

CoVERT: A Corpus of Fact-checked Biomedical COVID-19 Tweets

Isabelle Mohr, Amelie Wührl, Roman Klinger

Institut für Maschinelle Sprachverarbeitung, University of Stuttgart, Germany
{isabelle.mohr, amelie.wuehrl, roman.klinger}@ims.uni-stuttgart.de

Abstract

Over the course of the COVID-19 pandemic, large volumes of biomedical information concerning this new disease have been published on social media. Some of this information can pose a real danger to people’s health, particularly when false information is shared, for instance recommendations on how to treat diseases without professional medical advice. Therefore, automatic fact-checking resources and systems developed specifically for the medical domain are crucial. While existing fact-checking resources cover COVID-19-related information in news or quantify the amount of misinformation in tweets, there is no dataset providing fact-checked COVID-19-related Twitter posts with detailed annotations for biomedical entities, relations and relevant evidence. We contribute CoVERT, a fact-checked corpus of tweets with a focus on the domain of biomedicine and COVID-19-related (mis)information. The corpus consists of 300 tweets, each annotated with medical named entities and relations. We employ a novel crowdsourcing methodology to annotate all tweets with fact-checking labels and supporting evidence, which crowdworkers search for online. This methodology results in moderate inter-annotator agreement. Furthermore, we use the retrieved evidence extracts as part of a fact-checking pipeline, finding that the real-world evidence is more useful than the knowledge indirectly available in pretrained language models.

Keywords: fact-checking, bioNLP, social media mining, COVID-19

1. Introduction

The COVID-19 pandemic has elicited a global response in the scientific community, leading to a burst of new information regarding the pathophysiology of the virus, new treatments for infected patients as well as vaccines in production (Chahrour et al., 2020). At the same time, large volumes of information about COVID-19 have also been published on social media platforms (Yang et al., 2020; Kouzy et al., 2020; Singh et al., 2020; Shahi et al., 2021), which are known to reach large networks of individuals in short amounts of time (Kouzy et al., 2020). This reach is paramount to curbing the spread of a virus like SARS-CoV-2. Although much of the information circulating online is potentially useful, some of it, in particular misinformation and disinformation, can also pose a great danger (Soltaninejad, 2020). Incorrect health advice can not only put individuals at risk, false and misleading information may also be detrimental to efforts in controlling the pandemic. Therefore, the stream of information on social media sites requires critical thought and fact-checking.

Manual fact-checking, however, is time-consuming and expensive. One way to perform the task automatically is to use machine learning to classify a claim to be true or false. An intermediate step to exploit potential evidence is to automatically decide for claim-evidence pairs if the evidence supports or refutes the claim (or whether that information is insufficient to rule a verdict) (Vlachos and Riedel, 2014). The growing need for fact-checking of COVID-19-related information has been met with releases of corpora of fact-checked COVID-19 claims (Shahi and Nandini, 2020; Shahi et al., 2021), and databases containing known facts concerning COVID-19 and

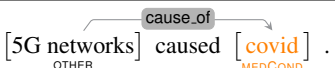
Variable	Values
NER Tweet	
Verdict	REFUTES
URL	https://www.muhealth.org/...
Evidence	There are two types of conspiracy associated with 5G-COVID-19. One version suggests that radiation from 5G lowers your immune system, which makes you more susceptible to the virus (Shultz, 2020). The idea that ...

Table 1: Example instance from the CoVERT corpus, with NER and crowdsourced fact-checking annotations, including textual evidence and the source URL.

related information (Reese et al., 2021; Domingo-Fernández et al., 2020). However, to the best of our knowledge, there is no resource that specifically addresses the truthfulness of biomedical information relating to COVID-19 circulating on Twitter.

We fill this gap by creating a corpus of fact-checked biomedical tweets, the Covid VERified Tweet (CoVERT) corpus¹, to facilitate the fact-checking task in this domain. The dataset consists of 300 tweets containing real-world claims about COVID-19. We annotate the tweets with two types of information: First, we annotate each tweet with the medical entities *Medical condition*, *Treatment*, *Symptom/Side-effect*, *Other* and annotate how they

¹The CoVERT corpus and supplementary material is available at <https://www.ims.uni-stuttgart.de/data/bioclaim>.

relate to each other (*cause_of*, *causative_agent_of*, *not_cause_of*). Secondly, crowdworkers verify the medical claim within the tweet. To the best of our knowledge, employing crowdworkers to perform biomedical fact-checking on Twitter posts has not been done before, though crowdworkers have completed subtasks for verifying COVID-19-related Reddit posts. Three annotators research a given claim and provide substantiating evidence from the web along with their verdict (SUPPORTS, REFUTES and NOT ENOUGH INFORMATION (NEI)). Table 1 shows an example instance from our dataset.

In an exploratory analysis of this dataset we find our novel methodology of crowdsourced fact-checking to be effective for this task, with moderate agreement between annotators on their verdicts. Finally, we explore to which extent textual evidence extracts provided by the annotators help inform fact-checking systems when making a prediction. We find that real evidence provides our pipeline with more useful information than what is available implicitly in the pretrained language model BERT (Devlin et al., 2019). However, the NOT ENOUGH INFORMATION class inherently has no substantiating evidence, making this class a challenge for this approach.

2. Related Work

In recent years, automated fact-checking has increasingly come into focus (Thorne and Vlachos, 2018). Systems can be grouped into approaches with and without access to external evidence. The former need to combine information from other texts or structured resources with the claim, while the latter rely on linguistic patterns that signal false information.

2.1. Fact-Checking without External Evidence

Previous research that automatically check claims without evidence make use of, i.a. emotion patterns (Giahanou et al., 2019) or surface-level linguistic properties (Rashkin et al., 2017). These can be categorized as stylometric approaches (Schuster et al., 2020). Based on the finding that language models capture relational information contained within the data they are trained on, Lee et al. (2020) include language models in the fact-checking pipeline itself. Tapping into the implicit knowledge aggregated over very large pretraining datasets, they use BERT (Devlin et al., 2019) to create an evidence text by unmasking an entity in the original claim. This is passed to an entailment model that predicts whether the language model’s evidence supports or refutes the claim.

