

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

COVIDLIES: Detecting COVID-19 Misinformation on Social Media

Permalink

<https://escholarship.org/uc/item/3tz4f3fd>

Author

Hossain, Tamanna T

Publication Date

2021

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

COVIDLIES: Detecting COVID-19 Misinformation on Social Media

THESIS

submitted in partial satisfaction of the requirements
for the degree of

MASTER OF SCIENCE

in Computer Science

by

Tamanna Hossain

Thesis Committee :
Professor Sameer Singh, Chair
Professor Stephan Mandt
Professor Sean Young

2021

TABLE OF CONTENTS

| | Page |
|--|------------|
| LIST OF FIGURES | iii |
| LIST OF TABLES | iv |
| ACKNOWLEDGMENTS | vi |
| ABSTRACT OF THE THESIS | vii |
| 1 Introduction | 1 |
| 2 Problem Setup | 4 |
| 3 Dataset Collection | 6 |
| 4 Performance of Benchmark Models | 10 |
| 4.0.1 Evaluation Metrics | 11 |
| 4.0.2 Misconception Retrieval | 11 |
| 4.0.3 Stance Detection | 13 |
| 5 Related Work | 22 |
| 6 Conclusions and Future Work | 24 |
| Bibliography | 25 |

LIST OF FIGURES

| | Page |
|--|------|
| 1.1 COVIDLIES Dataset. Given a <i>tweet</i> , we annotate whether any of the known <i>misconceptions</i> are expressed in the tweet, in particular, if the tweet spreads the misconception (e.g., they <i>Agree</i>), combats the spread of the misconception (e.g., they <i>Disagree</i>), or takes No Stance towards the misconception. | 2 |
| 1.2 Misconception Detection Pipeline consisting of two sub-tasks, (a) <i>Misconception Retrieval</i> that identifies the known misconceptions that are relevant to the given tweet, and (b) <i>Stance Detection</i> that identifies whether the tweet agrees, disagrees, or expresses no stance, for each of the relevant misconceptions. | 2 |
| 4.1 Confusion Matrices for stance detection task using SBERT (DA) models trained on MultiNLI. The second model uses BERTSCORE (DA) to first determine whether a misconception-tweet pair is Relevant or No Stance, and only Relevant pairs are further classified by SBERT (DA). | 19 |

LIST OF TABLES

| | Page |
|--|------|
| 3.1 Top Misconceptions by Class. Misconceptions with more than 80 tweets total are ranked by the percentage of tweets annotated for each class. The top three misconceptions for each class with the corresponding percentage that a paired tweet would be annotated as that respective class are shown. For example, for the misconception ‘Democrats are using the coronavirus situation to harm President Trump’, 65% of the tweets paired with this misinformation were annotated as <i>Agree</i> . | 7 |
| 3.2 Distribution of Labels in the annotations. | 8 |
| 4.1 Misconception Retrieval Performance. We present evaluation for misinformative tweets (e.g., tweets that <i>Agree</i> with one or more misconceptions), as well as combined evaluation on Relevant tweets (i.e., tweets that either <i>Agree</i> or <i>Disagree</i> with one or more misconceptions.) | 13 |
| 4.2 Misconception Retrieval Examples. We present examples to demonstrate the difference in performance between some of the semantic similarity models: (I) BERTSCORE (DA), (II) Avg. BERT (DA), (III) Vanilla BERTScore, and (IV) Avg. GloVe. ✓= The model retrieved the relevant misconception for a tweet with rank 1; ✗= The model did not score the relevant misconception for the tweet with rank 1. | 14 |
| 4.3 Stance Detection Performance. We present evaluation for classification of tweet-misconception pairs into <i>Agree</i> , <i>Disagree</i> , and, No Stance classes. Precision (P), Recall (R), and F1-Score (F1) are presented for each class as well as macro averaged values. DA indicates domain-adaptive pretraining on COVID-19 tweets. | 16 |
| 4.4 Stance Detection Examples. Presenting examples of cases where combining or domain adaptation lead to flipping prediction to the correct class (<i>Agree</i>). All models here are trained on MultiNLI. | 18 |
| 4.5 Minimum Functionality Tests (MFT). Success rates of our best zero- and few-shot models on simple tests with tweets constructed using perturbations of misconceptions. Examples of each linguistic capability test with removals / additions and the expected output are shown. All examples are failures of at least one of the two models. | 20 |
| 4.6 Invariance Tests (INV). Success rates of our best zero- and few-shot models on tests with tweet perturbations that should not lead to output changes. Examples of each linguistic capability test with removals / additions and the expected output are shown. All examples are failures of at least one of the two models. | 21 |

4.7 **Directional Tests (DIR)**. Success rates of our best zero- and few-shot models on tests with tweet perturbations that should lead to soecific output changes. Examples of each linguistic capability test with **removals**/**additions** and the expected output are shown. All examples are failures of at least one of the two models. 21

ACKNOWLEDGMENTS

I would like to thank the many people who made this work possible. First, I thank my Chair, Dr. Sameer Singh, and committee members, Dr. Stephan Mandt and Dr. Sean Young, for their guidance and support. Next, I thank the co-authors of *COVIDLies: Detecting COVID-19 Misinformation on Social Media*, the EMNLP 2020 NLP COVID-19 Workshop paper (*Best Paper Award*) on which this thesis is based: Robert L. Logan, Arjuna Ugarte, and Yoshitomo Matsubara.

Next, I would like to thank the anonymous reviewers of EMNLP 2020 NLP COVID-19 workshop for their comments and the authors of related work for publishing their code and data. I would like to acknowledge Nicole Woodruff, an undergraduate at UCLA, Aileen Guillen, a medical student at UCI, Sadhika Yamasani, and Victoria Rong, undergraduates at UCI, for volunteering to help annotate data for this project. I would like to thank Lidia Flores, staff research associate in Dr. Young's Lab, for her help in compiling tweets. I also thank Elena Kochkina and Maria Liakata of the Alan Turing Institute for helpful discussions and feedback. Finally, I would like to thank my parents, husband, family and friends for their love and support.

This material is based upon work sponsored in part by NSF award #IIS-1817183, in part by NSF award #CNS-1925741, and in part by the DARPA MCS program under Contract No. N660011924033 with the United States Office Of Naval Research.

ABSTRACT OF THE THESIS

COVIDLIES: Detecting COVID-19 Misinformation on Social Media

By

Tamanna Hossain

Master of Science in Computer Science

University of California, Irvine, 2021

Professor Sameer Singh, Chair

The ongoing pandemic has heightened the need for developing tools to flag COVID-19-related misinformation on the internet, specifically on social media such as Twitter. However, due to novel language and the rapid change of information, existing misinformation detection datasets are not effective for evaluating systems designed to detect misinformation on this topic. Misinformation detection can be divided into two sub-tasks: (i) retrieval of misconceptions relevant to posts being checked for veracity, and (ii) stance detection to identify whether the posts *Agree*, *Disagree*, or express *No Stance* towards the retrieved misconceptions. To facilitate research on this task, we release COVIDLIES¹, a dataset of 6591 expert-annotated tweets to evaluate the performance of misinformation detection systems on 62 different pieces of COVID-19 related misinformation. We evaluate existing NLP systems on this dataset, providing initial benchmarks and identifying key challenges for future models to improve upon. For the stance detection sub-task we provide benchmark models in zero-shot and few-shot settings. In addition to evaluation of the models using standard metrics, we also provide behaviour testing of the best models of each setting.

¹<https://ucinlp.github.io/covid19>

Chapter 1

Introduction

Detecting spread of misinformation such as, rumors, hoaxes, fake news, propaganda, spear phishing, and conspiracy theories, is an important task for natural language processing (Thorne et al., 2017; Shu et al., 2017; Thorne and Vlachos, 2018). Online social media networks provide particularly fertile ground for the spread of misinformation—they lack gate-keeping and regulations, users publish content without having to go through an editor, peer review, verification of qualification, or providing sources, and social networks tend to create “echo chambers” or closed networks of communication insulated from disagreements.

The COVID-19 pandemic has created a pressing need for tools to combat the spread of misinformation. Since the pandemic affects the global community, there is a wide audience seeking information about the topic, whose safety is threatened by adversarial agents invested in spreading misinformation for political and economic reasons. Furthermore, due to the complexity of medical and public health issues, it is also difficult to be completely accurate and factual, leading to disagreements that get exacerbated with misinformation. This difficulty is compounded by the rapid evolution of knowledge regarding the disease. As researchers learn more about the virus, statements that seemed true may turn out to be false, and vice versa. Detecting this spread of pandemic-related misinfor-

Tweet: “Coronavirus CV19 was a top secret biological warfare experiment. That is why it is only affecting the poor.”

