

# SciLander: Mapping the Scientific News Landscape

Maurício Gruppi<sup>1</sup>, Panayiotis Smeros<sup>2</sup>, Sibel Adalı<sup>1</sup>, Carlos Castillo<sup>3</sup>, Karl Aberer<sup>2</sup>

<sup>1</sup> Rensselaer Polytechnic Institute, USA

<sup>2</sup> École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

<sup>3</sup> Universitat Pompeu Fabra, Spain

gouvem@rpi.edu, panayiotis.smeros@epfl.ch, adalis@cs.rpi.edu, chato@acm.org, karl.aberer@epfl.ch

## Abstract

The COVID-19 pandemic has fueled the spread of misinformation on social media and the Web as a whole. The phenomenon dubbed ‘infodemic’ has taken the challenges of information veracity and trust to new heights by massively introducing seemingly scientific and technical elements into misleading content. Despite the existing body of work on modeling and predicting misinformation, the coverage of very complex scientific topics with inherent uncertainty and an evolving set of findings, such as COVID-19, provides many new challenges that are not easily solved by existing tools. To address these issues, we introduce SciLander, a method for learning representations of news sources reporting on science-based topics. SciLander extracts four heterogeneous indicators for the news sources; two generic indicators that capture (1) the copying of news stories between sources, and (2) the use of the same terms to mean different things (i.e., the semantic shift of terms), and two scientific indicators that capture (1) the usage of jargon and (2) the stance towards specific citations. We use these indicators as signals of source agreement, sampling pairs of positive (similar) and negative (dissimilar) samples, and combine them in a unified framework to train unsupervised news source embeddings with a triplet margin loss objective. We evaluate our method on a novel COVID-19 dataset containing nearly 1M news articles from 500 sources spanning a period of 18 months since the beginning of the pandemic in 2020. Our results show that the features learned by our model outperform state-of-the-art baseline methods on the task of news veracity classification. Furthermore, a clustering analysis suggests that the learned representations encode information about the reliability, political leaning, and partisanship bias of these sources.

## Introduction

The COVID-19 pandemic has resulted in a significant increase in information production and consumption at the same time. With this came a large increase in unreliable information, dubbed ‘infodemic’ (Buchanan 2020). This increase was also coupled with the growing scrutiny of media sources and purposeful amplification of any errors they made. As the readers sought correct, timely, and trustworthy information, many news and media sources worked hard to discredit others and create confusion (Van Bavel et al. 2020).

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

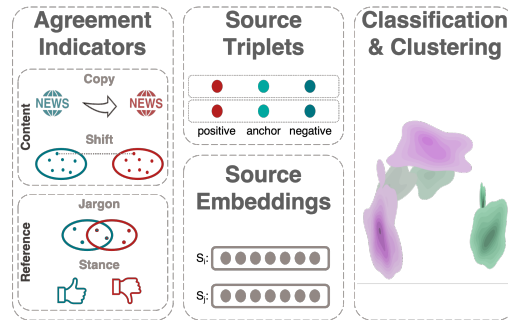


Figure 1: Overview of SciLander, including agreement indicator extraction (§ & §), triplet sampling and unsupervised source embeddings training (§), and evaluation on the downstream tasks of classification and clustering (§).

Governments and public health agencies have the responsibility to respond to the crisis and protect the public from misinformation by utilizing the power of social and news media (Castillo 2016). Yet, the same social and news media work as a catalyst for the infodemic, allowing disinformation to be dispersed on a large scale, regardless of the significant effort to hinder its spread (McKay and Tenove 2021).

Despite the existing body of work on modeling and predicting misinformation, coverage of a complex scientific topic with inherent uncertainty and evolving set of findings, such as COVID-19, provides many new challenges that are not easily solved by existing tools (Zarocostas 2020). On the article level, the evaluation of news stories may be challenging as they may contain information that cannot be easily verified. Furthermore, many sources may not have the necessary staffing for the proper communication of science-related topics, they may be known to have published incorrect information, this information may also have changed over time, or the source may have later corrected it.

Often, language-based methods fail in such a task because different sources may use the same terms to mean different things. Furthermore, many sources may use scientific references to back up their claims; however, the validity of these references is not easily verifiable. Being able to map out the consequential and systematic patterns of behavior of such sources in terms of both *content* and *references* would

be particularly useful in such scenarios (Chung, Nam, and Stefanone 2012). It would allow sources to be compared to other known sources in terms of their coverage, and develop explanations to the aspects in which they are similar to or different from each other.

To address these challenges, we introduce a novel method called SciLander. SciLander builds on a set of novel features, based on the deep processing of news articles published by a set of sources, producing a vector representation of these news sources. To build this, we incorporate measures of similarity and difference between the sources based on their citation behavior, the republishing of articles from each other, and their general language usage. In particular, we use the coverage of COVID-19 to show that this embedding has many desirable features that can help multiple downstream tasks.