Both the stylometric and the language model-based approaches are limiting. The first group fails to recognize well-presented false information that do not share the surface level characteristics of false news. When using a language model to extract evidence, the models also likely propagate biases learned from the

training data (Guo et al., 2022). Additionally, such models are biased towards particular prompts which are used to query and extract the implicit knowledge from the model (Sung et al., 2021).

2.2. Evidence-based Approaches

Both lines of research described in Section 2.1 are in contrast with the way journalists approach fact-checking, who rather seek substantiating external evidence. To go beyond the claim itself when making a verdict prediction, Vlachos and Riedel (2015) and Ciampaglia et al. (2015) use the structured nature of information available in knowledge bases, while others retrieve evidence extracts from resources like news articles and websites online (Popat et al., 2018; Hanselowski et al., 2019). This allows to fact-check claims with probabilistic estimates (Vlachos and Riedel, 2015) or to use a knowledge graph as a topology to predict how likely a claim is to be true (Ciampaglia et al., 2015).

In an effort to mirror the traditional process of journalistic fact-checking, Popat et al. (2018) retrieve substantiating articles that provide evidence for or against a claim. Further, they incorporate a credibility assessment of the claim’s source in their model.

In the related task of stance detection, Ferreira and Vlachos (2016) use news headlines as evidence for claims. Other approaches extract the summaries of articles rather than their headlines as evidence for claims (Hanselowski et al., 2019; Alhindi et al., 2018), filtering out sentences that are irrelevant in order to create more fine-grained substantiating evidence. Similarly, Wadden et al. (2021) make use of the evidence contained within the abstracts of research papers in order to fact-check scientific claims.

Retrieving evidence allows systems to make informed decisions based on external sources which helps overcome the limitations claim-focused models face. In addition, evidence can also provide the end-user with an explanation for the fact-checking verdict (Kotonya and Toni, 2020).

2.3. Fact-Checking Corpora

Recently, the task of fact-checking has received much attention (Thorne and Vlachos, 2018), resulting in an influx of published datasets. The Fact Extraction and VERification (FEVER) (Thorne et al., 2018) dataset is one of the largest available corpora, consisting of 185,445 claims generated from Wikipedia. Claims are manually annotated (labels: SUPPORTED, REFUTED, NOT ENOUGH INFO) and accompanied by an evidence sentence from the same original Wikipedia article. In the political domain, previously fact-checked claims from *PolitiFact* and *Channel4* have been collected (Vlachos and Riedel, 2014; Wang, 2017).

The COVID-19 pandemic has demonstrated the increasing need for fact-checking in the medical domain. The *FakeCovid* dataset (Shahi and Nandini, 2020) consists of 5,182 fact-checked news articles

mentioning COVID-19, collected from various fact-checking websites. This dataset however lacks substantiating evidence for claims. Kouzy et al. (2020) collect 673 COVID-19 related tweets of which 24.8 % contained misinformation. *COVID-Fact* (Saakyan et al., 2021) and *HealthVer* (Sarrouiti et al., 2021) both address the limitation of existing corpora consisting of synthetic or manually summarized claims. Both datasets provide user-generated Covid-19 related claims. *COVID-Fact* focuses on Reddit while *HealthVer* claims stem from excerpts that a search engine produces when queried with questions about Covid-19.

While most resources rely on existing fact-checks or expert annotators to label claims, few have explored crowdsourcing to alleviate the data bottleneck. Saakyan et al. (2021) present crowdworkers with a claim and 5 automatically collected sentences, from which supporting evidence should be selected, thereby adjudicating a fact-checking verdict. Further reducing the work of a crowdworker, Hanselowski et al. (2019) task crowdworkers with refining previously fact-checked claims from *Snopes* by marking a specific snippet as evidence within a given source text. Recently, Allen et al. (2021) successfully employed crowds in providing truthfulness ratings for headline and lede sentence pairs from suspicious articles on Facebook. They find that aggregated crowd judgments strongly correlate with judgments made by professional fact-checkers (Allen et al., 2021).

To the best of our knowledge there is no existing resource that facilitates fact-checking user-generated biomedical claims related to COVID-19 on Twitter. To explore this task, we focus specifically on creating a corpus and leverage the crowd in making veracity assessments and retrieving substantiating evidence to support their verdict.

3. Corpus Creation and Annotation

3.1. Data Collection and Preprocessing

We sample tweets from an in-house tweet repository based on medical terms from MeSH² and frequently occurring terms. We select data posted between January 2020 and June 2021 containing a mention of COVID-19 and one of the terms ‘effect’, ‘side-effect’, ‘vaccine’, ‘symptom’ or ‘treatment’. Further, we select only tweets that also contain the lexeme ‘caus’ (matching ‘causes’, ‘caused’, etc.) to narrow the search space to ‘causal’ relations.³ This results in a total sample of 38,251 COVID-19-related tweets.

As this set still contains many non-biomedical tweets (e.g. tweets discussing politics around COVID-19), we select a random starting set of 1118 tweets

²<https://www.nlm.nih.gov/mesh/meshhome.html>

³We use a boolean expression with inexact matching: (COVID OR corona) AND (effect OR side-effect OR vaccine OR symptom OR treatment) AND (caus*)




Figure 1: Excerpts of tweets that contain named entity and relation annotations.

and manually label them as either ‘biomedical’ or ‘non-biomedical’, resulting in 408 and 710 per class respectively. We train a bag-of-words-based feed-forward neural network⁴ to predict this binary class for the remaining tweets. We further filter tweets unlikely to contain a claim using the claim detection model for tweets from Wühl and Klinger (2021). After removing duplicates and retweets, we are left with a set of 3,785 biomedical tweets containing claims with causal relations, from which we randomly sample 300 tweets for our annotation tasks.

3.2. Annotation

3.2.1. Entity and Relation Annotation

We annotate entities and relations in the biomedical claim tweets for the downstream verification task. We perform manual annotations supported by the scispacy model for biomedical NER⁵ (Neumann et al., 2019). This model assigns a generic ‘Entity’ label which we manually categorize as one of the following types:

Medical Condition: Mentions of diseases, illnesses, ailments or disorders.

Treatment: Medical care given to patients for a disease or illness.

Symptom/Side-effect: Secondary physical effects of a medical treatment or condition.