Misconception: “Coronavirus is genetically engineered.”

Label: Agree

Tweet: “I just read a tweet where someone claimed Coronavirus was actually a result of 5g exposure. These idiots walk among us.”

Misconception: “Coronavirus is caused by 5G.”

Label: Disagree

Tweet: “CDC: Coronavirus spreads rapidly in dense populations with public transit and regular social gatherings.”

Misconception: “Coronavirus can only survive in cold temperatures.”

Label: No Stance

Figure 1.1: **COVIDLIES Dataset.** Given a *tweet*, we annotate whether any of the known *misconceptions* are expressed in the tweet, in particular, if the tweet spreads the misconception (e.g., they *Agree*), combats the spread of the misconception (e.g., they *Disagree*), or takes *No Stance* towards the misconception.

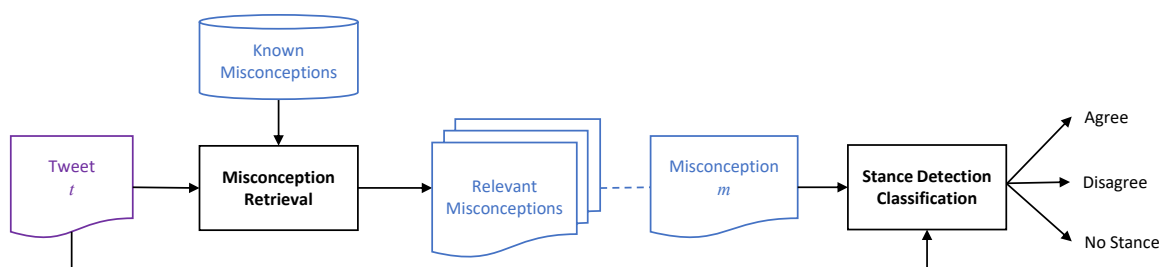


Figure 1.2: **Misconception Detection Pipeline** consisting of two sub-tasks, (a) *Misconception Retrieval* that identifies the known misconceptions that are relevant to the given tweet, and (b) *Stance Detection* that identifies whether the tweet agrees, disagrees, or expresses no stance, for each of the relevant misconceptions.

mation, thus, has become a critical problem, receiving significant attention from government and public health organizations (WHO, 2020), social media platforms (TechCrunch, 2020), and news agencies (BBC, 2020; CNN, 2020; New York Times, 2020).

In this paper, we introduce the COVIDLIES dataset for misconception detection on Twitter. COVIDLIES comprises of 62 common misconceptions about COVID-19 along with 6591 related tweets, identified and annotated by researchers from the UCI School of Medicine. Given a tweet, we annotate whether any of the *known* misconceptions, curated by the researchers, are expressed by the tweet. If

they are not, then they are considered No Stance. If they are, we further identify whether the tweet propagates the misconception (*Agree*) or is informative by contradicting it (*Disagree*). Example misconception-tweet pairs for each label are illustrated in Figure 1.1.

We provide benchmark results for each of these sub-tasks. First, we evaluate text similarity models on their ability to detect whether a tweet is relevant to a given misconception (a.k.a *misconception retrieval*). Next we evaluate zero-shot and few-shot models for the ability to detect the stance of each towards retrieved misconceptions (a.k.a *stance detection*.) For the zero-shot setting we train on the pre-existing tasks of natural language inference (NLI) and fact verification. For the few-shot setting we train on COVID-19 Health Risk Assessment task combined with a dataset of COVID-19 tweet-misconception pairs annotated for stance by researchers from the UCI School of Medicine. Our results show that existing models struggle at both tasks (49.7 Hits@1 for retrieval and 32.5 macro F1 on zero-shot stance detection), however improve considerably after domain adaptation (Gururangan et al. (2020); 74.3 Hits@1 for retrieval and 46.3 macro F1 on zero-shot stance detection). We see some further improvement on stance detection when using the few-shot setting (macro F1 of 49.3).

While our initial results using domain adaptation and few-shot learning are encouraging, they leave much room for improvement. There is still much work that needs to be done before NLP systems can be seriously considered for combating COVID-19-related misinformation, and we hope COVIDLIES will be useful to help researchers understand when such systems are ready to be deployed.

Chapter 2

Problem Setup

We assume access to a collection of positively phrased *known* misconceptions $M = \{m_1, \dots, m_{|M|}\}$, e.g., “Wearing masks does not prevent spread of COVID-19.” is a misconception. As we describe later, the set of misconceptions in this work are vetted, curated, and maintained by medical researchers. Given a collection of tweets, $T = \{t_1, \dots, t_{|T|}\}$, the task is to determine, for each input t , whether there exists a misconception $m \in M$ that is being discussed, and if so, whether the discussion propagates the misconception (i.e., identifies m as true, and thus is spreading the misconception) or refutes the misconception (i.e., identifies m as false). This task is naturally separated into the following steps (shown in Figure 1.2):

1. **Misconception Retrieval:** Given t return a subset $M_t \subseteq M$ of relevant misconceptions.
2. **Stance Detection:** For each (m, t) pair ($m \in M_t$), predict whether the m and t *Agree*, *Disagree*, or t takes *No Stance* with respect to m .

Due to limited availability of labeled data specific to this problem, we expect that models will need to be supervised on other, related tasks. For misconception retrieval, for example, relevant misconceptions can be ranked by measuring the semantic similarity between the tweet and each misconception, e.g., using cosine similarity between average word embeddings or more recent

transformer-based methods such as BERTSCORE (Zhang et al., 2019). For the stance detection sub-task we perform zero-shot learning by training on the pre-existing tasks of natural language inference (NLI) and fact verification. We also perform few-shot learning by training on COVID-19 Health Risk Assessment task combined with a dataset of COVID-19 tweet-misconception pairs annotated for stance by researchers from the UCI School of Medicine.

Chapter 3

Dataset Collection

Due to novel language used to describe the disease and its associated misconceptions, existing misinformation detection dataset are unlikely to be effective for evaluating systems designed to detect COVID-19-related misinformation on social media. To facilitate research on this problem, we collect an evaluation dataset, COVIDLIES; the collection process is described below.

Misconceptions We extract misconceptions from a Wikipedia article about misinformation related to the COVID-19 pandemic (Wikipedia, 2020). The extracted statements are manually examined, and statements that are not misinformation are removed. We manually rephrase the misinformation statements to a positive expression of that misinformation, e.g. “Some conspiracy theorists also alleged that the coronavirus outbreak was cover-up for a 5G-related illness” is shortened to “Coronavirus is caused by 5G”.

Tweets Our source of tweets is the collection of COVID-19-related tweets identified by Chen et al. (2020). We only use tweets from March and April 2020, and filter out non-English tweets.

| Class | Misconception | % |
|-----------|---|-------|
| Agree | SARS-CoV-2 can survive for weeks on surfaces. | 48.0 |
| | Drinking cow urine and applying cow dung on the body can cure covid-19. | 37.0 |
| | COVID-19 is only as deadly as the seasonal flu. | 22.1 |
| Disagree | COVID-19 is only as deadly as the seasonal flu. | 51.0 |
| | We're very close to a vaccine. | 12.0 |
| | Holy communion cannot be the cause of the spread of coronavirus. | 10.9 |
| No Stance | Acetic acid is effective against coronavirus. | 100.0 |
| | Cannabis protects against COVID-19. | 100.0 |
| | Clapping will kill coronavirus. | 100.0 |

Table 3.1: **Top Misconceptions by Class.** Misconceptions with more than 80 tweets total are ranked by the percentage of tweets annotated for each class. The top three misconceptions for each class with the corresponding percentage that a paired tweet would be annotated as that respective class are shown. For example, for the misconception ‘Democrats are using the coronavirus situation to harm President Trump’, 65% of the tweets paired with this misinformation were annotated as *Agree*.

Annotation Process To help identify tweets related to our list of misconceptions, we use BERTSCORE (Zhang et al., 2019) to compute a similarity metric on tweet-misconception pairs. For each given misconception, the 100 most similar tweets are selected for annotation. Each of these tweet-misconception pairs is manually labeled by researchers in the UCI School of Medicine as either: *Agree* (tweet is a positive expression of the misconception), *Disagree* (tweet contradicts/disagrees with the misconception), or *No Stance* (tweet is neutral or not relevant to the misconception).