**Our Contribution.** The technical contributions we introduce are the following:

- We propose four news agreement indicators for sources: i) the shared content or republished articles, ii) the semantic shift of terms in the common vocabulary, iii) the usage of scientific jargon, and iv) the citation stance of the news sources (§ & §);
- We combine these indicators in a unified framework for training unsupervised news source embeddings (§);
- We evaluate our method using a dataset of news publications related to COVID-19. Sources in this dataset are labeled with respect to reliability and political leaning;
- We compare our method to strong baselines on the problem of veracity classification of news sources and show a significant gain in performance when combining the indicators proposed in this paper;
- We test the applicability of our method in an online learning experiment, showing that it can be used to learn features from sources even if little data is available or if new coming sources are presented in the landscape;
- We show that the learned features encode information about the sources' reliability level, partisanship bias, and political leaning through a clustering analysis experiment.

## Related Work

We distinguish three levels of granularity for misinformation in news and social media: claims, articles, and sources. This is a broad research area where results are scattered through multiple disciplines and venues; below we present studies relevant to each of the three aforementioned levels.

**Claims & Articles Veracity.** Many of the computational methods for veracity assessment of news articles employ machine learning techniques in supervised binary or multi-label classification settings (Baly et al. 2019; Reis et al. 2019; Zhang et al. 2018; Yang et al. 2019; Vishwakarma, Varshney, and Yadav 2019). In this setting, a news article is given as the input to a model and it must predict whether the article contains false information. Other studies aim at detecting the veracity of information at a more granular level by working with claims and rumors (Zubiaga et al. 2018; Shaar et al. 2020; Hansen et al. 2019; Jiang et al. 2020;

Smeros, Castillo, and Aberer 2021). This approach consists in detecting fragments of text, e.g., sentences or paragraphs, worthy of fact-checking. Thus, a single document or news article may contain several claims, some of which may be inaccurate or deceiving.

**Source Veracity.** Source-based approaches are holistic approaches that evaluate the quality of a news source as a whole, without focusing on individual claims or articles extracted from it. Baly et al. (2018, 2019); Li and Goldwasser (2019) highlight the importance of features beyond text to evaluate the veracity of news sources, such as the presence in social media and the existence of a Wikipedia page about a source. Furthermore, Shu, Wang, and Liu (2019) explore the interactions between users, authors, and sources, while Gruppi, Horne, and Adalı (2021) observe content sharing trends among news publishers. Finally, Bourgeois, Rappaz, and Aberer (2018); Rappaz, Bourgeois, and Aberer (2019) study the selection bias in the topic coverage of news sources by exploring the co-references of these sources to the same news events, while Ribeiro et al. (2018) infer the biases of news sources by utilizing their advertiser insights into the demographics of their social media audience.

Both claim- and article-level veracity assessments require data labeling at a very large scale (e.g., individual claims or articles labeled as *reliable* or *unreliable*) and heavily rely on text-specific features these short pieces of text provide. Our approach is, to the best of our knowledge, the first approach that aggregates information about the writing style and citation behavior of news sources to learn unsupervised source representations, that is aware of the science-related content published by them.

## Corpus

Our study targets the reliability of sources when reporting news related to science. Thus, we use a corpus of news articles targeted on the emerging scientific topic of COVID-19, and a corpus of scientific references, also targeted on COVID-19. We summarize the basic statistics of both corpora in Table 1.

**NELA-GT-2020.** The collection of news articles contains a total of 1.78 million articles published by 519 sources (Gruppi, Horne, and Adalı 2021). Each article in the dataset contains a title, full text, name of the publishing source, and publication timestamp. We use a subset containing only articles related to COVID-19, resulting in 991,116 news articles from 493 sources, published over 18 months, between January 1st 2020 and July 1st 2021. We obtain this subset by applying keyword-based filtering using the COVID-19 terminology from Shugars et al. (Shugars et al. 2021), selecting articles that contain at least one COVID-related keyword in the title or body text.

**Media Bias/Fact Check Labels.** We retrieve labels for sources in the corpus from the news assessment agency Media Bias/Fact Check<sup>1</sup>. We obtain the *political leaning* of news sources, represented by direction (left or right) and magnitude (mild, moderate, extreme). These are encoded

<sup>1</sup><https://mediabiasfactcheck.com>

NELA-GT-2020	
Total Articles	~1.8M
COVID-19 Articles	~1M
Total Sources	493
Labeled Sources	316
Reliable Sources	122
Unreliable Sources	194
Partisan Sources	162
Scientific References	
COVID-19 Papers (CORD-19)	~300K
Scientific Domains (SciLens)	~1K
References in NELA-GT-2020	~200K