Other: Relevant entities that do not fall into the above mentioned categories.

We discard non-biomedical entities that the model incorrectly identified and only keep entities that are directly relevant to the causal phrase in the tweet claim. Figure 1 shows examples of annotated tweet claims. To test the quality of this manual revision, a second annotator labels a random subset of 100 tweets. Both annotators are female, with a background in (computational) linguistics, aged 27 and 28 years old, respectively.⁶

The same annotator who assigned the entity labels further annotates each instance with relations. We use the relations *cause_of* (UR214) and *causative_agent_of* (UR173) from UMLS, a medical terminology,

⁴Implementation details can be found in the Appendix, Table 8.

⁵See `en_core_sci_lg` from <https://github.com/allenai/scispacy>.

⁶We provide the annotation guidelines and example fact-checking annotation environment together with the corpus.

classification and coding system (Bodenreider, 2004), as well as the relation *not_cause_of* (UR214!). The entity and relation annotations are relevant for creating prompts for language models.

3.2.2. Fact-checking Annotation

Annotation Environment. We use *Google Forms*⁷ to collect fact-checking verdicts and substantiating evidence for the tweet claims and recruit annotators via *Prolific*⁸. We present 10 tweets with instructions as one annotation task (An example annotation environment for one tweet can be found in the Appendix). For each tweet, we ask annotators to search for the claim’s key terms using Google Search to find supporting or refuting sources. Sources should be credible and reputable (extensions ‘.gov’ and ‘.mil’ and general news, medical and scientific articles), and avoid resources that are known to be satirical/comedic like *The Onion*⁹. If no resource could be found after searching for two minutes, annotators labeled the claim as non-verifiable.

If relevant evidence is found, annotators state whether this evidence *supports* or *refutes* the claim. This leads to three labels for claim-evidence pairs, answering the question “Could you find a resource that confirms or refutes the claim?”:

SUPPORTS. Yes, the resource CONFIRMS this claim.

REFUTES. Yes, the resource REFUTES this claim.

NEI. No, I could not find any resource that confirms or refutes this claim.

Annotators are tasked to provide the URL together with the relevant evidence text snippet.

To ensure that participants follow instructions carefully, each annotation task consisting of 10 tweets includes an attention check and reject submissions that fail this check.

Crowdsourcing Platform and Payment. We use Prolific to recruit annotators. We filter them to meet the following criteria: currently living in either the UK, US, Ireland or Germany, between 18 and 45 years old, English first language speaker, have achieved at least an undergraduate degree (BA/BSc/other) in a subject related to Biomedicine, Biochemistry, Biology or Medicine, and do not have any literacy difficulties. Each tweet is annotated by three annotators. The recommended time of completion is 20 minutes. Each annotator is paid £7.50 per hour, which corresponds to £0.25 for each tweet. To increase awareness of our task, we implement a bonus system. We manually assess all submission and award annotators with a bonus of £0.50 if the annotations are of good quality, i.e., at least 8 out of 10 annotations are coherent and the given evidence substantiates the fact-checking label. The expenses of this study amount to £314.52.

Entity Category	IA F ₁ Score
Treatment	0.95
Symptom/Side-effect	0.68
Medical Condition	0.94
Other	0.93
Weighted Macro Average	0.88

Table 2: Inter-annotator (IA) agreement F₁ scores on the named entities of 100 tweets.

4. Corpus Analysis

4.1. Annotator Statistics

78 annotators participated in the fact-checking annotation. All of the crowdsourced annotators were students at the time of participation. Even though this was not a strict requirement, 78.2% have completed an undergraduate and 19.2% a graduate degree. Most annotators (96.1%) are between the age of 19 and 30, and 87.2% of annotators identified as female, 12.8% as male. The average time taken to complete the annotation task (a set of 10 tweets) was 25:38 minutes, which is slightly longer than was intended by the instructions. We presume that participants actually spent 2 minutes verifying each tweet, and needed additional time to read and understand instructions.

4.2. Annotator Agreement and Adjudication

4.2.1. Entity and Relation Annotation

Table 2 presents the inter-annotator (IA) F₁ scores achieved between the two annotators manually revising the automatic named entity annotations. We consider one annotation as gold annotations and treat the other annotator’s labels as predictions (Hripcsak and Rothschild, 2005). The agreement between the two annotators is high, with inter-annotator F₁ scores above 0.9 in all but one category (Symptom/Side-effect). The macro-average (weighting by number of samples per category) is 0.88 F₁. Note that the annotation task did not include segmentation.

4.2.2. Crowdsourced Fact-checking

Verification Task. We calculate Cohen’s κ per annotation task (set of 10 tweets), where each task has three annotators (thus, three pair-wise comparisons) and take the average of all pairs over all tasks (a total of 30) as a final agreement score. The average Cohen’s κ on the verification labels is 0.44 indicating moderate agreement. However, to adjudicate the final label for each tweet, we select the majority decision (where at least 2 annotators agree) as the gold label. Out of all tweets, only in two cases did annotators not agree at all on the verification labels. Here, we assign the label NOT ENOUGH INFO. For the remaining tweets, at least two annotators agree on the verification label.

Most disagreements (55) are between the pair of labels SUPPORTS and NOT ENOUGH INFO. There are only 42 disagreements between the pair SUPPORTS and

⁷<https://www.google.com/forms/about/>

⁸<https://prolific.co/>

⁹<https://www.theonion.com/>

Entity / Relation Class	# Instances
Medical Condition	424
Treatment	102
Other	102
Symptom/Side-effect	94
cause_of	253
not_cause_of	30
causative_agent_of	17

Table 3: Instances per entity and relation class in the CoVERT dataset.

REFUTES. The pair of labels with fewest disagreements are REFUTES and NEI (29). This means two annotators are most likely to disagree on an instance where one annotator labels the tweet as SUPPORTS while another annotates it as NEI, while disagreements on REFUTES and NEI are not as common.

Evidence Retrieval. To gauge how well the crowdworkers agreed on the evidence they used to substantiate their verdicts, we compare the URLs they provided during the annotation process. Out of all 300 tweets, for 78 tweets (26% of total) two or more annotators responded with a link to the same resource. Notably, even when a link to the same resource was provided by two different annotators, we observe 5 cases where annotators disagree on the interpretation of the resource in supporting or refuting the claim. For instance, the claim “Covid-19 can cause hearing impairment, tinnitus and dizziness” for which annotators provided the same evidence URL (from <https://www.healthyhearing.com/>) leads to disagreement regarding SUPPORTS vs. REFUTES verdicts.

