. In total 19,954 web pages have been collected. The list of websites and the full content of each website constitute this additional dataset used for document retrieval. This dataset is enhanced with some additional websites used only in the document retrieval experiments, detailed in Section 5.1.

Method. Information sources were indexed by creating a Pyserini Lucene index and PyGaggle was used to implement a re-ranker model on the results. The documents were split into paragraphs of 300 tokens segmented with a BERT tokenizer.

To retrieve the information we first used a BM25 score. Additionally, we tested the effect of multi-stage retrieval by re-ranking the initial results using MonoBERT (Nogueira et al., 2019) and MonoT5 models, and query expansion using RM3 pseudo-relevance feedback (Abdul-Jaleel et al., 2004) on the BM25 results (Lin, 2019; Yang et al., 2019).

MonoBERT uses a BERT model trained using as inputs the query and each of the documents to be re-ranked encoded together (`[CLS] query [SEP] doc [SEP]`), and then the `[CLS]` output token is passed to a single layer fully-connected network that produces the probability of the document being relevant to the query.

4.2 Sentence Retrieval

For each claim, once documents are retrieved using BM25 and MonoT5 re-ranking of the top 100 BM25 results, we then further retrieve the N most similar sentences obtained from the 10 most relevant documents. The relevance of the sentences is

calculated using cosine similarity in relation to the original claim. The similarity is obtained with the pre-trained model MiniLM-L12-v2 (Wang et al., 2020b), using Sentence-Transformers⁷ (Reimers and Gurevych, 2019) to encode the sentences.

4.3 Veracity Assessment

We propose two veracity assessment approaches built on the NLI results of claim-evidence pairs. For each of the most similar sentences (pieces of evidence) retrieved for a claim, we apply the pre-trained NLI model RoBERTa-large-MNLI⁸ (Liu et al., 2019). This model acts as a cross-encoder on pairs of sentences, trained to detect the relationship between the two sentences: *contradiction*, *neutrality*, or *entailment*. The model is trained on the Multi-Genre Natural Language Inference (MultiNLI) dataset (Williams et al., 2018). The inference results are then used in our proposed approaches described below.

NLI-SAN. The first approach, named NLI-SAN, incorporates the inference results of claim-evidence pairs into a Self-Attention Network (SAN) (See Figure 1a). First, a claim is paired with each piece of retrieved relevant evidence. Each pair (c, e_i) is fed into a RoBERTa-large⁸ model, and the last hidden layer output S_i is used as its representation. Additionally, each pair is also fed to the mentioned RoBERTa-large-MNLI⁸ model obtaining I_i , a triplet containing the probability of *contradiction*, *neutrality*, or *entailment*.

$$\begin{aligned} S_i &= \text{RoBERTa}(c, e_i) \\ I_i &= \text{RoBERTa}_{\text{NLI}}(c, e_i) \end{aligned} \quad (1)$$

⁵<https://www.crummy.com/software/BeautifulSoup/>

⁶<https://scrapy.org/>

⁷<https://github.com/UKPLab/sentence-transformers>

⁸<https://huggingface.co/>

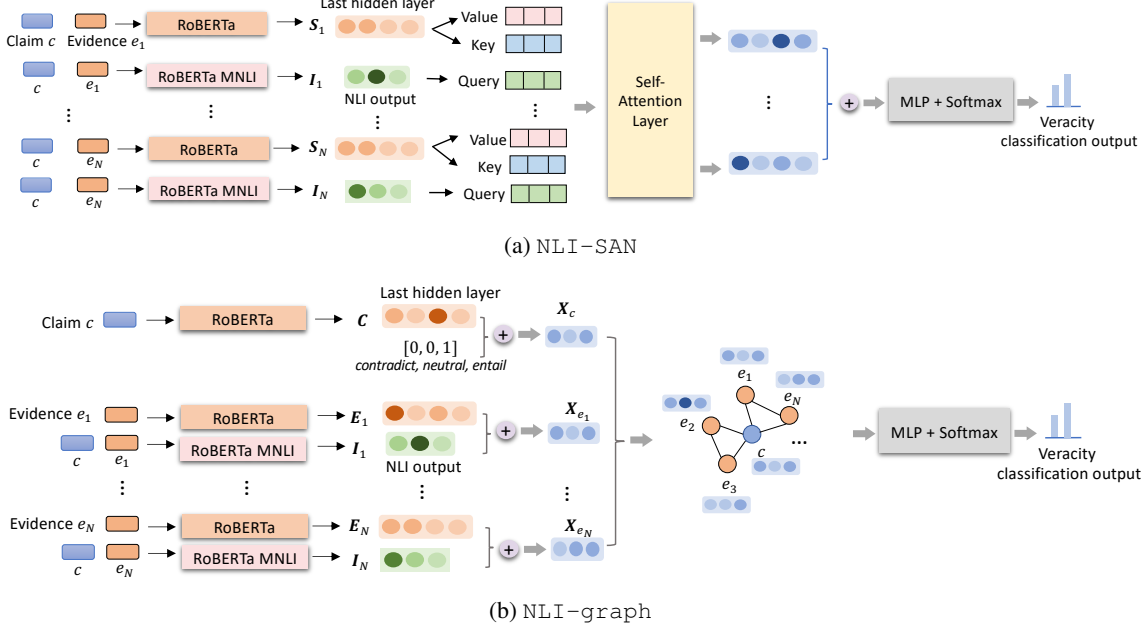


Figure 1: Proposed veracity classification models. \oplus means concatenation.