Table 1: Summary of the used corpora. We see that more than half of the articles in NELA-GT-2020 are related to the topic of COVID-19. The labels for reliable, unreliable and partisan sources are obtained from Media Bias/Fact Check.

as integer numbers in  $[-3, 3]$ , negative values indicate left-bias, positive values indicate right bias, and 0 represent center sources. Furthermore, we obtain a *conspiracy-theory* label, a binary indicator denoting whether a source publishes conspiracy theories and/or pseudoscience content. These are often highly unreliable sources and may or may not exhibit political leaning. Finally, we obtain *factual reporting*, an integer score from 0 to 5 assigned to each source, where 0 indicates the least credible score and 5 is the most credible score. A source that constantly publishes misleading content, fails to fact-check its publications, and does not disclose an editorial board tends to be associated with a lower *factual reporting* score.

Based on the *factual reporting* score, we divide news sources into two reliability classes, namely the *Reliable News Sources* and the *Unreliable News Sources*. The rules defining each class are described as follows:

- *Reliable News Sources*: sources whose *factual reporting* score is greater than 2.
- *Unreliable News Sources*: sources flagged as conspiracy-theory news producers or sources whose *factual reporting* score is less than or equal to 2.

**Scientific References.** We enhance the news collection described above, by extracting the external scientific references of news articles, i.e., the outgoing hyperlinks from the main body of the news articles. We also extract the context of each reference, i.e., the passage of the news article that surrounds this reference. We consider two repositories of references provided by *CORD-19* and *SciLens*.

One of the most prominent collection of papers related to *COVID-19*, consisting of peer-reviewed papers as well as preprints and other historical coronavirus research, is *CORD-19* (Wang et al. 2020). We use the *2021-06-14* release of *CORD-19*, which contains a total of 310,833 papers.

The second source of scientific references comes from *SciLens* (Smeros, Castillo, and Aberer 2019). *SciLens* provides a list of the top-1000 university domains (as indicated by CWUR.org), enhanced with a manually curated list of open-access publishers and grey literature databases. Indeed, these scientific references are more prevalent in news



Figure 2: Example of a subgraph of the Content Sharing Network where nodes, representing sources, are connected by directed edges denoting the direction of the copied content between sources. Node color indicates the reliability class of the source (green for *Reliable*, purple for *Unreliable*), and edge width indicates the amount of content copied.

than the *CORD-19* papers, because their writing style and terminology used is typically more oriented towards a non-expert audience.

## Content Indicators

In this section, we introduce two content-based indicators that we use to align news sources. Particularly, we introduce an indicator regarding the shared content and an indicator regarding the semantic shift of terms between sources.

### Copy Indicator

Content Sharing Network (CSN) is a model of content replication by sources in the news landscape. The sharing of news articles has been shown to be a common factor between news sources that adopt similar narratives around certain topics, which also correlates with the credibility of these sources (Horne, Nørregaard, and Adalı 2019). Figure 2 illustrates how sources are related in a CSN, where articles are copied from source to source.

The CSN is modeled as a directed graph where nodes represent news sources and edges indicate sources that copy articles verbatim from one another. Edges weights are proportional to the amount of content copied between the connected sources. The adjacency matrix  $C$  of such network represents the affinity between the news sources. We obtain this matrix using the method proposed by Horne et al. (2019) which consists of computing document vector representations for news articles using a TF-IDF bag-of-words representation. Articles are considered verbatim copies of each other if the cosine similarity between their vectors is greater than a threshold of 0.85, and the direction of the copying is determined by the publication date of the article. The similarity threshold is defined following the recommendations from Horne, Nørregaard, and Adalı (2019).

The final adjacency matrix is obtained by aggregating all copied articles at the source level. Thus, a directed edge from node  $i$  to  $j$  exists if source  $j$  copies articles from source  $i$ . The complement of the degree of relatedness distance between sources  $i$  and  $j$ , is given as a function of the weight of the edge  $(i, j)$  and is defined as:

$$d_{copy}(i, j) = 1 - \frac{|A_i \cap A_j|}{|A_j|}$$

Source		Usage
Modern Mama	Alternative	{...} <b>these specific herbs have strong antiviral actions</b> , including against other strains of coronavirus.
Healthy Living	Holistic	{...} Garlic is known to have potent antibacterial, <b>antiviral</b> , antifungal and antiprotozoal abilities.
The Guardian		{...} overwhelming emergency departments and causing governments to overspend on <b>antiviral medications</b> .
The Post	Washington	{...} although the <b>antiviral drug</b> remdesivir has been shown to help some patients {...}

Table 2: Semantic shift of the term “antiviral”. We observe a contextual shift of the word. In the top two cases, the term is used to describe alternative medicine with herbs, while in the bottom two cases, the term is used with its ordinary (scientific) connotation.

where  $A_i$  and  $A_j$  are articles published by sources  $i$  and  $j$ ; thus, their intersection should contain articles from source  $i$  copied by source  $j$ . The value of  $d_{cpy}$  increases as fewer articles are copied from  $i$  to  $j$  and decreases as more articles in  $j$  are copied from  $i$ .