4.3. Corpus Statistics

4.3.1. Entities and Relations

The resulting CoVERT corpus has a total of 722 entity annotations, with the number of instances per entity and relation class enumerated in Table 3. The largest observed entity class is ‘Medical Condition’, while the smallest is ‘Symptom/Side-effect’. As for the relations, we most frequently observe the *cause_of* relation.

4.3.2. Crowdsourced Fact-checking

Verification Task. The distribution of labels in the corpus after adjudication is as follows: 198 instances of SUPPORTS, 66 instances of REFUTES and 36 instances of NEI. The SUPPORTS class has the largest number of examples, while the NEI class has the least.

Evidence Retrieval. We collect a total of 659 URLs from the annotation task. Figure 2 displays the 20 most frequently referenced domain names and their counts, which make up 66% of all mentioned URLs. 188 domain names are unique, but 150 of these domains occur less than 5 times in the collected data.




Figure 2: Top 20 most frequent domain names occurring in URL links to evidence resources.

The most frequently mentioned domain names are medical and health science related domains. Most of these are also generally deemed reputable and credible sources of information such as the CDC (Center for Disease Control and Prevention) and NIH (National Institutes of Health). This assures the annotators were following the annotation instructions carefully, leading to a good quality of annotations.

5. Experiments

We want to investigate the extent to which access to external evidence impacts the prediction of fact-checking verdicts for COVID-19-related claims from Twitter. To explore this, we employ a fact-checking pipeline proposed by Lee et al. (2020). Their approach does not access external evidence. Instead it extracts evidence from the implicit knowledge contained in the language model BERT (Devlin et al., 2019). We adapt this system such that it can access external fact-checking evidence (i.e., evidence annotators extracted for CoVERT) and compare the verdict prediction performance.

Before investigating this as our primary research objective, we perform two preliminary experiments. As the claims within the CoVERT data are Twitter-based instead of Wikipedia-based (FEVER dataset (Thorne et al., 2018)), we first explore to which extent BERT and BioBERT (Lee et al., 2019) contain domain-specific knowledge discussed in the CoVERT dataset. For this, we probe each model with the BioLAMA probe (Sung et al., 2021). In addition, we want to understand how capable BERT is to “generate” evidence sentences as suggested by Lee et al. (2020) in our setting to be used in the fact-checking pipeline. We therefore analyze the predictions BERT makes when unmasking entities in the tweets from our dataset. We outline the methodology and results for both preliminary experiments and the fact-checking pipeline in the following section.

5.1. Methods

5.1.1. Probing Language Models for Domain-specific Knowledge

To gauge how well implicit knowledge is stored in BERT and BioBERT, and how well these language models lend themselves as a source of evidence like Lee et al. (2020) suggest, we investigate to which extent information from the CoVERT dataset is contained within these two language models. We employ the BioLAMA probe suggested by Sung et al. (2021) which allows us to see whether a language model is able to correctly predict masked object entities in a constrained setting. BioLAMA originally probes BERT and BioBERT for their inherent relational knowledge using factual triples sourced from CTD, UMLS and Wikidata. Comparing probing results for CoVERT data to those attained using the original dataset of biomedical factual triples in Sung et al. (2021) gives insights into the domain of biomedical tweets and allows us to conclude which language model best captures the domain-specific knowledge within CoVERT.

BioLAMA generates fill-in-the-blank cloze statements or ‘prompts’ like “Hepatitis has symptoms such as [Y].”, where [Y] is the masked object, unmasked as ‘abdominal pain’ by BioBERT. In our experiment, we use the manual approach for generating prompts as described in Sung et al. (2021). We follow their evaluation methodology and use top- k accuracy as the evaluation metric. This is equal to 1 for an instance if any of the top- k object entities match an object in the gold annotated object list. If there are no matches, the score for this instance is 0. This binary setting allows to calculate accuracy as the number of correct predictions divided by all predictions.

5.1.2. Language Model Capacity for Unmasking Claim Entities

In addition to probing the knowledge within the language model, we further want to understand its capacity to create ‘evidence texts’ for modelling the fact-checking task without access to external evidence. Therefore, we conduct an analysis of the predictions BERT makes when masking object entities in the biomedical claim tweets. For each of the three fact-checking categories, we extract and analyze the probabilities with which the object entity is predicted. We report the probability as 0 if the correct entity does not appear in the top 1000 predictions.

5.1.3. Verdict Prediction With and Without External Evidence

We investigate whether verdict prediction benefits from access to external evidence, or whether this evidence can be replaced with unmasked claims from a language model and still achieve similar results. To do this, we re-implement the fact-checking pipeline by Lee et al. (2020). Figure 3 shows a diagram of the pipeline. The input to the pipeline consists of a $\langle \text{text}, \text{hypothesis} \rangle$ pair,

Hypotheses	Text	“Covid causes <u>heart inflammation</u> in like 70% of people (including asymptomatic). That seems like a high rate of heart inflammation.”
	LM	“Covid causes <u>inflammation</u> in like 70% of people (including asymptomatic). That seems like a high rate of heart inflammation.”
	Web	“The SARS-CoV-2 virus can damage the heart in several ways. For example, the virus may directly invade or inflame the heart muscle, and it may indirectly harm the heart by disrupting the balance between oxygen supply and demand.”

Table 4: Example instance for verdict prediction. The *text* is the original tweet, while the *hypothesis* is either the tweet unmasked by the language model (LM) or the evidence extract retrieved from the web. The underlined entity is unmasked by the language model.

which is passed to the textual entailment (TE) model from AllenNLP (Gardner et al., 2018). We take the last layer of the pretrained entailment model (before the softmax) to obtain “entailment features”, which are passed on to a multi-layer perceptron (MLP) for the final verification prediction. The MLP component is originally trained on the FEVER 2018 training set (Thorne et al., 2018), and is referred to as MLP-FEVER in our experiments.