The sentence representation is combined with the NLI output through a Self Attention Network (SAN) (Galassi et al., 2020; Bahdanau et al., 2015).

The RoBERTa-encoded claim-evidence representation \mathbf{S}_i with length $n_S = n_K = n_V$ is mapped onto a Key $\mathbf{K} \in \mathbb{R}^{n_K \times d_K}$ and a Value $\mathbf{V} \in \mathbb{R}^{n_V \times d_V}$, while the NLI output \mathbf{I}_i of each claim-evidence pair is mapped onto a Query $\mathbf{Q} \in \mathbb{R}^{n_Q \times d_Q}$. The representation dimensionality is $d_K = d_V = d_Q = 1024$. The attention function is defined as:

$$\text{Att}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}(\mathbf{Q}\mathbf{K}^\top / \sqrt{d})\mathbf{V} \quad (2)$$

While standard attention mechanisms use only the sentence representation information for the Key, Value and Query, here the inference information is used in the Query. This *attention* mechanism is applied to each of the claim-evidence pairs, and the outputs are concatenated into an output \mathbf{O}_{SAN} that is passed through a Multi-Layer Perceptron (MLP) with hidden size d_h and a Softmax layer to generate the veracity classification output.

$$\hat{y} = \text{softmax}(\text{MLP}_{\text{ReLU}}(\mathbf{O}_{\text{SAN}})) \quad (3)$$

NLI-graph. We propose an alternative approach based on Graph Convolutional Networks (GCN). First, for each claim-evidence pair, we derive RoBERTa-encoded representations for the claims and evidence separately (using the pooled output of the last layer) and obtain NLI results of the pairs

as before.

$$\mathbf{C}_i = \text{RoBERTa}(c); \quad \mathbf{E}_i = \text{RoBERTa}(e_i) \quad (4)$$

$$\mathbf{I}_i = \text{RoBERTa}_{\text{NLI}}(c, e_i) \quad (5)$$

Next, we build an evidence network in which the central node is the claim and the rest of the nodes are the evidence. Two nodes are linked if their similarity value exceeds a pre-defined threshold, which is empirically set to 0.9 by comparing the results of the experimental evaluation described in the following section using different thresholds. The similarity is considered between claim and evidence, but also between pieces of evidence. Similarity calculation is performed following the same approach as in Section 4.2. The features considered in each evidence node are the concatenation of \mathbf{E}_i and \mathbf{I}_i . For the claim node we use its representation \mathbf{C}_i and a unity vector $(0, 0, 1)$ for the inference. The network is implemented with the package PyTorch Geometric (Fey and Lenssen, 2019), using in the first layer the GCNConv operator (Kipf and Welling, 2016) with 50 output channels and self-loops to the nodes, represented by:

$$\mathbf{X}' = \hat{\mathbf{D}}^{-1/2} \hat{\mathbf{A}} \hat{\mathbf{D}}^{-1/2} \mathbf{X} \mathbf{W}, \quad (6)$$

where \mathbf{X} is the matrix of node feature vectors, $\hat{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ denotes the adjacency matrix with inserted self-loops, $\hat{D}_{ii} = \sum_{j=0} \hat{A}_{ij}$ its diagonal degree matrix, and \mathbf{W} is a trainable weight matrix.

Once the node representation is updated via GCN, all the node representations are averaged

and passed to the MLP and the Softmax layer to generate the final veracity classification output.

$$\hat{y} = \text{softmax}(\text{MLP}_{\text{ReLU}}(\text{O}_{\text{graph}})) \quad (7)$$

5 Experiments

In this section, we perform a twofold evaluation: We first evaluate our document retrieval methods (presented in §4.1) on obtaining information relevant to the dataset claims from a database of COVID-19 related websites. We subsequently present an evaluation of the veracity assessment approaches for the claims (described in §4.3).

5.1 Document Retrieval

In order to evaluate our document retrieval methods, we need the gold-standard relevant document for each claim. Therefore, in the documents dataset described in section 4.1 we additionally include the web content referenced in each of the information sources used to compile our claim dataset:

The CoronaVirus Alliance Database. All web pages from the websites referenced as fact-checking sources for the claims have been downloaded from 151 different domains.

CoAID dataset. We downloaded the websites used as fact-checking sources of false claims and the websites where correct information on true claims is gathered from 68 different domains.

MM-COVID. We collected both fact-checking sources and reliable information related to the claims of this dataset from 58 web domains.

CovidLies dataset. We include the web content used as fact-checking sources of the misconceptions from 39 domains.

We have not included web content from the TREC Challenges, as each of them is performed on a very large dataset specific to each challenge (CORD19 and Common Crawl corpus), as explained previously. Note that in our subsequent experiments, we have excluded all fact-checking websites to avoid finding directly the claim references. The results of the document retrieval are presented in Table 5. For each claim, the $\text{precision}@k$ is defined as 1 if the relevant result is retrieved in the top k list and 0 otherwise.