### Shift Indicator

We analyze how specific technical terms are used differently between news sources. Different uses of a certain term in two pieces of text can occur if that same term is used in a different context in each of the texts. Semantic shift is the process through which the usage of a given word drifts when compared across different sources. Specifically, we consider the lexical semantic shift, which posits the semantics of a word to be defined by its contextual relationships to other lexicons (Cruse et al. 1986). We argue that significant contextual shifts of topic-related words may serve as a signal of source disagreement, i.e., two sources using a certain target word in significantly different contexts may indicate that they use such words with different intents. An illustrative example is shown in Table 2. Note, in both examples, the word antiviral is still used to indicate “something that is effective against viruses”; however, the contexts give different connotations to what the antiviral product is.

Semantic shift has been used extensively in computational linguistics studies of language evolution (Hamilton, Leskovec, and Jurafsky 2016) and, more recently, in studies quantifying the linguistic differences across domains (Yin, Sachidananda, and Prabhakar 2018; Schlechtweg et al. 2019). In our method, we use semantic shift as an indicator of agreement among sources as it helps to uncover unique narratives created by unreliable sources, especially those based on conspiracy theories, deviating significantly from the narratives from reliable media.

The semantic shift between two sources  $i$  and  $j$  is measured by the deviation in the usage of words they have in common. Specifically, we define semantic shift as the agree-

gated distance between word embeddings for terms in the common vocabulary of sources  $i$  and  $j$ . However, because the word embeddings are trained independently from each other, they cannot be directly compared. For example, suppose that  $v_a$  and  $v_b$  denote word vectors for the word *virus* learned from the sources *The Washington Post* and *Global Research*, respectively. The cosine distance  $d_{cos}(v_a, v_b)$  is not meaningful unless we first create a mapping between the embedding spaces of each source. This mapping can be achieved by applying an orthogonal transformation to one of the embedding spaces to minimize the sum of the pairwise Euclidean distances between word vectors of the common vocabulary. Being orthogonal means that this transformation preserves the inner product of the embeddings in the transformed space; for that reason, this mapping is also called embedding *alignment* (Hamilton, Leskovec, and Jurafsky 2016; Joulin et al. 2018).

Finding the best alignment of two embedding spaces is not a trivial task. Learning a transformation from all the words in the common vocabulary is often undesired, as the objective of the mapping is to minimize the distance between every pair of word vectors, hence minimizing the distance between words that are potentially semantically distinct (Yin, Sachidananda, and Prabhakar 2018). To learn alignments between word embeddings, we employ the state-of-the-art self-supervised semantic shift (S4) method (Gruppi, Chen, and Adali 2021), which is designed to select the best words for generating a mapping between two embeddings. This procedure is applied to embeddings trained using Word2Vec (Mikolov et al. 2013).

Once we train and align the embeddings, we compute the semantic distance between sources  $i$  and  $j$  as the average cosine distance between the top 10% most frequent words in  $i$  and  $j$  (stop words excluded). Thus, the distance between sources  $i$  and  $j$  is defined as:

$$d_{sem}(i, j) = \frac{\sum_{v \in V_i \cap V_j} \cos(\text{emb}_i(v), \text{emb}_j(v))}{|V_i \cap V_j|}$$

where  $V_i$  and  $V_j$  are the vocabularies of sources  $i$  and  $j$ ,  $\text{emb}_i(v)$  and  $\text{emb}_j(v)$  compute the embeddings representation of word  $v$ , and  $\cos$  computes the cosine distance between the embeddings. Additionally,  $V_i$  and  $V_j$  may be replaced with subsets of the common vocabulary to avoid using every word in the analysis (e.g., filter for the most frequent words).

### Reference Indicators

In this section, we introduce the reference indicators that we used to align news sources. Particularly, we introduce two dedicated scientific indicators, namely, the usage of scientific jargon and the citation stance. These indicators are *reference indicators*, i.e., they define a distance among sources given a common (scientific) reference.

### Reference Context Extraction

To compute the *reference indicators*, we need the textual context of the references, i.e., the paragraph in which these references are cited. To extract this context, we: i) locate the

Source	Reference Context
TheNewYorker	In June, just three months into a historic health crisis, a survey by the Center for Disease Control and Prevention found that forty per cent of Americans were already struggling with at least one <b>mental-health issue</b> .
RedState	It is no wonder that many Americans have <b>lost their faith</b> throughout 2020. Too many leaders have been inconsistent in their actions minus their continued breaches of the public trust.
Reference	Title
CDC	<b>Mental Health</b> , Substance Use, and Suicidal Ideation During the COVID-19 Pandemic — United States, June 24–30, 2020.