For the *hypothesis* component of the input pair, we experiment with two types of inputs. First, we use the approach by Lee et al. (2020) and generate the *hypothesis* using the language model BERT (Devlin et al., 2019), which unmask an entity in the original text. Alternatively, we input the respective evidence snippet from the CoVERT dataset, effectively giving the pipeline access to external, real-world evidence. Refer to Table 4 for an example of such a $\langle \text{text}, \text{hypothesis} \rangle$ pair.

To further test the effect of evidence extracts on the pipeline, we additionally fine-tune the MLP component with text and evidence pairs, which we call MLP-Evidence in our experiments, using a 80/10/10 train-develop-test split of the CoVERT data.¹⁰

5.2. Results

5.2.1. BioLAMA Probing

We employ the BioLAMA probe to investigate whether BERT and BioBERT contain suitable domain-specific knowledge to serve as an evidence source during fact-checking. Table 5 reports results for the probe within the CoVERT dataset as well as the scores for the original BioLAMA probe (Sung et al., 2021). The

¹⁰Implementation details and hyperparameters of MLP-FEVER and MLP-Evidence can be found in the Appendix, Table 9.

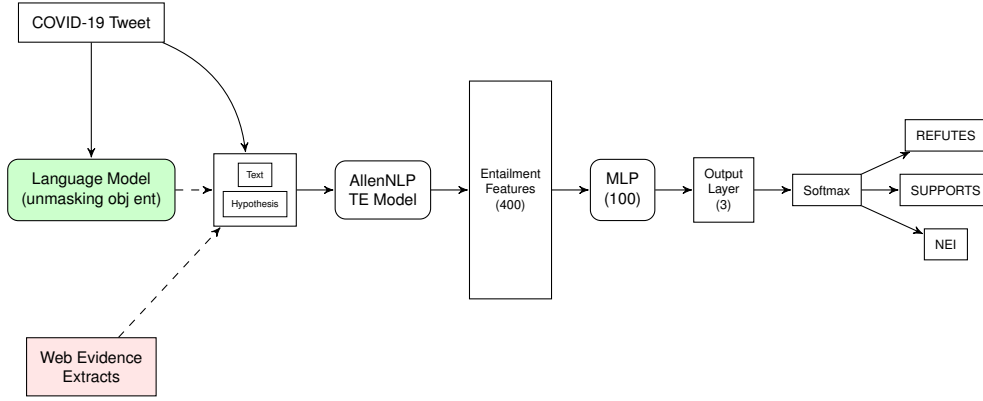


Figure 3: Full depiction of the fact-checking pipeline. ‘Hypothesis’ is either generated by the language model (green) or taken from crowdsourced annotations (red).

Dataset	BERT		BioBERT	
	Acc@1	Acc@5	Acc@1	Acc@5
SUPPORTS	4.6	10.79	1.59	5.56
REFUTES	3.62	6.57	2.17	5.0
NEI	0.0	7.07	0.0	25.18
CTD+UMLS +Wiki	0.86	3.08	1.75	6.09

Table 5: Accuracy scores (%) for BERT probed with BioLAMA using CoVERT (SUPPORTS, REFUTES and NEI) and the original BioLAMA probing collection (CTD, UMLS and Wikidata).

SUPPORTS class is most accurately modeled out of the three verdict classes (4.6 Acc@1), while the lowest score (0.0 Acc@1) is achieved for NEI.

Generally, BERT achieves higher scores on the CoVERT data than the original corpus of triples. The original BioLAMA probe using BioBERT on the CTD, UMLS and Wikidata collection achieved a score of 1.75 Acc@1 and 6.09 Acc@5. On the CoVERT corpus, BERT achieves higher scores than BioBERT, even though the corpus consists of biomedical information. We hypothesize that this is because the specialized biomedical language used in the BioBERT training data might not be a good match for the non-expert language commonly used in tweets. Since BERT appears better suited for our dataset, we use this language model in subsequent experiments.

5.2.2. Language Model Capacity for Unmasking Claim Entities

We further explore how capable the BERT language model is to unmask entities in biomedical claims. Table 6 reports the probabilities with which masked object entities are predicted by BERT. We report the results grouped by the fact-checking class (SUPPORTS, REFUTES and NEI) that each instance belongs to. Most entities are predicted with very low probability by BERT. There are differences between the three

Class	Min	Q1	Q2	Q3	Max
SUP.	0	0	0.00057	0.03258	0.98289
REF.	0	0	0	0.00210	0.98672
NEI	0	0	0.00003	0.00995	0.7943

Table 6: Five-number summary of probabilities with which masked object entities are predicted by BERT.

fact-checking categories, where SUPPORTS has a third quartile that is higher than that of REFUTES. Additionally, the NEI category has a maximum lower than that of SUPPORTS and REFUTES. However, the mean, median and first quartile of all three categories are all very close or equal to 0. The language model still predicts the object entity of incorrect claims with high probabilities. This indicates that predictions for this class have to be considered carefully when using it in the downstream task of evidence creation.

5.2.3. Verdict Prediction With and Without External Evidence

Our main research question is to investigate the impact of external evidence when predicting fact-checking verdicts for COVID-19-related claims in tweets. Table 7 presents the results from the verdict prediction experiment. We report the results for two prediction pipelines, namely MLP-FEVER (following Lee et al. (2020)) and our adaptation of this pipeline (MLP-Evidence) in which the MLP component is fine-tuned with evidence and text pairs from CoVERT, giving it access to external evidence. Each pipeline is evaluated on the FEVER 2018 test dataset (Thorne et al., 2018) consisting of 9999 instances, and on pairs of tweets and language model generated evidence (Tweet + LM) and tweets and CoVERT evidence (Tweet + Evidence).

Provided with hypotheses generated using BERT, MLP-FEVER achieves the highest F_1 score of 0.60. The same pipeline achieves an 0.49 F_1 on the FEVER dataset. When using the evidence extracts as

	FEVER			Tweet + LM			Tweet + Ev.		
	P	R	F ₁	P	R	F ₁	P	R	F ₁
MLP-FEVER	.52	.47	.49	.60	.61	.60	.61	.38	.46
MLP-Evidence	.64	.36	.46	.66	.68	.57	.68	.74	.69

Table 7: Results for the two fact-checking pipelines, MLP-FEVER and MLP-Evidence, evaluated on the FEVER 2018 dataset, Tweet + LM Pairs and Tweet + Evidence Pairs. The highest F₁ score is highlighted for each input set.

hypotheses, the performance drops to 0.46 F₁. Fine-tuning the MLP component with evidence extracts (MLP-Evidence) the performance is slightly lower on FEVER data and the Tweet + LM pairs (difference of .03 F₁, respectively), however, achieves substantially higher F₁ score of 0.69 on the CoVERT corpus.