We can see that by using BM25, it is possible in many cases to retrieve the relevant results at the very top of our searches. Combining BM25 with MonoBERT did not offer any improvement. It even introduced noise to the retrieval results, leading to inferior performance compared to using BM25 only on AP@5 and AP@10. MonoT5 appears

	AP@5	AP@10	AP@20	AP@100
BM25	0.54	0.56	0.58	0.62
BM25+MonoBERT	0.52	0.55	0.58	0.62
BM25+MonoT5	0.55	0.58	0.60	0.62
BM25+RM3+MonoT5	0.51	0.53	0.55	0.57

Table 5: Document retrieval results. Average precision for different cut-offs. For the MonoBERT and MonoT5 cases, 100 initial results are retrieved in the first retrieval stage before re-ranking.

to be more effective, consistently improving the retrieval results across all metrics. Moreover for this dataset the use of query expansion using RM3 pseudo-relevance feedback on the BM25 results does not improve the results.

5.2 Veracity Assessment Evaluation

Here we evaluate our proposed NLI-SAN and NLI-graph veracity assessment approaches. To gain a better insight into the benefits of the proposed architectures, we conducted additional experiments on the variants of the models including:

- NLI, using only the NLI outputs of the claim-evidence pairs. The outputs are concatenated and then passed through the final classification layer to generate veracity classification results.
- NLI+sent, this is the ablated version of NLI-SAN without the self-attention layer. Here, the RoBERTa-encoded claim-evidence representations are concatenated with the NLI results and then fed to the classification layer to produce the veracity classification output.
- NLI+PSent, this is similar to the previous ablated version, but using the pooled representation of the claim-evidence pair to concatenate with the NLI result.
- NLI-graph_{abl}, this is the ablated version of NLI-graph in which the node representation is the NLI result of the corresponding claim-evidence pair without its RoBERTa-encoded representation.

For NLI, NLI+sent and NLI-SAN, we consider the 5 most similar sentences for each claim, obtained from the 10 most relevant documents of the information source database. Those documents are retrieved using BM25 and MonoT5 re-ranking of the top 100 BM25 results. For NLI-graph, NLI-graph_{abl} and NLI+PSent, in order to have enough nodes to benefit from the network structure, the number of retrieved sentences is increased to 30 for each claim, selected as the 3

Model	False			True			Macro F1
	Precision	Recall	F1	Precision	Recall	F1	
GEAR (Zhou et al., 2019)	0.81	0.60	0.69	0.85	0.94	0.89	0.79
KGAT (Liu et al., 2020)	0.89	0.96	0.92	0.98	0.95	0.97	0.94
NLI	0.48	0.24	0.31	0.75	0.90	0.82	0.56
NLI+Sent	0.91	0.87	0.89	0.95	0.97	0.96	0.92
NLI+PSent	0.87	0.72	0.79	0.90	0.96	0.93	0.86
NLI-SAN	0.93	0.89	0.91	0.96	0.97	0.97	0.94
NLI-graph _{-abl}	0.50	0.33	0.39	0.77	0.87	0.81	0.60
NLI-graph	0.89	0.83	0.86	0.94	0.96	0.95	0.90

Table 6: Veracity classification results on the PANACEA SMALL dataset. The best result in each column is highlighted in bold.

Model	False			True			Macro F1
	Precision	Recall	F1	Precision	Recall	F1	
GEAR (Zhou et al., 2019)	0.88	0.88	0.88	0.93	0.94	0.94	0.91
KGAT (Liu et al., 2020)	0.95	0.98	0.96	0.99	0.98	0.98	0.97
NLI	0.52	0.27	0.36	0.69	0.86	0.76	0.56
NLI+Sent	0.94	0.94	0.94	0.97	0.97	0.97	0.95
NLI+PSent	0.89	0.77	0.82	0.88	0.95	0.91	0.86
NLI-SAN	0.95	0.95	0.95	0.97	0.98	0.97	0.96
NLI-graph _{-abl}	0.60	0.43	0.50	0.73	0.84	0.78	0.64
NLI-graph	0.94	0.91	0.93	0.95	0.97	0.96	0.94

Table 7: Veracity classification results on the PANACEA LARGE dataset. The best result in each column is highlighted in bold.

most similar sentences from the top 10 retrieved documents. The retrieval procedure is as in sections 4.1 and 4.2. Details of parameter settings can be found in Appendix B. We compare against the SOTA methods GEAR⁹(Zhou et al., 2019) and KGAT¹⁰(Liu et al., 2020), with settings as described by the authors.

For all approaches we perform 5-fold cross-validation and report the averaged results on the SMALL dataset in Table 6. By using the NLI information alone it is possible to obtain reasonable results for the *True* claims, however, this is not the case for the most relevant *False* claims. Once we add sentence representations the efficiency of the method increases significantly. Using NLI-SAN instead of simply concatenating contextualised claim-evidence representations and NLI outputs further improves the results. A similar observation can be made in the results generated by NLI-graph and its variants; the contextualised representations of claim-evidence pairs are much more important than merely using the corresponding NLI values. We also note that using the graph version NLI-graph obtains better scores

than a non-graph model with the same information NLI+PSent, however the scores are still lower than the NLI-SAN method. Our method performs on a par with KGAT, while being simpler, and outperforms GEAR.

Complementing the results for the SMALL dataset, Table 7 presents the results for the LARGE dataset. In general, we observe improved performance for all models across all metrics for both classes compared to the results on the SMALL dataset. The previous results in the SMALL dataset constitute a more challenging case, since the uniqueness of the claims is increased and therefore the veracity assessment models are not able to learn from similar claims when performing the assessment.

5.3 Discussion

Our results show that in document retrieval, we have obtained values of around 0.6 from a simple term scoring and re-ranking retrieval model. However, this baseline represents only a rough measure of quality using this technique, since we have only evaluated the retrieval of a single document specific to each claim; we have not evaluated the quality of other retrieved documents.