Annotation Quality To evaluate inter-rater reliability, we randomly chose a subset of 200 tweet-misconception pairs and had four researchers manually label the subset. Percent agreement between researchers was 79%. Fleiss Kappa score was 0.69 which indicates substantial agreement between researchers (0.61–0.8). Disagreements were discussed and resolved before continuing to label the remaining tweet-misconception pairs in the dataset.

Most disagreements came down to labeler interpretation. For example, given the misconception,

| Class | Count | Percentage |
|-----------|-------|------------|
| Agree | 288 | 4.37 % |
| Disagree | 154 | 2.34 % |
| No Stance | 6,149 | 93.29 % |

Table 3.2: **Distribution of Labels** in the annotations.

“Drinking large amounts of water will protect against coronavirus”, and a tweet of “It’s a good thing everyone is stocking up on water to survive the Coronavirus because the 128 OZ of Diet Coke, the double cheeseburgers, and radiation from our phones definitely won’t kill us first”, one researcher labeled the pair as No Stance because it does not address any protective benefits but another researcher labeled the pair as *Agree* because people were stocking up on water to survive Coronavirus, which, to that researcher, implied water was protective. After discussions among the researchers, we concluded the pair was No Stance as it did not implicitly address the benefits of water and the statement was stated in a sarcastic tone. Other labeling challenges included deciding whether or not links or images in the tweet should be taken into account, as these could potentially change context of the tweet. We concluded that we would only evaluate the text as is since the various models would not be able to take images and links into account.

Dataset Statistics The current dataset contains 62 misconceptions, along with 6591 annotated tweet-misconception pairs. Statistics about the distribution of labels are provided in Table 3.2. The distribution is heavily skewed, containing mostly No Stance tweets, and a higher proportion of *Agree* tweets than *Disagree*. The heavy skew towards No Stance tweets could be a due to the dataset construction methodology, specifically using BERTSCORE without fine-tuning to retrieve tweets per misconception. As we show in 4.0.2, domain adaptation significantly improves misconception matching. Further, presence of more *Agree* than *Disagree* tweets could be due to a bias in BERTSCORE towards scoring agreement higher.

Top misconceptions for each class are shown in Table 3.1. We only consider misconceptions with more than 80 annotated tweets, and rank the misconceptions for each class by the proportion of

tweets that are annotated as that class. We present the top three misconceptions for each class with their corresponding percentage. There are misconceptions for which 100% of the paired annotated tweets express No Stance, which we do not see for the other two classes. We notice there is a misconception with nearly 50% of paired tweets labeled as *Agree*; and the highest proportion of Disagree labeled tweets found for any misconception in the Disagree class was 51%.

COVIDLIES, however, is an evolving dataset; annotation is not yet complete for all 62 Wikipedia misconceptions matched to 100 tweets using BERTSCORE, and we are continually identifying additional misconceptions, as well as collecting more recent tweets for annotation. Further, we will gather more relevant tweets by using domain-adapted retrieval models, which, as we will see in the next section, considerably outperform the current approach to retrieval, BERTSCORE.

Chapter 4

Performance of Benchmark Models

Supervised classifiers have been used extensively for detecting misinformation (Wang, 2017; Karimi et al., 2018; Shu et al., 2017, 2019). However, existing tasks involve static or slowly evolving domains, and topics that do not require specific expertise to annotate. Gathering an annotated dataset large enough to be used for training a COVID-19 misinformation detector is difficult: the way misconceptions are expressed rapidly evolves, and identifying whether or not something is a misconception requires expertise in public health and medicine. Further, even the misconceptions themselves change over time as we learn more about the disease and the pandemic. Thus, it is desirable that COVID-19 misinformation detection systems are: (i) *data efficient*, e.g., trained with little to no supervision, and (ii) *flexible*, e.g., allow the addition, removal, or modification of the known misconceptions.

In this section, we investigate whether models trained for related tasks in natural language processing can be adapted to misinformation detection on the COVIDLIES dataset without additional training. We specifically focus on models that can be used to score two input sequences, i.e., tweet-misconception pairs. Because these models come pretrained on different tasks, they are naturally *data efficient*, and furthermore, due to their pairwise nature, are also *flexible* as modification of

supported misconceptions is performed at the input level. Our code, dataset, and a demo of our best performing system are all available at <https://ucinlp.github.io/covid19>.

4.0.1 Evaluation Metrics

In the misconception retrieval sub-task, for a each tweet, t , the goal is to retrieve all the misconceptions that the tweet refers to (i.e. may be labeled **Agree** or **Disagree** by the annotators). Note: for clearer description, we introduce a new “pseudo-label”, **Relevant**, to refer to misconceptions that either **Agree** or **Disagree** with a given tweet. We treat this as a ranking task, where for each tweet, t , the system ranks the list of misconceptions, M , in decreasing order of relevancy. We evaluate this ranking using the standard information retrieval metrics $\text{Hits}@k$ and Mean Reciprocal Rank (MRR) for each Relevant misconception $m \in M_t^*$.

The stance detection sub-task is a standard classification problem with three classes (**Agree**, **Disagree**, and **No Stance**). As such we perform evaluation by measuring the precision, recall, and F1-score of the predicted classes.

4.0.2 Misconception Retrieval

We evaluate a number of information retrieval and semantic similarity approaches for the misconception retrieval sub-task.

Information Retrieval We use two information retrieval approaches. The first approach uses TF-IDF vectorization of tweets and misconceptions. Cosine similarity is used to score each tweet-misconception pair. Misconceptions are retrieved for each tweet in decreasing order of this score. NLTK is used for tokenization and vectorization. The second approach uses the BM25 algorithm, a bag-of-words retrieval technique which retrieves documents in decreasing probability of relevance

of the query term. IDF and document lengths are used to determine probability of relevance. We use the pyserini implementation of BM25 to retrieve misconceptions for each tweet.

Semantic Similarity We obtain vectorized representations of tweets and misconceptions using word embeddings. We then use two approaches for computing the semantic similarity between them: (i) cosine similarity computed between average token embeddings, and (ii) BERTSCORE (Zhang et al., 2019), which involves computation over BERT token embeddings of the tweet and misconception to obtain an F1-score-like measurement that we use as a similarity score.

For the cosine similarity approach, we experiment with both non-contextualized and contextualized word embeddings. For non-contextualized word embeddings we use 300D GloVe trained on 2014-Wikipedia and Gigaword embeddings (Pennington et al., 2014). For contextualized embeddings we use a pretrained BERT-LARGE (Devlin et al., 2018) model. However, Since BERT is not trained on COVID-19-related text we also use COVID-Twitter-BERT¹ (Müller et al., 2020) which uses domain-adaptive pretraining (Gururangan et al., 2020) on 160M tweets about COVID-19. For sake of brevity, we will append the suffix (DA) to models that use COVID-Twitter-BERT instead of spelling out the full model name.

Results We present the performance of similarity models in Table 4.1. Average embedding, both with GloVe and BERT embeddings, perform the worst (and are fairly similar to each other). Although information retrieval based approaches, TF-IDF and BM25, considerably outperform the average embedding techniques, BERTSCORE captures the similarity as accurately as well. Domain adaptation, however, further improves the embedding-based similarity techniques, improving average BERT embeddings to be as good as others, while making BERTSCORE much more accurate than all other techniques. Thus we see that using domain adaptation and BERTSCORE are both important for performing accurate misconception retrieval.