The results show that fine-tuning the MLP component with evidence extracts, the pipeline achieves higher scores than with access to ‘evidence’ generated by BERT only, showing that the retrieved evidence extracts contribute to the performance of this pipeline.

5.3. Error Analysis and Discussion

To see how predictions change when the pipeline has access to language model generated evidence or CoVERT evidence extracts, we conduct a qualitative error analysis into predictions made by the fact-checking pipeline before and after fine-tuning.

We observe that MLP-FEVER with BERT-generated evidence mistakenly labels 4 out of 6 instances as REFUTES that are actually SUPPORTS instances. Inspecting these cases, we find that the language model generated evidence text is likely to produce features indicating *entailment*, as can be seen in Ex. *b*, Table 10 in the Appendix. Similarly, in 3 out of 15 SUPPORTS instances are mistakenly labelled REFUTES by MLP-FEVER, with example *c* in Table 10 showing that the evidence created by BERT is sufficiently dissimilar to result in *contradiction* entailment features in our pipeline, finally resulting in RERFUTES predictions. These instances show that BERT is creating evidence texts that are likely unuseful to the pipeline.

The pipeline struggles to make use of evidence extracts when the MLP component is only trained on the FEVER dataset. We inspect 10 instances where MLP-FEVER incorrectly predicts NEI instead of SUPPORTS and find that the evidence extracts are much longer (and sometimes a long paragraph), thus including more information than hypotheses originally seen in FEVER. Examples of this can be seen in Table 11, Ex. *d* and *f*. In both instances, the MLP-Evidence pipeline (fine-tuned with evidence from CoVERT) correctly predicts SUPPORTS, showing that the MLP component is able to interpret the input pairs correctly after fine-tuning. Although the overall F₁ score of the pipeline increases when fine-tuned with evidence extracts from CoVERT,

the total number of SUPPORTS class predictions has increased from 33 to 55, with 15 false positives, and very seldomly predicting the NEI class. This may be an inherent limitation of the evidence extracts, as they by definition do not contain any evidence texts for the NEI class, meaning this class is never modelled during fine-tuning. Additionally, the distribution across classes in the fine-tuning set is unbalanced (422 evidence extracts for SUPPORTS vs. 128 for REFUTES).

6. Conclusion and Future Work

We present CoVERT, the first fact-checked biomedical COVID-19-related tweet corpus, along with a novel approach to using evidence extracts as part of a verdict prediction pipeline. We outline how we leverage a crowd for fact-checking, finding moderate agreement among annotators and reliable annotations when aggregating verdict labels. Our extension of the verdict prediction pipeline (Lee et al., 2020) using evidence extracts from CoVERT indicates that the task can benefit from real-world evidence rather than only using ‘evidence’ generated from the implicit knowledge contained in language models.

Apart from facilitating evidence-based fact-checking directly, the CoVERT corpus additionally allows querying structured databases for evidence retrieval and verdict prediction, as it offers annotated entities and relations, which can be linked to ontologies.

Future work may consider if the pipeline benefits from more fine-grained evidence extracts. Evidence extracted at sentence level (Hanselowski et al., 2019) may contain information directly relevant to the claim, thereby facilitating richer entailment features.

7. Ethical Considerations

All participants in the crowdsourcing study agreed to participate at their own volition and signed a consent form at the outset of the study. Although we did not consider the task to present any harm to the crowd workers, it is possible that annotators were exposed to false information in the tweets and during their research.

It is important that the annotations in this dataset are not taken out of context with regard to the time-frame at which they were annotated. As biomedical knowledge, particularly with regard to SARS-COV-2, is continuously updated as new research is published, the evidence extracts and verdicts in the CoVERT dataset may be outdated in the future.

8. Acknowledgements

This research has been conducted as part of the FIBISS project which is funded by the German Research Council (DFG, project number: KL 2869/5-1). We thank Enrica Troiano for valuable discussions.