⁹<https://github.com/thunlp/GEAR>

¹⁰<https://github.com/thunlp/KernelGAT>

The distinction into *True* and *False* claims can be rather coarse-grained. We note that initially we considered a larger number of veracity labels, including more nuanced cases that could be interesting to analyse (see A.1). However, we have not found a clear separation between complex cases and it would seem that different fact checkers do not follow the same conventions when labelling such cases. The development of datasets especially focused on such nuanced cases may be therefore an important line of work in the future, together with the development of techniques for these more complex situations.

In analysing misclassified claims, we note some interesting cases. The scope and globality of the pandemic imply that similar issues are mentioned repeatedly on multiple occasions, yet claims to be verified may include nuances or specificities. This is challenging as it is easy to retrieve information that omits relevant nuances. E.g. The claim “*Baron Trump had COVID-19, Melania Trump says*” retrieves sentences such as “*Rudy Giuliani has tested positive for COVID-19, Trump says.*” with a similar structure and mentions but missing the key name. This type of situation could be addressed by using Named Entity Recognition (NER) methods that prioritise matching between the entities involved in the claim and the information sources. See e.g. (Taniguchi et al., 2018; Nooralahzadeh and Øvreliid, 2018).

Other interesting cases involve claims for which documents with adequate information are retrieved, but the sentences containing evidence cannot be identified because they are too different from the original claim. E.g. The claim “*Vice President of Bharat Biotech got a shot of the indigenous COV-AXIN vaccine*” retrieves correct documents on the issue. Similar sentences are retrieved such as “*Covaxin which is being developed by Bharat Biotech is the only indigenous vaccine that is approved for emergency use.*”. Despite being similar such retrieved sentences give no information about the claimed situation. In the retrieved document, the sentence “*The pharmaceutical company, has in a statement, denied the claim and said the image shows a routine blood test.*” contains the essential information to debunk the original claim but is missed by the sentence retrieval engine as it is very different from the claim (See Table A1 in Appendix C for other examples).

Such cases are more difficult to deal with, as the similarity between claim and evidence is cer-

tainly a good indicator of relevance. Nevertheless, these cases are very interesting for future work using more complex approaches. We have made an initial attempt to address this problem by representing claims and retrieved documents using Abstract Meaning Representation (Banarescu et al., 2013) in order to better select relevant information. Although the results were not satisfactory, it may be an interesting avenue for future exploration. Another line of future work is the design of strategies against adversarial attacks to mitigate possible risks to our system.

6 Conclusions

We have presented a novel dataset that aggregates a heterogeneous set of COVID-19 claims categorised as *True* or *False*. Aggregation of heterogeneous sources involved a careful deduplication process to ensure dataset quality. Fact-checking sources are provided for veracity assessment, as well as additional information sources for *True* claims. Additionally, claims are labelled with sub-types (Multimodal, Social Media, Questions, Numerical, and Named Entities).

We have performed a series of experiments using our dataset for information retrieval through direct retrieval and using a multi-stage re-ranker approach. We have proposed new NLI methods for claim veracity assessment, attention-based NLI-SAN and graph-based NLI-graph, achieving in our dataset competitive results with the GEAR and KGAT state-of-the-art models. We have also discussed challenging cases and provided ideas for future research directions.

Acknowledgements

This work was supported by the UK Engineering and Physical Sciences Research Council (grant no. EP/V048597/1, EP/T017112/1). ML and YH are supported by Turing AI Fellowships funded by the UK Research and Innovation (grant no. EP/V030302/1, EP/V020579/1).

References

- Nasreen Abdul-Jaleel, James Allan, W Bruce Croft, Fernando Diaz, Leah Larkey, Xiaoyan Li, Mark D Smucker, and Courtney Wade. 2004. Umass at trec 2004: Novelty and hard. *Computer Science Department Faculty Publication Series*, page 189.
- Muhammad Abdul-Mageed, AbdelRahim Elmadany, El Moatez Billah Nagoudi, Dinesh Pabbi, Kunal