¹<https://huggingface.co/digitalepidemiologylab/covid-twitter-bert>

| Model | Agree | | | | Relevant (Agree or Disagree) | | | |
|------------------------------------|-------|------|------|------|------------------------------|------|------|------|
| | H@1 | H@5 | H@10 | MRR | H@1 | H@5 | H@10 | MRR |
| Cosine Sim., TF-IDF | 42.7 | 66.0 | 76.7 | 0.54 | 40.5 | 64.5 | 78.0 | 0.52 |
| BM25 | 49.7 | 76.4 | 83.0 | 0.62 | 49.9 | 74.7 | 81.7 | 0.61 |
| Cosine Sim., Avg. GloVe | 11.5 | 34.0 | 62.8 | 0.25 | 15.5 | 40.5 | 64.3 | 0.29 |
| Cosine Sim., Avg. BERT Embds. | 11.5 | 35.4 | 55.2 | 0.24 | 11.8 | 33.8 | 53.2 | 0.24 |
| BERTSCORE | 37.2 | 63.9 | 84.7 | 0.51 | 42.3 | 68.2 | 84.5 | 0.56 |
| <i>with Domain Adaptation (DA)</i> | | | | | | | | |
| Cosine Sim., Avg. BERT Embds. | 55.2 | 73.3 | 80.2 | 0.65 | 52.5 | 71.7 | 78.9 | 0.62 |
| BERTSCORE | 76.7 | 94.4 | 97.2 | 0.84 | 74.3 | 90.2 | 93.9 | 0.81 |

Table 4.1: **Misconception Retrieval Performance.** We present evaluation for misinformative tweets (e.g., tweets that *Agree* with one or more misconceptions), as well as combined evaluation on Relevant tweets (i.e., tweets that either *Agree* or *Disagree* with one or more misconceptions.)

We illustrate the differences in the similarity models using example predictions in Table 4.2. The first example provides a challenging case of retrieval that requires taking both COVID-19 knowledge and contextual information (e.g. multiple sentences, ‘hot’ vs ‘cold’) into account, and thus only the BERTSCORE (DA) model is able to retrieve the correct misconception. The second example primarily requires domain knowledge that ‘coronavirus’ and ‘Sars-cov-2’ are very similar, and only domain-adapted models are able to score the correct misconception highest. The last example shows when contextual embeddings (BERT) outperform non-contextual embedding (GloVe).

4.0.3 Stance Detection

Due to the lack of adequately large datasets for stance detection with pairs of sentences (Mohammad et al., 2016; Ferreira and Vlachos, 2016; Gorrell et al., 2018), we cannot use existing datasets to train models for our setup. However, since classes in misinformation detection correspond to those in natural language inference (NLI) and fact verification, tasks with much larger training datasets, we instead experiment with zero-shot learning on these tasks. The COVID-19 Health Risk Assessment task released after the pandemic started, also allowed us to experiment with few-shot learning by combining it with COVID-19 tweet-misconception pairs annotated for stance by researchers at the

| I | II | III | IV | Example |
|---|----|-----|----|---|
| ✓ | ✗ | ✗ | ✗ | <p>Tweet: @lillyleiris I guess because weather is hot there . I’m in philipines just 3 cases here . Corona virus don’t like hot weather : Argentina, Thailand , Indonesia , philipines ... look more safe than others countrys</p> <p>Misconception: Coronavirus can only survive in cold temperatures.</p> |
| ✓ | ✓ | ✗ | ✗ | <p>Tweet: There is evidence that coronaviruses can live on inanimate surfaces for up to nine days, but it’s not yet clear how likely humans are to be infected by touching these surfaces. https://t.co/DJ99AAISWw</p> <p>Misconception: SARS-CoV-2 can survive for weeks on surfaces.</p> |
| ✓ | ✓ | ✓ | ✗ | <p>Tweet: Covid-19 is about 43 times more deadly if you get it, but China’s number of cases is leveling off at around 80K which is much less than the number of US flu cases. If the number of cases is kept small then Covid-19 will be minor compared to the flu.</p> <p>Misconception: COVID-19 is only as deadly as the seasonal flu.</p> |

Table 4.2: **Misconception Retrieval Examples.** We present examples to demonstrate the difference in performance between some of the semantic similarity models: **(I)** BERTSCORE (DA), **(II)** Avg. BERT (DA), **(III)** Vanilla BERTScore, and **(IV)** Avg. GloVe. ✓= The model retrieved the relevant misconception for a tweet with rank 1; ✗= The model did not score the relevant misconception for the tweet with rank 1.

UCI School of medicine.

Zero-shot We train linear classifiers on three common NLI datasets—SNLI (Bowman et al., 2015), MultiNLI (Williams et al., 2018), and MedNLI (Shivade, 2019). These classifiers use the following features, respectively: (i) concatenated unigram and bigram TF-IDF vectors for each input, (ii) concatenated average GloVe embeddings for each input, (iii) Bidirectional LSTM encoding, and (iv) the Sentence-BERT (SBERT) (Reimers and Gurevych, 2019) representation that uses siamese and triplet networks to obtain semantically meaningful sentence embeddings. Note that for (iii) and (iv), the transformer architectures (BiLSTM and SBERT) are jointly trained with the linear classifier. We also train SBERT on two fact verification datasets—FEVER (Thorne et al., 2018b), and SciFact (Wadden et al., 2020).

Few-shot For the few shot setting we train SBERT on a combination of datasets consisting of COVID-19 tweets and misinformative claims. Firstly, we use COVID-HeRA (Dharawat et al., 2020), a dataset for assessing health risks of COVID-19-related social media posts. It consists of pairs of tweets and misinformative claims or fake news headlines with label of health risk severity. Labels of *not severe*, *possibly severe*, and *highly severe* are all mapped to *Agree*, while *refute/rebutts* is mapped to *Disagree*. Records for No Stance are created by randomly matching tweets and misinformative claims or fake news headlines. We combine COVID-HeRA with 2,466 tweet-misconception pairs annotated for stance by researchers at the UCI School of Medicine. This dataset is constructed by matching 9 COVID-19 misconceptions from Poynter (Poynter, 2020) with related tweets using keyword searches via Talkwalker². Imbalance of classes is corrected by oversampling of the *Disagree* class in both COVID-HeRA and our stance annotations. This combined dataset consists of some COVID-19 misconceptions that are similar to those in COVIDLIES, eg. *'Hydroxychloroquine-azithromycin is a proven cure for COVID-19 patients.'* vs. *'Chloroquine can cure coronavirus'* in COVIDLIES. We will refer to this dataset as CH+PA (COVID-HeRA+Poynter Annotations).

BERTSCORE (DA) + Linear Classifier Since BERTSCORE with domain adaptation performs best at retrieval for relevant classes, we use it to improve stance detection. We combine BERTSCORE (DA) with our linear classifiers, initially classifying tweet-misconception pairs with high BERTSCORE scores (>0.4) as Relevant, subsequently using the linear model to determine whether the pair *Agree* or *Disagree*. We denote such “combined models” by inserting a plus sign between the retrieval model and linear model, e.g., BERTSCORE (DA) + BiLSTM denotes a model that uses BERTSCORE (DA) to determine retrieve relevant misconceptions and a BiLSTM model for classifying the stance.

Results Stance detection results in Table 4.3 show that, generally, most models do not perform well on the *Agree* and *Disagree* classes, which are minority classes in our dataset. On the other hand, performance on No Stance is high; quite a few models achieve an F1-score of 93.9% or higher.