9. Bibliographical References

- Alhindi, T., Petridis, S., and Muresan, S. (2018). Where is Your Evidence: Improving Fact-checking by Justification Modeling. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 85–90, Brussels, Belgium, November. Association for Computational Linguistics.
- Allen, J., Arechar, A. A., Pennycook, G., and Rand, D. G. (2021). Scaling up fact-checking using the wisdom of crowds. *Science Advances*, 7(36):eabf4393.
- Bodenreider, O. (2004). The Unified Medical Language System (UMLS): integrating biomedical terminology. doi: 10.1093/nar/gkh061.
- Chahrour, M., Assi, S., Bejjani, M., Nasrallah, A. A., Salhab, H., Fares, M., and Khachfe, H. H. (2020). A Bibliometric Analysis of COVID-19 Research Activity: A Call for Increased Output. *Cureus*.
- Ciampaglia, G. L., Shiralkar, P., Rocha, L. M., Bollen, J., Menczer, F., and Flammini, A. (2015). Computational Fact Checking from Knowledge Networks. *PLOS ONE*, 10(6):1–13, 06.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Domingo-Fernández, D., Bakshi, S., Schultz, B., Gadiya, Y., Karki, R., Raschka, T., Ebeling, C., Hofmann-Apitius, M., and Kodumullil, A. T. (2020). COVID-19 Knowledge Graph: a computable, multi-modal, cause-and-effect knowledge model of COVID-19 pathophysiology. *Bioinformatics*, 09. btaa834.
- Ferreira, W. and Vlachos, A. (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, San Diego, California, June. Association for Computational Linguistics.
- Gardner, M., Grus, J., Neumann, M., Tafjord, O., Dasigi, P., Liu, N. F., Peters, M., Schmitz, M., and Zettlemoyer, L. (2018). AllenNLP: A deep semantic natural language processing platform. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia, July. Association for Computational Linguistics.
- Giahannou, A., Rosso, P., and Crestani, F. A. (2019). Leveraging Emotional Signals for Credibility Detection. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR’19*, page 877–880, New York, NY, USA. Association for Computing Machinery.
- Guo, Z., Schlichtkrull, M., and Vlachos, A. (2022). A Survey on Automated Fact-Checking. *Transactions of the Association for Computational Linguistics*, 10:178–206, 02.
- Hanselowski, A., Stab, C., Schulz, C., Li, Z., and Gurevych, I. (2019). A Richly Annotated Corpus for Different Tasks in Automated Fact-Checking. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 493–503, Hong Kong, China, November. Association for Computational Linguistics.
- Hripcsak, G. and Rothschild, A. S. (2005). Agreement, the f-measure, and reliability in information retrieval. *Journal of the American Medical Informatics Association : JAMIA*, 12(3):296–298.
- Kotonya, N. and Toni, F. (2020). Explainable automated fact-checking: A survey. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5430–5443, Barcelona, Spain (Online), December. International Committee on Computational Linguistics.
- Kouzy, R., Abi Jaoude, J., Kraittem, A., El Alam, M., Karam, B., Adib, E., Zarka, J., Traboulsi, C., Akl, E., and Baddour, K. (2020). Coronavirus Goes Viral: Quantifying the COVID-19 Misinformation Epidemic on Twitter. *Cureus*, 12, 03.
- Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C. H., and Kang, J. (2019). Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, Sep.
- Lee, N., Li, B. Z., Wang, S., Yih, W.-t., Ma, H., and Khabsa, M. (2020). Language Models as Fact Checkers? In *Proceedings of the Third Workshop on Fact Extraction and VERification (FEVER)*, pages 36–41, Online, July. Association for Computational Linguistics.
- Neumann, M., King, D., Beltagy, I., and Ammar, W. (2019). ScispaCy: Fast and robust models for biomedical natural language processing. In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 319–327, Florence, Italy, August. Association for Computational Linguistics.
- Popat, K., Mukherjee, S., Yates, A., and Weikum, G. (2018). DeClarE: Debunking Fake News and False Claims using Evidence-Aware Deep Learning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 22–32, Brussels, Belgium, October-November. Association for Computational Linguistics.
- Rashkin, H., Choi, E., Jang, J. Y., Volkova, S., and Choi, Y. (2017). Truth of Varying Shades: Analyzing Language in Fake News

- and Political Fact-Checking. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2931–2937, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Reese, J. T., Unni, D., Callahan, T. J., Cappelletti, L., Ravanmehr, V., Carbon, S., Shefchek, K. A., Good, B. M., Balhoff, J. P., Fontana, T., Blau, H., Matentzoglou, N., Harris, N. L., Munoz-Torres, M. C., Haendel, M. A., Robinson, P. N., Joachimiak, M. P., and Mungall, C. J. (2021). KG-COVID-19: A Framework to Produce Customized Knowledge Graphs for COVID-19 Response. *Patterns*, 2(1):100155.
- Saakyan, A., Chakrabarty, T., and Muresan, S. (2021). COVID-fact: Fact extraction and verification of real-world claims on COVID-19 pandemic. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2116–2129, Online, August. Association for Computational Linguistics.
- Sarrouti, M., Ben Abacha, A., Mrabet, Y., and Demner-Fushman, D. (2021). Evidence-based fact-checking of health-related claims. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3499–3512, Punta Cana, Dominican Republic, November. Association for Computational Linguistics.
- Schuster, T., Schuster, R., Shah, D. J., and Barzilay, R. (2020). The Limitations of Stylometry for Detecting Machine-Generated Fake News. *Computational Linguistics*, 46(2):499–510, 06.
- Shahi, G. K. and Nandini, D. (2020). FakeCovid - A Multilingual Cross-domain Fact Check News Dataset for COVID-19. In *CySoc 2020 International Workshop on Cyber Social Threats, ICWSM 2020*.
- Shahi, G. K., Dirkson, A., and Majchrzak, T. A. (2021). An exploratory study of COVID-19 misinformation on twitter. *Online Social Networks and Media*, 22.
- Singh, L., Bansal, S., Bode, L., Budak, C., Chi, G., Kawintiranon, K., Padden, C., Vanarsdall, R., Vraga, E., and Wang, Y. (2020). A first look at covid-19 information and misinformation sharing on twitter. ArXiv preprint at <https://arxiv.org/abs/2003.13907>.
- Soltaninejad, K. (2020). Methanol Mass Poisoning Outbreak: A Consequence of COVID-19 Pandemic and Misleading Messages on Social Media. *Int J Occup Environ Med*, 11(3):148–150, 07.
- Sung, M., Lee, J., Yi, S., Jeon, M., Kim, S., and Kang, J. (2021). Can Language Models be Biomedical Knowledge Bases? In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4723–4734, Online and Punta Cana, Dominican Republic, November. Association for Computational Linguistics.
- Thorne, J. and Vlachos, A. (2018). Automated Fact Checking: Task Formulations, Methods and Future Directions. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3346–3359, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Thorne, J., Vlachos, A., Cocarascu, O., Christodoulopoulos, C., and Mittal, A. (2018). The fact extraction and VERification (FEVER) shared task. In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 1–9, Brussels, Belgium, November. Association for Computational Linguistics.
- Vlachos, A. and Riedel, S. (2014). Fact Checking: Task definition and dataset construction. In *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*, pages 18–22, Baltimore, MD, USA, June. Association for Computational Linguistics.
- Vlachos, A. and Riedel, S. (2015). Identification and verification of simple claims about statistical properties. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2596–2601, Lisbon, Portugal, September. Association for Computational Linguistics.
- Wadden, D., Lo, K., Wang, L. L., Cohan, A., Beltagy, I., and Hajishirzi, H. (2021). Longchecker: Improving scientific claim verification by modeling full-abstract context. ArXiv preprint at <https://arxiv.org/abs/2112.01640>.
- Wang, W. Y. (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, Vancouver, Canada, July. Association for Computational Linguistics.
- Wühlr, A. and Klinger, R. (2021). Claim detection in biomedical Twitter posts. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 131–142, Online, June. Association for Computational Linguistics.
- Yang, K., Torres-Lugo, C., and Menczer, F. (2020). Prevalence of low-credibility information on twitter during the COVID-19 outbreak. In *Workshop Proceedings of the 14th International AAAI Conference on Web and Social Media*.