- Verma, and Rannie Lin. 2021. [Mega-COV: A billion-scale dataset of 100+ languages for COVID-19](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3402–3420, Online. Association for Computational Linguistics.
- Rami Aly, Zhijiang Guo, Michael Sejr Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, Oana Cocarascu, and Arpit Mittal. 2021. [The fact extraction and VERification over unstructured and structured information \(FEVEROUS\) shared task](#). In *Proceedings of the Fourth Workshop on Fact Extraction and VERification (FEVER)*, pages 1–13, Dominican Republic. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking. In *Proceedings of the 7th linguistic annotation workshop and interoperability with discourse*, pages 178–186.
- Cambridge journals Coronavirus Free Access Collection. 2020. Cambridge journals coronavirus free access collection. <https://www.cambridge.org/core/browse-subjects/medicine/coronavirus-free-access-collection>.
- Emily Chen, Kristina Lerman, and Emilio Ferrara. 2020. Tracking social media discourse about the covid-19 pandemic: Development of a public coronavirus twitter data set. *JMIR Public Health and Surveillance*, 6(2):e19273.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. Enhanced lstm for natural language inference. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1657–1668.
- Qingyu Chen, Alexis Allot, and Zhiyong Lu. 2021. Lit-covid: an open database of covid-19 literature. *Nucleic acids research*, 49(D1):D1534–D1540.
- COVID-19 Data Portal (EU). 2020. Covid-19 data portal (eu). <https://www.covid19dataportal.org/>.
- Fabio Crestani, Mounia Lalmas, Cornelis J Van Rijsbergen, and Iain Campbell. 1998. “is this document relevant?... probably” a survey of probabilistic models in information retrieval. *ACM Computing Surveys (CSUR)*, 30(4):528–552.
- Limeng Cui and Dongwon Lee. 2020. Coaid: Covid-19 healthcare misinformation dataset. *arXiv preprint arXiv:2006.00885*.
- Enyan Dai, Yiwei Sun, and Suhang Wang. 2020. Ginger cannot cure cancer: Battling fake health news with a comprehensive data repository. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, pages 853–862.
- Dimitar Dimitrov, Erdal Baran, Pavlos Fafalios, Ran Yu, Xiaofei Zhu, Matthäus Zloch, and Stefan Dietze. 2020. Tweetscov19-a knowledge base of semantically annotated tweets about the covid-19 pandemic. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 2991–2998.
- Arianna D’Ulizia, Maria Chiara Caschera, Fernando Ferri, and Patrizia Grifoni. 2021. Fake news detection: a survey of evaluation datasets. *PeerJ Computer Science*, 7:e518.
- Elsevier journals Novel Coronavirus Information Center. 2020. Elsevier journals novel coronavirus information center. <https://www.elsevier.com/connect/coronavirus-information-center>.
- William Ferreira and Andreas Vlachos. 2016. Emergent: a novel data-set for stance classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: Human language technologies*, pages 1163–1168.
- Matthias Fey and Jan E. Lenssen. 2019. Fast graph representation learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- Andrea Galassi, Marco Lippi, and Paolo Torrioni. 2020. Attention in natural language processing. *IEEE Transactions on Neural Networks and Learning Systems*.
- Reza Ghaeini, Sadid A Hasan, Vivek Datla, Joey Liu, Kathy Lee, Ashequl Qadir, Yuan Ling, Aaditya Prakash, Xiaoli Fern, and Oladimeji Farri. 2018. Dr-bilstm: Dependent reading bidirectional lstm for natural language inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1460–1469.
- Andreas Hanselowski, Hao Zhang, Zile Li, Daniil Sorokin, Benjamin Schiller, Claudia Schulz, and Iryna Gurevych. 2018. Ukp-athene: Multi-sentence textual entailment for claim verification. *EMNLP 2018*, page 103.
- 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*, Online. Association for Computational Linguistics.

- Xiaolei Huang, Amelia Jamison, David Broniatowski, Sandra Quinn, and Mark Dredze. 2020. [Coronavirus twitter data: A collection of covid-19 tweets with automated annotations](http://twitterdata.covid19dataresources.org/index). [Http://twitterdata.covid19dataresources.org/index](http://twitterdata.covid19dataresources.org/index).
- Daniel Kerchner and Laura Wrubel. 2020. Coronavirus tweet ids. *Harvard Dataverse*.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Neema Kotonya and Francesca Toni. 2020. Explainable automated fact-checking for public health claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7740–7754.
- Rabindra Lamsal. 2021. Design and analysis of a large-scale covid-19 tweets dataset. *Applied Intelligence*, 51(5):2790–2804.
- Tianda Li, Xiaodan Zhu, Quan Liu, Qian Chen, Zhigang Chen, and Si Wei. 2019. Several experiments on investigating pretraining and knowledge-enhanced models for natural language inference. *arXiv preprint arXiv:1904.12104*.
- Yichuan Li, Bohan Jiang, Kai Shu, and Huan Liu. 2020. [Mm-covid: A multilingual and multimodal data repository for combating covid-19 disinformation](#).
- Jimmy Lin. 2019. The neural hype and comparisons against weak baselines. In *ACM SIGIR Forum*, volume 52, pages 40–51. ACM New York, NY, USA.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Zhenghao Liu, Chenyan Xiong, Maosong Sun, and Zhiyuan Liu. 2020. Fine-grained fact verification with kernel graph attention network. In *The 58th annual meeting of the Association for Computational Linguistics (ACL)*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Jackson Luken, Nanjiang Jiang, and Marie-Catherine de Marneffe. 2018. Qed: A fact verification system for the fever shared task. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 156–160.
- MedRN medical research network SSRN Coronavirus Infectious Disease Research Hub. 2020. Medrn medical research network ssrn coronavirus infectious disease research hub. <https://www.ssrn.com/index.cfm/en/coronavirus/>.
- Shahan Ali Memon and Kathleen M Carley. 2020. Characterizing covid-19 misinformation communities using a novel twitter dataset. In *CEUR Workshop Proceedings*, volume 2699.
- Yixin Nie, Haonan Chen, and Mohit Bansal. 2019. Combining fact extraction and verification with neural semantic matching networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6859–6866.
- Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. 2020. Document ranking with a pre-trained sequence-to-sequence model. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 708–718.
- Rodrigo Nogueira, Wei Yang, Kyunghyun Cho, and Jimmy Lin. 2019. Multi-stage document ranking with bert. *arXiv preprint arXiv:1910.14424*.
- Farhad Nooralahzadeh and Lilja Øvrelid. 2018. Siriusltg: An entity linking approach to fact extraction and verification. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 119–123.
- Jeppe Nørregaard, Benjamin D Horne, and Sibel Adalı. 2019. Nela-gt-2018: A large multi-labelled news dataset for the study of misinformation in news articles. In *Proceedings of the international AAAI conference on web and social media*, volume 13, pages 630–638.
- Oxford journals resources on COVID-19. 2020. Oxford journals resources on covid-19. <https://academic.oup.com/journals/pages/coronavirus>.
- Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2249–2255.
- Ronak Pradeep, Xueguang Ma, Rodrigo Nogueira, and Jimmy Lin. 2021. Vera: Prediction techniques for reducing harmful misinformation in consumer health search. In *Proceedings of the 44th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2021)*.
- Umair Qazi, Muhammad Imran, and Ferda Ofli. 2020. Geocov19: a dataset of hundreds of millions of multilingual covid-19 tweets with location information. *SIGSPATIAL Special*, 12(1):6–15.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. 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)*, pages 3982–3992.