Table 3: Usage of scientific jargon when citing a report by *CDC* (Czeisler et al. 2020). We highlight that the citation context of *TheNewYorker* is semantically closer to the report than the citation context of *RedState*.

references by parsing the raw HTML page of each news article of our data collection, and ii) traverse the structural tree of the page to discover the most fine-grained text passage that contains the reference. Currently, we do not support end-notes within articles, i.e., anchors at the bottom of articles where all the scientific references are listed, because it is a journalistic practice rarely appearing in our corpus.

### Jargon Indicator

This indicator quantifies the scientific nature of the context in which a reference is used. To estimate this indicator, we need a lexicon of terms (*jargon\_terms* in the following) that are considered jargon in the scientific domain of our corpus. Since, as we explain in §, our corpus contains news articles related to COVID-19, we use the vocabulary of *CDC A-Z Index*<sup>2</sup>, manually enhanced with common COVID-19 terminology. After applying standard cleaning (e.g., punctuation removal), we compute the following distance:

$$d_{jar}(i, j) = |ctx_r(i) \cap ctx_r(j) \cap jargon\_terms|$$

where  $ctx_r(i)$  and  $ctx_r(j)$  are the terms in the citation contexts of sources  $i$  and  $j$  for each common reference  $r$ .

We note that we do not aggregate for all common references between sources  $i$  and  $j$ ; hence, we do not limit to a single distance between these sources. In this way, we encode the co-citation volume between sources  $i$  and  $j$ , which is useful for our triplet sampling strategy (details in ). After computing  $d_{jar}(i, j)$ , we apply *Min-Max Normalization* in the interval  $[0, 1]$  to comply with the previously-defined distances. As we observe in Table 3, even such a simplistic metric is able to capture cases in which news sources completely distort the scientific message of the cited reference.

### Stance Indicator

This indicator quantifies the sentiment charge of the context in which a reference is cited. To measure this sentiment

<sup>2</sup><https://www.cdc.gov/az>

Source	Reference Context
FiveThirtyEight	{...} based on current CDC guidelines {...} experts said that undercounting (deaths) was still more likely than overcounting.
The Truth About Cancer	Perhaps worst, the <b>CDC has continued to lie</b> about the death count by artificially inflating it. CDC guidelines for determining COVID-19 deaths include: Anyone who tests positive, even if they died from other causes. Anyone who had COVID-19 symptoms, even if they aren't tested.
Reference	Title
CDC	Guidance for Certifying Deaths Due to Coronavirus Disease 2019 (COVID-19).

Table 4: Stance of news sources when citing (using the underlined hypertext) a webinar by *CDC* (Anderson et al. 2020). We highlight that *The Truth About Cancer* uses more emotionally loaded words than *FiveThirtyEight*.

charge, we use the *Multi-Genre Natural Language Inference* model *BART* for zero-shot classification (Lewis et al. 2020). This model<sup>3</sup> computes the probability that we infer a certain *hypothesis* given a *premise*. Thus, the model needs no explicit training on the downstream task of stance classification since the desired classes are provided implicitly in the *hypothesis*. After experimenting with various templates for *premise* and *hypothesis*, we report the ones that yield the most reliable results:

*premise* = **reference context**

*hypothesis* = “The stance of this example is **negative**”

The output of this model is a value in the interval  $[0, 1]$ , denoting the probability a given premise implies our hypothesis. We note that, by using this premise and hypothesis, we treat *neutral* and *positive* stances similarly, i.e., as *non-negative* stances, because we want to highlight extremely negative stances (Table 4). Using this model we compute the following distance:

$$d_{ref}(i, j) = |stance(ctx_r(i)) - stance(ctx_r(j))|$$

where  $stance(\cdot)$  computes the stance of the citation contexts of sources  $i$  and  $j$  for each common reference  $r$ .

Similarly as above, after computing  $d_{ref}(i, j)$ , we apply *Min-Max Normalization* in the interval  $[0, 1]$ . As we observe in Table 4, this indicator distinguishes between the sentiment of sources towards a common reference.

## Unsupervised Source Embeddings

The previous section described the heterogeneous indicators that we extract from each news source. In this section, we describe how we combine these indicators in a unified framework to learn unsupervised representations of news sources. The triplet loss function aims at coupling different parts of the input spaces (here, our indicators) into a single representation (Weinberger and Saul 2009). The triplets sampling

<sup>3</sup><https://huggingface.co/facebook/bart-large-mnli>

and embeddings training methods employed in this framework are well-established methods (Hoffer and Ailon 2015) used mainly in learning-to-rank recommendation systems (Chen et al. 2016; Wang et al. 2021).