²<https://www.talkwalker.com/>

| Model | Macro Avg. | | | Agree | | | Disagree | | | No Stance | | |
|-----------------------------|------------|------|------|-------|------|------|----------|------|------|-----------|------|------|
| | P | R | F1 | P | R | F1 | P | R | F1 | P | R | F1 |
| Zero-shot | | | | | | | | | | | | |
| <i>Trained on SNLI</i> | | | | | | | | | | | | |
| Linear, Bag-of-Words | 32.7 | 33.3 | 20.8 | 2.6 | 12.1 | 4.3 | 3.1 | 51.5 | 5.8 | 92.3 | 36.4 | 52.2 |
| Linear, Avg. GloVe | 33.3 | 28.4 | 23.6 | 5.5 | 17.3 | 8.3 | 1.5 | 24.0 | 2.8 | 92.9 | 44.0 | 59.8 |
| BiLSTM | 34.6 | 40.7 | 21.7 | 4.5 | 15.9 | 7.1 | 3.6 | 71.3 | 6.9 | 95.5 | 34.9 | 51.1 |
| SBERT | 34.4 | 31.2 | 22.8 | 8.7 | 20.1 | 12.1 | 1.8 | 36.3 | 3.4 | 92.6 | 37.1 | 53.0 |
| SBERT (DA) | 35.0 | 30.4 | 20.0 | 15.0 | 23.2 | 18.2 | 1.7 | 43.3 | 3.2 | 88.5 | 24.7 | 38.6 |
| BERTSCORE (DA) + BiLSTM | 40.3 | 45.1 | 40.7 | 12.7 | 6.9 | 8.9 | 13.1 | 35.7 | 19.2 | 95.2 | 92.7 | 93.9 |
| BERTSCORE (DA) + SBERT (DA) | 44.6 | 45.0 | 43.3 | 29.3 | 18.3 | 22.6 | 9.3 | 24.0 | 13.4 | 95.2 | 92.7 | 93.9 |
| <i>Trained on MultiNLI</i> | | | | | | | | | | | | |
| Linear, Bag-of-Words | 34.1 | 38.3 | 15.3 | 2.9 | 38.4 | 5.4 | 6.2 | 59.1 | 11.3 | 93.2 | 17.3 | 29.2 |
| Linear, Avg. GloVe | 31.7 | 28.0 | 16.0 | 3.0 | 27.7 | 5.4 | 2.4 | 32.2 | 4.5 | 89.6 | 24.1 | 38.0 |
| BiLSTM | 32.2 | 32.9 | 32.5 | 3.5 | 5.2 | 4.2 | 0.0 | 0.0 | 0.0 | 93.1 | 93.5 | 93.3 |
| SBERT | 36.4 | 47.0 | 27.8 | 10.1 | 40.8 | 16.2 | 3.7 | 56.1 | 6.9 | 95.5 | 44.0 | 60.2 |
| SBERT (DA) | 48.4 | 51.7 | 39.6 | 45.9 | 29.1 | 35.6 | 4.6 | 64.3 | 8.5 | 94.8 | 61.8 | 74.8 |
| BERTSCORE (DA) + BiLSTM | 37.1 | 42.4 | 38.6 | 16.1 | 34.6 | 22.0 | 0.0 | 0.0 | 0.0 | 95.2 | 92.7 | 93.9 |
| BERTSCORE (DA) + SBERT (DA) | 50.8 | 49.9 | 46.9 | 44.3 | 20.1 | 27.6 | 12.8 | 36.8 | 19.0 | 95.2 | 92.7 | 93.9 |
| <i>Trained on MedNLI</i> | | | | | | | | | | | | |
| Linear, Bag-of-Words | 35.0 | 39.3 | 19.9 | 3.7 | 50.9 | 6.9 | 7.6 | 41.5 | 12.9 | 93.7 | 25.5 | 40.0 |
| Linear, Avg. GloVe | 36.0 | 45.8 | 21.0 | 5.3 | 62.3 | 9.8 | 5.5 | 47.4 | 9.9 | 97.3 | 27.8 | 43.2 |
| BiLSTM | 32.3 | 32.6 | 22.4 | 4.3 | 54.0 | 7.9 | 0.0 | 0.0 | 0.0 | 92.7 | 43.8 | 59.5 |
| SBERT | 35.4 | 39.4 | 15.6 | 5.3 | 90.3 | 10.0 | 3.4 | 8.8 | 4.9 | 97.6 | 19.2 | 32.0 |
| SBERT (DA) | 35.8 | 46.7 | 21.4 | 6.9 | 59.2 | 12.4 | 3.9 | 52.0 | 7.2 | 96.7 | 28.9 | 44.5 |
| BERTSCORE (DA) + BiLSTM | 40.4 | 42.4 | 38.9 | 16.0 | 33.9 | 21.8 | 10.0 | 0.6 | 1.1 | 95.2 | 92.7 | 93.9 |
| BERTSCORE (DA) + SBERT (DA) | 44.3 | 48.1 | 45.7 | 21.9 | 26.0 | 23.8 | 15.7 | 25.7 | 19.5 | 95.2 | 92.7 | 93.9 |
| <i>Trained on FEVER</i> | | | | | | | | | | | | |
| SBERT (DA) | 42.5 | 40.9 | 41.1 | 25.3 | 23.9 | 24.6 | 7.5 | 1.8 | 2.8 | 94.7 | 97.0 | 95.8 |
| BERTSCORE (DA) + SBERT (DA) | 39.4 | 41.4 | 40.2 | 15.9 | 22.8 | 18.8 | 7.2 | 8.8 | 7.9 | 95.2 | 92.7 | 93.9 |
| <i>Trained on SciFact</i> | | | | | | | | | | | | |
| SBERT (DA) | 33.0 | 31.8 | 30.0 | 3.6 | 20.8 | 6.2 | 2.8 | 1.8 | 2.2 | 92.6 | 73.0 | 81.7 |
| BERTSCORE (DA) + SBERT (DA) | 42.9 | 44.1 | 42.1 | 17.1 | 32.5 | 22.4 | 16.4 | 7.0 | 9.8 | 95.2 | 92.7 | 93.9 |
| Few-shot | | | | | | | | | | | | |
| <i>Trained on CH+P</i> | | | | | | | | | | | | |
| SBERT (DA) | 41.2 | 53.2 | 42.4 | 14.0 | 55.4 | 22.4 | 14.1 | 24.0 | 17.8 | 95.3 | 80.3 | 87.1 |
| BERTSCORE (DA) + SBERT (DA) | 48.8 | 50.8 | 49.3 | 18.5 | 29.8 | 22.8 | 32.7 | 29.8 | 31.2 | 95.2 | 92.7 | 93.9 |

Table 4.3: **Stance Detection Performance.** We present evaluation for classification of tweet-misconception pairs into Agree, Disagree, and, No Stance classes. Precision (P), Recall (R), and F1-Score (F1) are presented for each class as well as macro averaged values. DA indicates domain-adaptive pretraining on COVID-19 tweets.

For the Zero-shot setting, BERTSCORE (DA) + SBERT (DA) (on MultiNLI) achieves the highest F1 (27.6) for the *Agree* class, while also obtaining the highest macro averaged Precision (50.8%), Recall (49.9%) and F1 (46.9). The few-shot setting slightly improves upon the macro averaged Recall(50.8%) and F1(49.3%). We notice that the combined BERTSCORE (DA) + linear classifier approach improves F1 across classes for many models.

Examples of stance predictions in Table 4.4 illustrate the differences between some of the zero-shot MultiNLI models. The first example demonstrates that knowledge about the domain vocabulary helps domain adapted models in predicting the correct stance, as it did for retrieval. The remaining two examples both show the advantage of the combined BERTSCORE (DA) + NLI approach, in particular, demonstrating that retrieval models are effective at identifying relevant misconceptions, which the NLI models are then able to correctly classify the stance of. As we can see from the confusion matrices of SBERT (DA) vs. BERTSCORE (DA) + SBERT(DA) (Figure 4.1), combining BERTSCORE (DA) with NLI greatly decreases the misclassification of *Disagree* as No Stance and No Stance as *Disagree*.

Behaviour Testing Standard evaluation metrics like we have used above can overestimate the real world performance of NLP models, and do not reveal enough information about situations where the models are failing or how to fix them. To evaluate our stance detection models more rigorously we use the matrix of linguistic capabilities and test types provided by CheckList (Ribeiro et al., 2020) for behaviour testing of NLP models. We evaluate our best zero- and few-shot models: BERTSCORE (DA) + SBERT(DA) trained on MNLI and CH+PA respectively. We test the following linguistic capabilities: Robustness (typos, irrelevant changes, paraphrases), Negation, Vocabulary, NER (named entity recognition), Temporal (understanding order of events), and SRL (semantic role labelling). We perform three types of types: (i) Minimum Functionality Tests (MFT), which are simple 'sanity checks' of targetted capabilities; (ii) Invariance Tests (INV) consisting of perturbations which should not change model output; and (iii) Directional Expectation Tests (DIR)

| Models | Labels | Example |
|-----------------------------|-----------|---|
| SBERT | Agree | Tweet: @JWrightforCA34 You are an out of control danger to the people. You’re spreading deliberate lies. Coronavirus 19 is NOT man made. It (like other Conrona Viruses) was passed from animal to human. Wild game eating from a local market in Wuhan is what’s most |
| SBERT (DA) | Disagree | Misconception: Coronavirus is genetically engineered. |
| BiLSTM | No Stance | Tweet: @Acyn The corona virus can live on a surface for up to 9 days. Just saying. |
| BERTSCORE (DA) + BiLSTM | Agree | Misconception: SARS-CoV-2 can survive for weeks on surfaces. |
| SBERT (DA) | No Stance | Tweet: @alexsalvinews Alex. Check out Dean Koontz, The Eyes of Darkness. 1981. He predicts the Wuhan-400 virus. He said in ”around” 2020, a pneumonia-like virus will be spread worldwide. |
| BERTSCORE (DA) + SBERT (DA) | Agree | Misconception: Dean Koontz predicted the pandemic in his 1981 novel The Eyes of Darkness. |

Table 4.4: **Stance Detection Examples.** Presenting examples of cases where combining or domain adaptation lead to flipping prediction to the correct class (*Agree*). All models here are trained on MultiNLI.

consisting of perturbations which should change model output in a specific way.