- Kirk Roberts, Tasmeeer Alam, Steven Bedrick, Dina Demner-Fushman, Kyle Lo, Ian Soboroff, Ellen Voorhees, Lucy Lu Wang, and William R Hersh. 2020. Trec-covid: rationale and structure of an information retrieval shared task for covid-19. *Journal of the American Medical Informatics Association*, 27(9):1431–1436.
- Stephen Robertson and Hugo Zaragoza. 2009. *The probabilistic relevance framework: BM25 and beyond*. Now Publishers Inc.
- Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. 1995. Okapi at trec-3. *Nist Special Publication Sp*, 109:109.
- Gautam Kishore Shahi, Anne Dirkson, and Tim A Majchrzak. 2021. An exploratory study of covid-19 misinformation on twitter. *Online social networks and media*, 22:100104.
- Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 2020. Fakenewsnet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big data*, 8(3):171–188.
- Amir Soleimani, Christof Monz, and Marcel Worring. 2020. Bert for evidence retrieval and claim verification. *Advances in Information Retrieval*, 12036:359.
- Dominik Stambach and Guenter Neumann. 2019. Team domlin: Exploiting evidence enhancement for the fever shared task. In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*, pages 105–109.
- Motoki Taniguchi, Tomoki Taniguchi, Takumi Takahashi, Yasuhide Miura, and Tomoko Ohkuma. 2018. Integrating entity linking and evidence ranking for fact extraction and verification. In *Proceedings of the First Workshop on Fact Extraction and Verification (FEVER)*, pages 124–126.
- The Lancet COVID-19 content collection. 2020. The lancet covid-19 content collection. <https://www.thelancet.com/coronavirus/collection>.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018a. FEVER: a large-scale dataset for fact extraction and VERification. In *NAACL-HLT*.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2018b. The FEVER2.0 shared task. In *Proceedings of the Second Workshop on Fact Extraction and VERification (FEVER)*.
- Ellen Voorhees, Tasmeeer Alam, Steven Bedrick, Dina Demner-Fushman, William R Hersh, Kyle Lo, Kirk Roberts, Ian Soboroff, and Lucy Lu Wang. 2021. Trec-covid: constructing a pandemic information retrieval test collection. In *ACM SIGIR Forum*, volume 54, pages 1–12. ACM New York, NY, USA.
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or fiction: Verifying scientific claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7534–7550.
- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Doug Burdick, Darrin Eide, Kathryn Funk, Yannis Katsis, Rodney Michael Kinney, Yunyao Li, Ziyang Liu, William Merrill, Paul Mooney, Dewey A. Murdick, Devvret Rishi, Jerry Sheehan, Zhihong Shen, Brandon Stilson, Alex D. Wade, Kuansan Wang, Nancy Xin Ru Wang, Christopher Wilhelm, Boya Xie, Douglas M. Raymond, Daniel S. Weld, Oren Etzioni, and Sebastian Kohlmeier. 2020a. **CORD-19: The COVID-19 open research dataset**. In *Proceedings of the 1st Workshop on NLP for COVID-19 at ACL 2020*, Online. Association for Computational Linguistics.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020b. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *arXiv preprint arXiv:2002.10957*.
- 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 (Volume 2: Short Papers)*, pages 422–426.
- Ralph Weischedel, Martha Palmer, Mitchell Marcus, Eduard Hovy, Sameer Pradhan, Lance Ramshaw, Nianwen Xue, Ann Taylor, Jeff Kaufman, Michelle Franchini, et al. 2013. Ontonotes release 5.0 ldc2013t19. *Linguistic Data Consortium, Philadelphia, PA*, 23.
- WHO database of publications on coronavirus. 2020. Who database of publications on coronavirus. <https://search.bvsalud.org/global-literature-on-novel-coronavirus-2019-ncov/>.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- Wei Yang, Kuang Lu, Peilin Yang, and Jimmy Lin. 2019. Critically examining the “neural hype” weak baselines and the additivity of effectiveness gains from neural ranking models. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, pages 1129–1132.
- Takuma Yoneda, Jeff Mitchell, Johannes Welbl, Pontus Stenertorp, and Sebastian Riedel. 2018. Ucl machine reading group: Four factor framework for fact finding

- (hexaf). In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 97–102.
- Xia Zeng, Amani S Abumansour, and Arkaitz Zubiaga. 2021. Automated fact-checking: A survey. *Language and Linguistics Compass*, 15(10):e12438.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.
- Wanjun Zhong, Jingjing Xu, Duyu Tang, Zenan Xu, Nan Duan, Ming Zhou, Jiahai Wang, and Jian Yin. 2020. Reasoning over semantic-level graph for fact checking. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6170–6180.
- Jie Zhou, Xu Han, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2019. Gear: Graph-based evidence aggregating and reasoning for fact verification. In *The 57th annual meeting of the Association for Computational Linguistics (ACL)*.
- 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 & Knowledge Management*, pages 3205–3212.