### Triplets Sampling

Our goal is, using the distances defined by the indicators, to discover pairs of similar sources and pairs of dissimilar sources. By joining these two sets of pairs, we create triplets of the form (*anchor*, *positive*, *negative*), where *anchor* is the common element of the pairs, *positive* is the element similar to the *anchor*, and *negative* is the element dissimilar to the *anchor*. For simplicity, in the following, we will refer to these triplets as (*a*, *p*, *n*).

We note that these triplets may not occur from the same indicator, i.e., the positive pair may occur from an indicator that is more appropriate for capturing the affinity between sources, and the negative pair may occur from an indicator that is more appropriate for capturing the disparity between sources. In our experimental evaluation (§), we evaluate each indicator in its ability to produce good positive and negative pairs as well as full triplets.

**Positive Pair Sampling.** We use the distances computed for each indicator to generate pairs of similar sources. For all indicators we introduce in § & §, short distance denotes similarity. Given an indicator *f* (*copy*, *shift*, *jargon*, or *reference*), we generate a positive pair of similar sources *i*, *j* with a probability inversely proportional to the distance between *i* and *j*:

$$pp_f(i, j) = \frac{d_f^{-1}(i, j)}{\sum_k d_f^{-1}(i, k)} \forall j \neq i$$

We draw *l* positive samples from this distribution for each indicator and each source in the dataset, producing a total of *l* positive source pairs (*a*, *p*).

**Negative Pair Sampling.** For negative sampling, we employ two strategies. For some indicators (e.g., the stance indicator), a large distance between sources denotes opposing sentiment, thus disagreement (e.g., the sources in Table 4). Hence, we use the inverse distribution we used for generating positive pairs to generate negative pairs:

$$np_f(i, j) = 1 - pp_f(i, j) \forall j \neq i$$

Similarly as above, we draw *l* negative samples from this distribution for each indicator and each source in the dataset, producing a total of *l* negative source pairs (*a*, *n*).

Nonetheless, there are indicators (e.g., the copy indicator) for which a large distance between sources does not necessarily denote disagreement; it only denotes the absence of agreement. In these cases, we draw the negative pairs uniformly from the set of sources.

Finally, we employ a cleaning heuristic to increase the accuracy of our triplets (detailed experiment in §). Specifically, we make sure that we do not select a negative pair (*a*, *n*) which we have already selected as positive pair (*a*, *p*):

$$(a, p) \wedge (a, n) \Rightarrow p \neq n$$

### Embeddings Training

Once we extract all the triplets, we use them for training a dense representation model for news sources with the *Triplet Margin Loss* (Balntas et al. 2016). The learning objective of *Triplet Margin Loss* is to minimize the distance between an anchor and a positive sample while maximizing the distance between the anchor and the negative sample.

The procedure we employ is the following. First, we initialize the embeddings for all the sources into a low-dimensional, dense vector space by randomly setting the weights in the embedding layer following a normal distribution  $\mathbb{N}(0, 1)$ . Then, given the input triplets (*a*, *p*, *n*), we train these embeddings by minimizing the loss function *L*:

$$L(a, p, n) = \max\{d(a, p) - d(a, n) + M, 0\}$$

where *d* is the distance function, and *M* is the margin parameter that controls the gap between positive and negative distances. The larger *M* is, the larger is the gap between *d(a, p)* and *d(a, n)*. We train the embeddings over several epochs until convergence and then use them as the representation of the news sources.

The parameters of this method are the margin *M*, the distance function *d*, and the size of the output vectors *s*. We release the optimal training parameters as well as the trained sources embeddings in our code release (§).

## Experiments

Our experimental evaluation is three-fold; first, we evaluate the indicators individually, then we evaluate the source embeddings on the downstream task of source reliability classification, and finally, we perform an unsupervised clustering where we analyze the patterns in the news sources captured by the learned features. In the following experiments, the labels from Media Bias/Fact Check are used, and word embeddings for the semantic shift are trained using Word2Vec with dimension 100, context window of 10, and minimum word count of 20. The parameters for SciLander are margin *M* = 1, vector size *s* = 50, and distance *d* used in the loss function is the cosine distance.

### Indicator Coverage

In our first experiment, we measure the overlap of the introduced indicators in terms of source and triplet coverage. We also measure the accuracy of the triplets computed by these indicators.

We define the source coverage (*sc*) and the triplet coverage (*tc*) between two indicators *i*, *j* as follows:

$$sc(i, j) = \frac{|src(i) \cap src(j)|}{|src(i)|}, tc(i, j) = \frac{|trpl(i) \cap trpl(j)|}{|trpl(i)|}$$

where *src*(.) and *trpl*(.) compute the distinct set of sources and triplets covered by a given indicator. We note that *sc* and *tc* are non-symmetric; consequently, the heatmaps in Figure 3 are also non-symmetric.