For all tests, we use the COVID-19 misconceptions from COVIDLIES. For MFTs we construct simple tweets based on perturbations of the misconceptions, eg. introducing a simple typo: Salt water → wtaer protects from coronavirus. From the results of the tests in 4.5 we see that there are linguistic capabilities that the zero-shot model is more competent at than the few-shot model and vice versa. Notably, the few-shot model is more robust to typos and positive paraphrases possibly because it is trained on informally written content (i.e. tweets). However, an alarming incompetency of the few-shot model is when the constructed tweet is identical to the misconception— the output should obviously be *Agree*, which the zero-shot model is able to correctly predict 100% of the time, while the the few-shot model only 91.9% of the time (it predicts *Disagree* in some cases). A possible explanation for why this happens is that the tweets in COVID-Hera tend to repeat the headlines or claims they are paired with verbatim, even in the *refute/rebutts* (i.e. *Disagree*) category. Both models

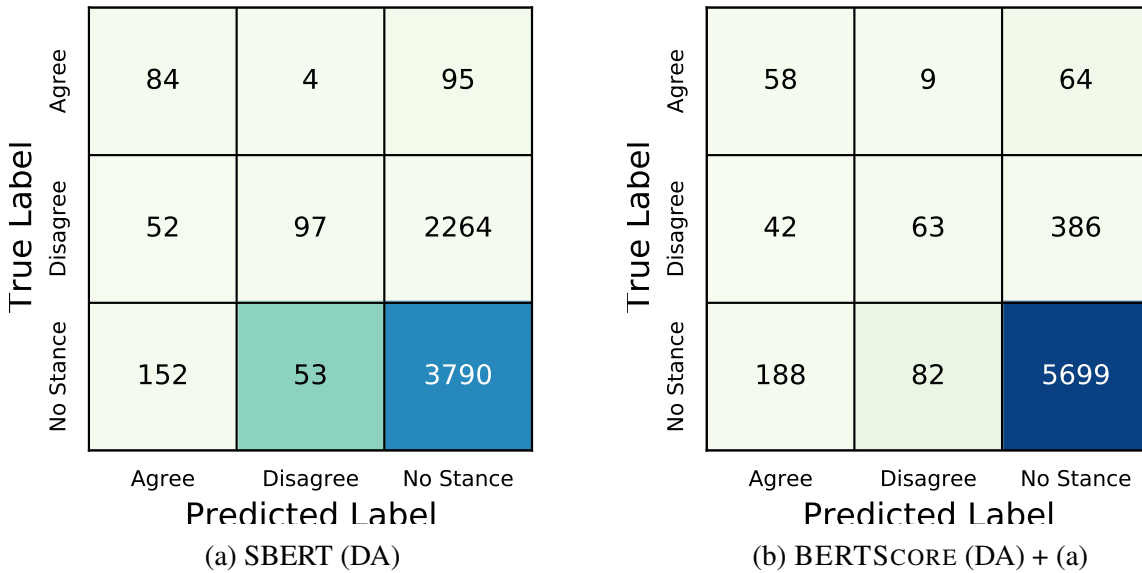


Figure 4.1: **Confusion Matrices** for stance detection task using SBERT (DA) models trained on MultiNLI. The second model uses BERTSCORE (DA) to first determine whether a misconception-tweet pair is Relevant or No Stance, and only Relevant pairs are further classified by SBERT (DA).

failed spectacularly at NER tests (0% success rates) when perturbing country names and numbers in misconceptions to construct tweets, i.e., they always fail to predict No Stance in these cases.

For INV and DIR tests we randomly sample 45 misconception-tweet pairs per class from COVIDLIES that both models are able to correct the label for. For INV tests we perturbed tweets in ways that should not lead to output changes, and for DIR tests perturbations were to induce a specific label switch. INV tests were more successful (usually greater than 70% success rate; Table 4.6) than DIR tests (less than 20% success rate; Table 4.7) for both models .

| | Test Description | Success Rate | | Example test cases and expected behavior | |
|-----------------|---|--------------|------|--|---|
| | | Zero | Few | Misconception | Constructed Tweet |
| Robust. | Identical | 100.0 | 91.9 | COVID-19 is only as deadly as the seasonal flu. | COVID-19 is only as deadly as the seasonal flu. Agree |
| | Positive variation | 54.9 | 75.9 | Cocaine cures coronavirus. | Using coke will protect you against COVID. Agree |
| | One typo | 91.9 | 91.9 | Salt water protects from coronavirus. | Salt water → wtacr protects from coronavirus. Agree |
| | Three typos | 61.3 | 85.5 | Cannabis protects against COVID-19. | Cannabis → Cannabips protects → rotects against → aagins COVID-19. Agree |
| | Five typos | 58.1 | 82.3 | Avocado and mint tea cures coronavirus. | Avocado and → an mint → dmit tea → ntea cures → curse coronavirus → coronavisu . Agree |
| | Contraction | 100.0 | 42.9 | Shaking hands cannot spread coronavirus. | Shaking hands cannot → can't spread coronavirus. Agree |
| | Punctuation | 100.0 | 93.6 | Jamaat cannot become a source of infection. | Jamaat cannot become a source of infection. Agree |
| Negation | Insert <i>not</i> | 100.0 | 79.0 | Lions were freed to keep people off the streets in Moscow. | Lions were not freed to keep people off the streets in Moscow. Disagree |
| | Insert <i>I disagree with this</i> | 50.0 | 61.3 | Africans are more resistant to coronavirus. | Africans are more resistant to coronavirus. I disagree with this. Disagree |
| | Negative variations | 70.5 | 45.0 | Coronavirus is caused by 5G. | Viruses cannot travel on radio waves/mobile networks |
| Vocab. | Insert <i>I agree with this</i> | 100.0 | 95.2 | Jamaat cannot become a source of infection. | Jamaat cannot become a source of infection. I agree with this. Agree |
| | Insert <i>definitely</i> | 100.0 | 90.3 | COVID-19 is only as deadly as the seasonal flu. | COVID-19 is definitely only as deadly as the seasonal flu. Agree |
| NER | Perturb country | 0.0 | 0.0 | A pack of elephants descended on a village under quarantine in China's Yunnan, got drunk on corn wine, and passed out in a tea garden. | A pack of elephants descended on a village under quarantine in China's → Brazil's Yunnan, got drunk on corn wine, and passed out in a tea garden. No Stance |
| | Perturb number | 0.0 | 0.0 | Bank of England £20 banknotes contain a picture of a 5G mast and the SARS-CoV-2 virus. | Bank of England £ 20 → 15 banknotes contain a picture of a 5G mast and the SARS-CoV-2 virus. No Stance |
| SRL | Parsing misconception in (question, "Yes") form | 96.8 | 88.7 | Coronavirus can only survive in cold temperatures. | Do I think that coronavirus can only survive in cold temperatures? Yes. Agree |
| | Parsing misconception in (question, "No") form | 66.1 | 82.3 | Dean Koontz predicted the pandemic in his 1981 novel The Eyes of Darkness. | Do I think that dean Koontz predicted the pandemic in his 1981 novel The Eyes of Darkness? No. Disagree |
| | Author stance is more important than others | 77.4 | 95.1 | Africans are more resistant to coronavirus. | Some people think that africans are more resistant to coronavirus. However, I don't think so. Disagree |
| Temporal | Stance change over time, present should prevail | 69.3 | 72.6 | SARS-CoV-2 can survive for weeks on surfaces. | I used to think sARS-CoV-2 can survive for weeks on surfaces. Now I don't think so anymore. Disagree |
| | Stance change over time, present should prevail | 61.3 | 87.1 | Shaking hands cannot spread coronavirus. | I used to disagree that shaking hands cannot spread coronavirus. Now I agree with this. Agree |
| Tone | Infer author stance from tone towards others | 59.7 | 91.9 | Kenneth Copeland on Victory Channel can cure COVID-19 through the TV. | Wow some idiots believe kenneth Copeland on Victory Channel can cure COVID-19 through the TV. |

Table 4.5: **Minimum Functionality Tests (MFT)**. Success rates of our best zero- and few-shot models on simple tests with tweets constructed using perturbations of misconceptions. Examples of each linguistic capability test with **removals**/**additions** and the expected output are shown. All examples are failures of at least one of the two models.