A Data Sources

Here we present detailed information of the data sources introduced in section 3.1.

It is worth noting that for the construction of our dataset, we have only included sources or datasets that contain explicit veracity labels of specific claims, thus we have not included collections of tweets related to COVID that do not have veracity labels (Chen et al., 2020; Lamsal, 2021; Abdul-Mageed et al., 2021; Huang et al., 2020; Dimitrov et al., 2020; Kerchner and Wrubel, 2020; Qazi et al., 2020). We have not included claims without independent fact-checking sources (Memon and Carley, 2020; Shahi et al., 2021) and information sources without formulated claims such as the collections of scholarly articles (Wang et al., 2020a; Chen et al., 2021), news articles (Zhou et al., 2020), or articles obtained through specific repositories as (COVID-19 Data Portal, EU; WHO database of publications on coronavirus; Elsevier journals Novel Coronavirus Information Center; Cambridge journals Coronavirus Free Access Collection; The Lancet COVID-19 content collection; Oxford journals resources on COVID-19; MedRN medical research network SSRN Coronavirus Infectious Disease Research Hub).

The data sources that we have used for the construction of our dataset are:

- **The CoronaVirusFacts/DatosCoronaVirus Alliance Database**¹¹. Published by Poynter¹², this online publication combines fact-checking articles from more than 100 fact-checkers from all over the world, being the largest journalist fact-checking collaboration on the topic worldwide¹³. The publication is presented as an online portal, thus we had to develop scripts to crawl the content and extract the relevant claims, categories, and information sources.
- **CoAID dataset**¹⁴. The dataset (Cui and Lee, 2020) contains fake news from fact-checking websites and real news from health information websites, health clinics, and public institutions. Unlike most other datasets, it contains a wide selection of true claims.

¹¹<https://www.poynter.org/ifcn-covid-19-misinformation/>

¹²www.poynter.org

¹³<https://www.poynter.org/coronavirusfactsalliance/>

¹⁴<https://github.com/cuilimeng/CoAID>

- **MM-COVID**¹⁵. The multilingual dataset (Li et al., 2020) contains fake and true news collected from Poynter and Snopes¹⁶, being a good complement to the first data source.
- **CovidLies dataset**¹⁷. The dataset (Hossain et al., 2020) contains a curated list of common misconceptions about COVID appearing in social media, carefully reviewed to contain very relevant and unique claims unlike other automatically collected datasets.
- **TREC Health Misinformation track**¹⁸. Research challenge using claims on the health domain focused on information retrieval from general websites through the Common Crawl corpus¹⁹. This dataset is specialized in a very specific domain, and has been used for a very different application than the previous data sources.
- **TREC COVID challenge**²⁰. Research challenge (Voorhees et al., 2021; Roberts et al., 2020) using claims on the health domain focused on information retrieval from scholarly peer-reviewed journals through the COVID19 dataset (Wang et al., 2020a), the largest existing compilation of such articles. Similar to the last source, but focused on scholarly papers unlike the other sources.

A.1 Pre-processing

A separate pre-processing step was carried out for each of the selected data sources:

The CoronaVirusFacts/CoronaVirus Alliance Database. The data was downloaded on 13 February 2021. From the 11,647 entries initially obtained, entries with no fact-checking source and categories with less than 10 entries were removed. The different fact-checkers used different categories to label the claims, although in most of the cases the difference was mainly in terms of spelling. Initially we identified the following common categories: *False* (including FALSE, FALSO, Fake, false, false and misleading, Two Pinocchios, Misinformation

¹⁵<https://github.com/bigheiniu/MM-COVID>

¹⁶www.snopes.com

¹⁷<https://github.com/ucinlp/covid19-data>

¹⁸<https://trec-health-misinfo.github.io/>

¹⁹<https://commoncrawl.org/>

²⁰<https://ir.nist.gov/covidSubmit/data.html>

/ Conspiracy theory, Not true, false headline, MANIPULATED, Unproven), *Misleading* (MISleading, MISLEADING, misLEADING, MiSLEADING, misleading, Misleading/False, Misleading), *Missing Context* (Missing context, Needs Context, missing context), *No Evidence* (NO EVIDENCE, No evidence, No Evidence), *Mostly False* (Mostly False, Mostly false, MOSTLY FALSE, mostly false, Mainly false), *Partially False* (Partially False, Partly false, Partially false, partially false, partly false), *Partially True* (PARTLY TRUE, Partially correct, Partially true, Partly true, HALF TRUE, HALF TRUTH, half true), and *Mostly True* (Mostly true, MOSTLY TRUE, mainly correct). Next, we conducted a manual inspection of the different categories. We found that the categories *Misleading*, *Missing Context*, *No Evidence*, *Mostly False*, and *Partially False* had no homogeneous and clear definition through the different fact-checking media. Each group contains claims fitting the definition mixed with claims that are simply false (e.g. of false claims under other labels: “Misleading: Only people from South Korea have Covid-19 antibodies”; “Partially False: The vaccines contain substances such as arsenic or uranium according to scientific studies”; “Mostly False: Pope Francis contracted coronavirus”). Therefore, we decided to not use these nuanced categories but group them in the general *False* category. Additionally, we removed the 25 claims from the categories *Partially True* and *Mostly True*, since they contained both *True* and *False* claims.