Appendix

Neural Network for Biomedical Filtering

Parameter/Variable	Setting
Input Size	5000 (BOW vector)
Hidden Size	100
Output Layer Size	2
Activation Function	Sigmoid
Optimizer	Adam
Learning Rate	0.01
Loss Function	Cross Entropy
Training epochs	30
# Train/Dev/Test Data	894/111/111

Table 8: Implementation details and training parameters of the feed-forward neural network for biomedical filtering.

MLP Component of the Verdict Prediction Pipeline

Parameter/Variable	MLP-FEVER	MLP-Evidence
	Setting	Setting
Input Size	400	400
Hidden Size	100	100
Output Layer Size	3	3
Activation Function	ReLU	ReLU
Optimizer	Adam	Adam
Learning Rate	0.001	0.01
Loss Function	Cross Entropy	Cross Entropy
Max training epochs	200	120
Patience	30	10
Batch size	32	32
# Train/Dev/Test Data	116 359/14 544/14 544	439/6/6

Table 9: Implementation details and training/fine-tuning parameters of the MLP component in the verdict prediction pipeline. MLP-FEVER: training on FEVER 2018 dataset, MLP-Evidence: fine-tuning on CoVERT evidence extracts.

Error Analysis: Predictions on Language Model ‘Evidence’

ID	Tweet	Classification				
		Entity		Gold	Prediction	
		Masked	Pred.		MLP-FEVER	MLP-Evidence
<i>a</i>	Stop calling it a vaccine!! Vaccines contain the same germs that cause <u>disease</u> .	disease	measles	R	R	S
<i>b</i>	Covid-19 vaccines initiates an early and progressive clotting of blood in the lungs (pulmonary thrombosis) which impairs blood supply and gas exchange at lungs, leading to respiratory failure, which in majority of cases cause <u>death</u> .	death	death	R	S	S
<i>c</i>	If it's unclear it can't be reported. The death has to be contributed from Covid because Covid causes the respiratory issues etc that cause the death, Covid caused it. Like lung cancer if death because their lung stopped functioning because of cancer they died from lung cancer.	respiratory issues	cancer	S	R	S

Table 10: Instances where the pipeline failed to correctly classify tweets given evidence generated by BERT. R: REFUTES, S: SUPPORTS, N: NOT ENOUGH INFORMATION, MLP-FEVER: training on FEVER 2018 dataset, MLP-Evidence: fine-tuning on CoVERT evidence extracts. Masked entities in the tweet are underlined.

Predictions on Evidence Extracts

Tweet	ID	Evidence	Gold	Prediction	
				MLP-FEVER	MLP-Evidence
Stop calling it a vaccine!! Vaccines contain the same germs that cause disease.	<i>a</i>	mRNA vaccines teach our cells how to make a protein, or even just a piece of a protein that triggers an immune response inside our bodies.	R	S	S
	<i>b</i>	a preparation that is administered (as by injection) to stimulate the body's immune response against a specific infectious agent or disease: such as a preparation of genetic material (such as a strand of synthesized messenger RNA) that is used by the cells of the body to produce an antigenic substance (such as a fragment of virus spike protein)	R	S	S
Covid causes heart inflammation in like 70% of people (including asymptomatic). That seems like a high rate of heart inflammation..	<i>c</i>	COVID has been associated with a higher incidence of heart inflammation in adolescents and young adults.	S	S	S
	<i>d</i>	The SARS-CoV-2 virus can damage the heart in several ways. For example, the virus may directly invade or inflame the heart muscle, and it may indirectly harm the heart by disrupting the balance between oxygen supply and demand.	S	N	S
If it's unclear it can't be reported. The death has to be contributed from Covid because Covid causes the respiratory issues etc that cause the death, Covid caused it. Like lung cancer if death because their lung stopped functioning because of cancer they died from lung cancer.	<i>e</i>	Most people infected with the COVID-19 virus will experience mild to moderate respiratory illness	S	S	S
	<i>f</i>	COVID-19 is a respiratory disease, one that especially reaches into your respiratory tract, which includes your lungs. COVID-19 can cause a range of breathing problems, from mild to critical.	S	N	S

Table 11: Instances where the pipeline failed to correctly classify Tweet/Evidence pairs. R: REFUTES, S: SUPPORTS, N: NOT ENOUGH INFORMATION. MLP-FEVER: pipeline with MLP trained only on FEVER, MLP-Evidence: pipeline with MLP fine-tuned on evidence extracts from CoVERT. Source of evidence can be found in the corpus file.

Example Fact-Checking Annotation Environment

Instructions:

Fact-check the emboldened claim by

1. reading the tweet carefully
2. taking note of the claim in bold text
3. entering its key terms into Google Search
4. finding a reputable source that confirms or refutes the claim
5. mark the claim as being confirmed or refuted in the multiple choice question
6. enter the URL to the reputable source you have found into the given answer slot
7. copy and paste the segment of text from the source that explicitly confirms or refutes the claim

Each claim should take no more than 2 minutes to completely. The bonus is only awarded if a sensible URL and supporting text are pasted into the relevant fields. If you have investigated the Google Search results and are not able to find a confirmation/disconfirmation, please select the multiple choice option “No, I could not find any resource that confirms or refutes this claim”.

Tweet:

Stop calling it a vaccine! **Vaccines contain the same germs that cause disease.** (For example, measles vaccines contains measles virus, and Hib vaccine contains Hib bacteria.) But they have been either killed or weakened to the point that they don’t make you sick. Covid shot doesnt

Could you find a resource that confirms or refutes the claim in bold text?

- ☐ Yes, the resource CONFIRMS this claim
- ☐ Yes, the resource REFUTES this claim
- ☐ No, I could not find any resource that confirms or refutes this claim

Enter the URL to the reputable resource you have found that substantiates or refutes the claim from the question above. If you marked “No, I could not find any resource . . .”, you do not need to add a URL.

Enter URL here

Enter a short extract from this source that confirms or refutes the above claim.

Enter text here