| Test Description | Success Rate | | Example Test Cases and Expected Behaviour | |
|--|--------------|------|--|--|
| | Zero | Few | Misconception | Altered Tweet |
| Robust. One typo | 91.1 | 89.6 | Africans are more resistant to coronavirus. | Social media folks: Africans → African: are immune to Corona virus Corona Virus: https://t.co/AFkDYSpTUF Agree |
| | 80.7 | 81.5 | Using namaste as a greeting prevents transmission of covid-19. | Namaste, not handshakes → ahndshakes, to → top prevent → revent Covid-19 https://t.co/tkTocDvWxy Agree |
| | 68.2 | 69.6 | Ginger ale will cure the Coronavirus | Ginger ale → al will → ciwl cure → ucre the → th Coronavirus → Croonavirus Agree |
| | 95.6 | 93.3 | COVID-19 is only as deadly as the seasonal flu. | @DSPArchives Yes because the flu is just as deadly as the coronavirus don't → do not you know Agree |
| | 93.3 | 91.1 | Coronavirus is genetically engineered. | The corona virus is a government made virus that is being used for population control Agree |
| Vocab. Insert <i>I agree with this</i> | 72.2 | 68.9 | Ginger ale will cure the Coronavirus | I disagree with this: Ginger ale will cure the Coronavirus Disagree |
| | 93.9 | 94.4 | Coronavirus is caused by 5G. | I just read a the tweet where someone claimed Coronavirus was actually the result of 5g exposure. These idiots walk among us. Disagree |
| SRL Parsing tweet in (question, "yes") form | 87.8 | 85.6 | Coronavirus is genetically engineered. | Do I think this? Coronavirus CV19 was a top secret biological warfare experiment. That is why it is only affecting the poor Yes. Agree |
| | 72.2 | 71.1 | Using namaste as a greeting prevents transmission of covid-19. | I used to disagree with this: Namaste, not handshakes, to prevent Covid-19 https://t.co/tkTocDvWxy Now I don't disagree anymore. Agree |

Table 4.6: **Invariance Tests (INV)**. Success rates of our best zero- and few-shot models on tests with tweet perturbations that should not lead to output changes. Examples of each linguistic capability test with **removals**/**additions** and the expected output are shown. All examples are failures of at least one of the two models.

| Test Description | Success Rate | | Example test cases and expected behavior | |
|---|--------------|------|--|--|
| | Zero | Few | Misconception | Altered Tweet |
| Negation Insert <i>"I disagree with this"</i> | 8.9 | 8.9 | Boiled ginger can cure coronavirus. | I disagree with this: Ginger ale will cure the Coronavirus Agree → Disagree |
| | 14.4 | 16.7 | Dean Koontz predicted the pandemic in his 1981 novel The Eyes of Darkness. | This is a hoax: A Dean Koontz novel written in 1981 predicted the outbreak of the coronavirus! https://t.co/zhSPngR9DM Agree → Disagree |
| SRL Parsing misconception in (question, "No!") form | 7.8 | 15.6 | Dean Koontz predicted the pandemic in his 1981 novel The Eyes of Darkness. | Do I think this? A Dean Koontz novel written in 1981 predicted the outbreak of the coronavirus! https://t.co/zhSPngR9DM No Agree → to Disagree |
| Temporal Stance change over time, present should prevail | 7.78 | 12.2 | COVID-19 is only as deadly as the seasonal flu. | I used to agree with this: COVID-19 is NOT like the FLU https://t.co/m2XbxkpcDC Now I don't think so anymore. Disagree → Agree |

Table 4.7: **Directional Tests (DIR)**. Success rates of our best zero- and few-shot models on tests with tweet perturbations that should lead to specific output changes. Examples of each linguistic capability test with **removals**/**additions** and the expected output are shown. All examples are failures of at least one of the two models.

Chapter 5

Related Work

COVID-19: In the social sciences, there have been recent efforts to quantify COVID-19 misinformation on social media (Brennen et al., 2020; Kouzy et al., 2020), as well experimental efforts to prevent propagation of misinformation (Pennycook et al., 2020). At the same time, members of the NLP community have been working on developing tools for the automatic detection of COVID-19-related misinformation online. Serrano et al. (2020) detect YouTube videos spreading conspiracy theories using features of user comments, and Dharawat et al. (2020) classify tweets by the severity of health risks associated with them. McQuillan et al. (2020) study the behaviour of COVID-19 misinformation networks on Twitter using mapping, topic modeling, bridging centrality, and divergence. Penn Medicine launched a chatbot to provide patients with accurate information about the virus (VolppKevin et al., 2020), and a crowdsourced chatbot, Jennifer, is also available to answer questions about the pandemic (Li et al., 2020). We are the first to frame COVID-19 misinformation detection as a two-stage task of misconception retrieval and pairwise classification of stance, and add to this body of work by providing a dataset and benchmark models for automated identification of misinformation.

Misinformation Detection: There are several datasets for misinformation detection with binary veracity labels, for example, FakeNewsNet (Shu et al., 2017, 2019, 2020) consisting of news articles, Some Like It Hoax (Tacchini et al., 2017) consisting of Facebook posts, and PHEME (Zubiaga et al., 2018) containing twitter threads.

Misinformation detection is also closely related to fact-checking since both tasks aim to assess the veracity of claims. FEVER (Thorne et al., 2018a, 2019) is a dataset of claims and evidence pairs with Supported, Refuted or NotEnoughInfo labels to facilitate research in automated fact checking. This is similar to Emergent (Ferreira and Vlachos, 2016), a stance classification dataset consisting of rumored claims and associated news articles with labels of For, Against, or Observing the claim. Stance detection is also the focus of the Fake News Challenge (FNC-1)¹ consisting of pairs of news article headlines and body texts with Agrees, Disagrees, Discusses, and Unrelated labels.

Our proposed models for detecting misinformation by using classifiers fall within the framework of detecting misinformation using content features (Volkova et al., 2017; Wei and Wan, 2017). Other approaches include using crowd behaviour (Tschitschek et al., 2018; Mendoza et al., 2010), reliability of the source (Lumezanu et al., 2012; Li et al., 2015), knowledge graphs (Ciampaglia et al., 2015), or a combination of these approaches (Castillo et al., 2011; Kumar et al., 2016). Adapting these techniques to COVID-19 misinformation is a promising direction for future work.

¹<http://www.fakenewschallenge.org/>

Chapter 6

Conclusions and Future Work

The ongoing COVID-19 pandemic has been accompanied by a corresponding ‘infodemic’ of misinformation about the virus. It is important to develop tools to automatically detect misinformation online, especially on social media sites where the volume and speed of the spread are high. However, rapidly evolving information and novel language make existing misinformation detection datasets and models ineffective for detecting COVID-19 misinformation. In this work, to initiate research on this important and timely topic, we introduced COVIDLIES, a benchmark dataset containing known COVID-19 misconceptions accompanied with tweets that **Agree**, **Disagree**, or express **No Stance** for each misconception, annotated by experts. Our code, dataset, and a demo of our best performing system are publicly available at <https://ucinlp.github.io/covid19>.

Given a tweet, we formulate the task of detecting misinformation as retrieving relevant misconceptions, and classifying whether the tweet supports or refutes it. We evaluate a number of approaches for this task, including common semantic similarity models for retrieval, accurate models trained on a variety of NLI datasets, and domain adaptation by pretraining language models on a corpus of COVID-19 tweets. We demonstrate domain adaptation significantly improves results for both sub-tasks of misinformation detection. We also show that it is feasible to detect the stance of tweets

towards misconceptions using both zero-shot and few-shot settings. We showed that few-shot learning slightly improves stance detection when evaluating using standard aggregate performance metrics. However, further behaviour testing using CheckList (Ribeiro et al., 2020) leaves an unclear picture of which setting is better. Both settings have considerable scope for improvement. Future work will involve exploring improved performance on stance detection, preferably without reliance of methods that require a large amount of data collection (like domain-adaptation) since they are not quickly available in an emerging crisis. We plan to continually expand our annotated dataset by including posts from other domains such as news articles and Reddit, and misconceptions from sources beyond Wikipedia, such as Poynter (2020). We invite researchers to build COVID-19 misinformation detection systems and evaluate their performance using the presented dataset.