To measure the accuracy of the computed triplets, we use the metric *Area Under the Receiver Operating Characteristics* (AUROC), which measures the *True Positive Rate* over the *False Positive Rate*. We also break down the AUROC of the triplets into i) the *AUROC<sub>p</sub>* of the positive part of



the triplets  $(a, p)$ , ii) the  $AUROC_n$  of the negative part of the triplets  $(a, n)$ , and iii) the  $AUROC_f$  of the full triplets  $(a, p, n)$ . Specifically, for each individual AUROC, we consider the following as true positives:

$$\begin{aligned}
 AUROC_p &: \{(a, p) \text{ s.t. } label(a) = label(p)\} \\
 AUROC_n &: \{(a, n) \text{ s.t. } label(a) \neq label(n)\} \\
 AUROC_f &: \{(a, p, n) \text{ s.t. } label(a) = label(p) \\
 &\quad \wedge label(a) \neq label(n)\}
 \end{aligned}$$

As we observe in Figure 3, although the sources covered by some indicators heavily overlap, the contributed triplets are quite unique. Indicatively, the stance indicator covers 27.5% of the sources, totally overlapping with the copy indicator. However, the contributed triplets of the stance indicator are different from the contributed triplets of all the other indicators and also more accurate. Indeed, we see that there is a trade-off between the source coverage of the indicators and the AUROC. Hence, the more specific the indicator is (e.g., the stance indicator), the better AUROC it has.

Finally, we observe that the overall AUROC for positive and negative pairs ( $AUROC_p$  and  $AUROC_n$ , respectively) are above the 50% baseline of a random positive (or negative) pair selection is truly positive (or negative).

It should be noted that the AUROC for complete triplets ( $AUROC_f$ ) is lower than 50%. This happens because the choice of the final triplets involves two independent decisions: the choice of the positive sample, and the choice of the negative sample. As noted above, each choice has a chance of success of 50% if chosen at random. Thus, for a triplet to be correctly selected, the random baseline is that a correct positive pair is chosen *and* a correct negative pair is chosen, which results in a  $0.5 \times 0.5 = 0.25$ , or 25% baseline chance. As we see in the following experiments, the model for training source embedding is robust to noisy triplets as it yields highly accurate results in all the downstream tasks we use it.

## Offline Source Classification

In this experiment, we evaluate the computed embeddings on a downstream classification task. We assume that, for all sources in our corpus, we have (offline) access to a significant fraction of their history of published articles.

**Baselines.** For this task, we implement baselines using *Stylistic Text Features*, *Contextualized Embeddings*, and *Co-citation Embeddings*, as well as combinations of the above.

*Stylistic Text Features.* We utilize stylistic text features from Horne et al. (Horne and Adali 2017) aggregated at the source level as representations. These features include, among others, the number of: part of speech tags, punctuation symbols, and capitalized words, which are the features that are typically used in news classifiers.

*Contextualized Embeddings.* We compute BERT (Devlin et al. 2019) embeddings for a total of 32 tokens from the title and the opening paragraph of the article, and average them for each source. Similarly, we compute SciBERT (Beltagy, Lo, and Cohan 2019) instead of BERT embeddings, which have been shown to lead to better performance in tasks involving scientific text. The configuration parameters of both

The figure consists of two heatmaps and a summary table. The top-left heatmap shows source coverage overlap between indicators: copy (100.00%, 72.46%, 39.13%, 27.54%), shift (80.91%, 100.00%, 41.42%, 28.16%), jargon (98.54%, 93.43%, 100.00%, 64.23%), and stance (100.00%, 91.58%, 92.63%, 100.00%). The top-right heatmap shows triplet coverage: copy (100.00%, 0.16%, 0.13%, 0.26%), shift (0.11%, 100.00%, 0.02%, 0.03%), jargon (0.79%, 0.17%, 100.00%, 7.03%), and stance (1.38%, 0.25%, 6.23%, 100.00%). The bottom table summarizes the AUROC values for each indicator and overall.

Indicator	$AUROC_p$	$AUROC_n$	$AUROC_f$	#sources
copy	72.7%	51.0%	36.3%	257
shift	61.9%	60.8%	41.8%	308
stance	89.7%	73.3%	68.3%	87
jargon	81.9%	51.0%	42.9%	126
overall	77.0%	69.7%	57.5%	316

Figure 3: Overlap of indicators in terms of source coverage (top left) and triplet coverage (top right); AUROC of the positive part, negative part, and full triplets (bottom). Although the sources covered by most indicators heavily overlap, their triplets are quite unique. Also, there is a trade-off between the source coverage of the indicators and their AUROC.

BERT and SciBERT are those suggested in a widely used release of this model (Wolf et al. 2019).

*Co-Citation Embeddings.* We compute a co-citation graph of sources based on their scientific references. We weight this graph either uniformly for each common reference, or by emphasizing the uniquely used references, using their TF-IDF score. In the overall graph, we run node2vec (Grover and Leskovec 2016) to extract source embeddings.