CoAID dataset. The datasets NewsRealCOVID-19, NewsFakeCOVID-19, and ClaimFakeCOVID-19 were selected. The additional available dataset contains claims already existing in other datasets, formulated in this case as questions, and thus was not included. The selected datasets contain *True* and *False* claims.

MM-COVID. The claims were obtained from the `English_news` part of the dataset since we are only interested in English claims. 3,409 claims were collected. Claims in other languages appeared in the file, therefore we did a language filtering using polyglot²¹. Additionally, claims without fact-checking sources were deleted. It contains *True* and *False* claims.

CovidLies dataset. The available claims have been manually revised by eliminating duplicates, resulting in a total of 62 misconception claims. It contains *False* claims.

²¹<https://github.com/aboSamoor/polyglot>

TREC Health Misinformation track. The claims used in the track were obtained and reformulated manually by us as affirmative claims (e.g., “*Can vitamin D cure COVID-19?*” was changed to “*Vitamin D cures COVID-19?*”) for consistency with the rest of the data sources and to allow claim veracity assessment. *True* and *False* claims are used.

TREC COVID challenge. The claims used in the challenge were obtained and reformulated manually by us as full sentences using the explanations related to each query (e.g., for a given query “*coronavirus immunity*”, and its explanation “*will SARS-CoV2 infected people develop immunity?*”, we form the following claim, “*coronavirus infected people develop immunity*”). *True* and *False* claims are used.

The above processed data sources were combined to provide 20,689 initial claims.

A.2 Claim Categorisation

The claims were analysed to identify types of claims that may be of particular interest, either for inclusion or exclusion depending on the type of analysis. The following types were identified: (1) Multimodal; (2) Social media references; (3) Claims including questions; (4) Claims including numerical content; (5) Named entities, including: PERSON – People, including fictional; ORGANIZATION – Companies, agencies, institutions, etc.; GPE – Countries, cities, states; FACILITY – Buildings, highways, etc. These entities have been detected using a RoBERTa base English model (Liu et al., 2019) trained on the OntoNotes Release 5.0 dataset (Weischedel et al., 2013) using Spacy²².

B Parameter Setting

In our veracity assessment experiments, the parameters of the initial RoBERTa models are frozen during the training. The inputs are padded and truncated to the longest sequence, and a ReLU function is used as the activation function for the hidden layer. The GCNConv outputs are padded to the longest graph size. The loss function used is cross-entropy. The size of the hidden layer is 50, the batch size is 30, and the training is performed for 100 epochs for NLI-SAN and its variants, and 200 epochs for NLI-graph and its variants. The optimizer used is AdamW (Loshchilov and Hutter, 2019) with $\beta_1 = 0.9$, $\beta_2 = 0.999$, a weight decay of 0.01, and a learning rate of 10^{-2} for NLI, 10^{-4} for NLI+Sent and NLI-SAN, and

²²<https://spacy.io/>

Claim	Description
<i>"Sugar causes a cytokine storm in the lungs that promotes COVID-19"</i>	Retrieved documents are relating COVID and its cytokine storm effects, but without the specific mention of sugar, which does not cause a cytokine storm.
<i>"Barron Trump had COVID-19, Melania Trump says"</i>	Retrieved sentences such as <i>"Rudy Giuliani has tested positive for COVID-19, Trump says."</i> with a similar structure and mentions but mistaking the family members and missing the key name.
<i>"Prince Charles tested positive for COVID-19 after meeting Bollywood singer Kanika Kapoor."</i>	Documents mentioning Prince Charles positive COVID tests are obtained, but without any mentions to the singer.
<i>"Vice President of Bharat Biotech got a shot of the indigenous COVAXIN vaccine"</i>	Correct documents on the issue are retrieved. Similar sentences are retrieved such as <i>"Covaxin which is being developed by Bharat Biotech is the only indigenous vaccine that is approved for emergency use."</i> or <i>"Bharat Biotech's Covaxin is the first Indian vaccine to receive approval to conduct Phase I/Phase II trials."</i> . However, being similar they give no information about the claimed situation. In the retrieved document, the sentence <i>"The pharmaceutical company, has in a statement, denied the claim and said the image shows a routine blood test."</i> contains the essential information to debunk the original claim. But it is missed by the sentence retrieval engine as it is very different from the claim.
<i>"Masks can be sanitized in microwave"</i>	Correct documents are retrieved with similar sentences such as <i>"Claiming masks can be sanitized in microwave resurfaces"</i> . However, sentences such as <i>"The study authors cautioned health care workers against trying to clean masks this way. Microwaves melted the masks, making them useless."</i> or <i>"He also warns people against using microwaves or ovens to heat their masks."</i> that are present in the retrieved documents but are not similar enough to the claim are missed.

Table A1: Examples of errors in document or sentence retrieval.

a learning rates of 10^{-4} for NLI-graph, 10^{-3} for NLI-graph_{abl}, and 10^{-5} for NLI+PSent, these last three with a step size of 0.1 after 100 epochs.

C Additional Examples of Document or Sentence Retrieval Errors

Here we expand on the examples mentioned in Section 5.3 related to difficulties in the document or sentence retrieval parts of the process. Table A1 presents in more detail the cases previously mentioned, and includes new examples.