Bibliography

- BBC. 2020. 'Hundreds dead' because of Covid-19 misinformation. <https://www.bbc.com/news/world-53755067> [Accessed on June 30, 2020].
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- J Scott Brennan, Felix M Simon, Philip N Howard, and Rasmus Kleis Nielsen. 2020. Types, sources, and claims of COVID-19 misinformation. *Reuters Institute*.
- Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on twitter. In *Proceedings of the 20th international conference on World wide web*, pages 675–684.
- 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.
- Giovanni Luca Ciampaglia, Prashant Shiralkar, Luis M Rocha, Johan Bollen, Filippo Menczer, and Alessandro Flammini. 2015. Computational fact checking from knowledge networks. *PloS one*, 10(6):e0128193.
- CNN. 2020. Covid-19 misinformation is spreading in 25 different languages, leading to deaths and injuries, study finds. <https://www.cnn.com/2020/08/10/health/covid-misinformation-spread-study-wellness/index.html> [Accessed on June 30, 2020].
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Arkin R Dharawat, Ismini Lourentzou, Alex Morales, and ChengXiang Zhai. 2020. Drink bleach or do what now? Covid-HeRA: A dataset for risk-informed health decision making in the presence of COVID19 misinformation. <https://openreview.net/forum?id=PmY1SNmJIEC> [Accessed on July 1, 2020].
- 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.

- Genevieve Gorrell, Kalina Bontcheva, Leon Derczynski, Elena Kochkina, Maria Liakata, and Arkaitz Zubiaga. 2018. Rumoureal 2019: Determining rumour veracity and support for rumours. *arXiv preprint arXiv:1809.06683*.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don't stop pretraining: Adapt language models to domains and tasks. In *Proceedings of ACL*.
- Hamid Karimi, Proteek Roy, Sari Saba-Sadiya, and Jiliang Tang. 2018. Multi-source multi-class fake news detection. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1546–1557.
- Ramez Kouzy, Joseph Abi Jaoude, Afif Kraitem, Molly B El Alam, Basil Karam, Elio Adib, Jabra Zarka, Cindy Traboulsi, Elie W Akl, and Khalil Baddour. 2020. Coronavirus goes viral: quantifying the COVID-19 misinformation epidemic on Twitter. *Cureus*, 12(3).
- Srijan Kumar, Robert West, and Jure Leskovec. 2016. Disinformation on the Web: Impact, Characteristics, and Detection of Wikipedia Hoaxes. In *Proceedings of the 25th international conference on World Wide Web*, pages 591–602.
- Yaliang Li, Qi Li, Jing Gao, Lu Su, Bo Zhao, Wei Fan, and Jiawei Han. 2015. On the discovery of evolving truth. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 675–684.
- Yunyao Li, Tyrone Grandison, Patricia Silveyra, Ali Douraghy, Xinyu Guan, Thomas Kieselbach, Chengkai Li, and Haiqi Zhang. 2020. Jennifer for COVID-19: An NLP-Powered Chatbot Built for the People and by the People to Combat Misinformation. *ACL 2020 Workshop NLP-COVID*.
- Cristian Lumezanu, Nick Feamster, and Hans Klein. 2012. # bias: Measuring the tweeting behavior of propagandists. In *Sixth International AAAI Conference on Weblogs and Social Media*.
- Liz McQuillan, Erin McAweeney, Alicia Bargar, and Alex Ruch. 2020. Cultural convergence: Insights into the behavior of misinformation networks on twitter. *arXiv preprint arXiv:2007.03443*.
- Marcelo Mendoza, Barbara Poblete, and Carlos Castillo. 2010. Twitter under crisis: Can we trust what we RT? In *Proceedings of the first workshop on social media analytics*, pages 71–79.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. SemEval-2016 Task 6: Detecting Stance in Tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41.
- Martin Müller, Marcel Salathé, and Per E Kummervold. 2020. COVID-Twitter-BERT: A Natural Language Processing Model to Analyse COVID-19 Content on Twitter. *arXiv preprint arXiv:2005.07503*.
- New York Times. 2020. Coronavirus doctors battle another scourge: Misinformation. <https://www.nytimes.com/2020/08/17/technology/coronavirus-disinformation-doctors.html> [Accessed on June 30, 2020].

- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Gordon Pennycook, Jonathon McPhetres, Yunhao Zhang, and David Rand. 2020. Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy nudge intervention. *PsyArXiv Preprints*, 10.
- Poynter. 2020. The coronavirusfacts/datoscoronavirus alliance database. <https://www.poynter.org/ifcn-covid-19-misinformation/> [Accessed on June 30, 2020].
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings Using Siamese BERT-Networks. *arXiv preprint arXiv:1908.10084*.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond accuracy: Behavioral testing of nlp models with checklist. *arXiv preprint arXiv:2005.04118*.
- Juan Carlos Medina Serrano, Orestis Papakyriakopoulos, and Simon Hegelich. 2020. NLP-based Feature Extraction for the Detection of COVID-19 Misinformation Videos on YouTube. *ACL 2020 Workshop NLP-COVID*.
- Shivade. 2019. MedNLI - A Natural Language Inference Dataset For The Clinical Domain. doi.org/10.13026/C2RS98.
- 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.
- Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake news detection on social media: A data mining perspective. *ACM SIGKDD explorations newsletter*, 19(1):22–36.
- Kai Shu, Suhang Wang, and Huan Liu. 2019. Beyond news contents: The role of social context for fake news detection. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 312–320.
- Eugenio Tacchini, Gabriele Ballarin, Marco L Della Vedova, Stefano Moret, and Luca de Alfaro. 2017. Some like it hoax: Automated fake news detection in social networks. *arXiv preprint arXiv:1704.07506*.
- TechCrunch. 2020. Facebook, reddit, google, linkedin, microsoft, twitter and youtube issue joint statement on misinformation. *TechCrunch*. <https://techcrunch.com/2020/03/16/facebook-reddit-google-linkedin-microsoft-twitter-and-youtube-issue-joint-statement-on-misinformation/> [Accessed on June 30, 2020].
- James Thorne, Mingjie Chen, Giorgos Myriantous, Jiashu Pu, Xiaoxuan Wang, and Andreas Vlachos. 2017. Fake news stance detection using stacked ensemble of classifiers. In *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*, pages 80–83, Copenhagen, Denmark. Association for Computational Linguistics.

- James Thorne and Andreas Vlachos. 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. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2018a. The fact extraction and verification (fever) shared task. *arXiv preprint arXiv:1811.10971*.
- 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)*.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2019. The FEVER2.0 Shared Task. In *Second Workshop on Fact Extraction and VERification (FEVER)*, pages 1–6.
- Sebastian Tschiatschek, Adish Singla, Manuel Gomez Rodriguez, Arpit Merchant, and Andreas Krause. 2018. Fake news detection in social networks via crowd signals. In *Companion Proceedings of the The Web Conference 2018*, pages 517–524.
- Svitlana Volkova, Kyle Shaffer, Jin Yea Jang, and Nathan Hodas. 2017. Separating facts from fiction: Linguistic models to classify suspicious and trusted news posts on twitter. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 647–653.
- G VolppKevin et al. 2020. Asked and answered: Building a chatbot to address covid-19-related concerns. *NEJM Catalyst Innovations in Care Delivery*.
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or fiction: Verifying scientific claims. *arXiv preprint arXiv:2004.14974*.
- William Yang Wang. 2017. "Liar, Liar Pants on Fire": A New Benchmark Dataset for Fake News Detection. *arXiv preprint arXiv:1705.00648*.
- Wei Wei and Xiaojun Wan. 2017. Learning to identify ambiguous and misleading news headlines. *arXiv preprint arXiv:1705.06031*.
- WHO. 2020. Novel coronavirus(2019-ncov) situation report - 13. <https://www.who.int/docs/default-source/coronaviruse/situation-reports/20200202-sitrep-13-ncov-v3.pdf> [Accessed on June 30, 2020].
- Wikipedia. 2020. Misinformation Related to the COVID-19 pandemic. https://en.wikipedia.org/wiki/Misinformation_related_to_the_COVID-19_pandemic [Accessed on June 30, 2020].
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *North American Chapter of the Association for Computational Linguistics (NAACL)*.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. BERTScore: Evaluating Text Generation with BERT. *arXiv preprint arXiv:1904.09675*.

Arkaitz Zubiaga, Ahmet Aker, Kalina Bontcheva, Maria Liakata, and Rob Procter. 2018. Detection and resolution of rumours in social media: A survey. *ACM Computing Surveys (CSUR)*, 51(2):1–36.