*Joint Embeddings.* The *Contextualized Embeddings* and the *Co-Citation Embeddings* capture two different modalities of news sources; their content and citation behavior. Thus, we create a joint representation by concatenating the two embeddings. Since the dimensionality of the joint embeddings is high, we apply Principal Component Analysis to reduce it and compare it with other baseline representations.

**Evaluation.** We test the usefulness of the learned representations in the problem of source veracity classification. We use the embeddings computed by i) SciLander trained on all indicators, ii) SciLander trained only on content indicators (shift or copy), and iii) the aforementioned baseline models, to train a Nearest Neighbors classifier in a 10-fold cross-validation setting. Figure 4 shows the F1 score of each model for increasing values of  $k$ .

Relying uniquely on textual features limits classifiers to a restricted set of signals. Our framework combines stylistic, semantic, and behavioral indicators to produce a representation that improves the separation of reliable and unreliable sources. Thus, compared to traditional baselines such as stylistic features or features extracted by BERT, our embeddings show significant performance improvement. Our method obtains the best F1 score (87%) for  $k = 37$ .




Figure 4: F1 scores using  $k$ -nearest neighbors classifiers over the source embeddings representations computed by SciLander and the various baselines described in §. SciLander obtains the best F1 score (87%) for  $k=37$ .

### Online Source Classification

In this experiment, we assume that we have two types of sources: i) offline (known) sources, for which we have access to a significant fraction of their publication history, and ii) online (newcomer) sources, for which we have access to a limited fraction of their publication history. As assessing articles from newcomer sources might be a time-consuming task, we inspect the lowest fraction of articles that is needed to accurately classify these sources.

The procedure that we employ is the following: i) we train embeddings for the *offline sources* (as we explain in §); ii) we freeze these embeddings for the *offline sources*; iii) we train embeddings for *online sources*, in the already shaped by the *offline sources* embeddings space.

We conduct the experiment on a 10-fold cross-validation setting. In Figure 5, we report the learning curve (F1 score) for increasing fractions of articles from newcomer sources in the same classification task described in §. We note that the temporal axis is not in chronological order but sampled randomly from the entire corpus (e.g., we sample articles representing a 3-month publishing activity of an online source from the entire publishing activity of that source). In that way, each temporal interval is independent of external events (e.g., the development of the vaccines), which affects the activity of most sources. As we observe in Figure 5, SciLander is able to reliably ( $F1 > 85\%$ ) classify sources, using only three months of their publishing activity.

### Source Clustering Analysis

We conduct an unsupervised clustering experiment to investigate potential trends revealed by the features learned by SciLander. Using the same embeddings from the previous experiments (50 dimensions,  $M = 1$ ), we apply DBSCAN clustering to the source vectors with the cosine distance as distance metric, minimum distance parameter  $\epsilon = 0.1$  and minimum cluster size  $n = 1$ . The resulting clusters are shown in Figure 6; each of the 7 clusters is shown in different color shades and labeled from  $A$  to  $G$ .

We characterize the clusters quantitatively with respect to the density of unreliable sources, political leaning, and the level of partisanship bias aggregated across the news sources




Figure 5: Learning curve (F1 score) for increasing fractions of articles from newcomer sources. SciLander is able to reliably ( $F1 > 85\%$ ) classify sources using only 3 months of their publishing activity.

within them. For each cluster, we compute the proportion of unreliable sources to the total number of sources in the cluster. Figure 7a shows the density of unreliable sources within each cluster. This result suggests that the source embeddings carry information about source credibility when grouping them, even though credibility labels or related features were unknown to the model during training.

Clusters  $C$  and  $E$  contain no unreliable sources and hold mostly mainstream news sources such as The Washington Post, Vox, National Public Radio (NPR), and the Chicago Tribune. The clusters containing the largest proportions of unreliable sources are the clusters  $A$ ,  $B$ , and  $G$ , and most sources in these clusters are websites that propagate conspiracy theories and promote pseudoscience. Details on the discovered clusters are shown in Table 5.

These results show that the SciLander embeddings are able to group sources based on similar reliability. Multiple clusters of relatively high purity with respect to reliability are created, some reliable (75%-100% reliable sources), some unreliable (0%-30% reliable sources).

We compute the overall political leaning of a cluster by averaging the political leaning scores of the sources within that cluster. Partisanship bias is obtained by the absolute value of leaning, scaled to a value in  $[0, 1]$ , with 0 indicating that there is no partisanship bias in the cluster, and 1 indicating the maximum partisanship bias, where all sources in the cluster exhibit a strong political leaning. The partisanship bias describes the agreement between the political leanings of sources within the cluster, and the magnitude of such leanings. The distribution of political leanings and partisanship bias are shown in Figures 7b and 7c. There is a noticeable disparity between the partisanship bias found in the two biggest unreliable clusters  $A$  and  $B$ . Sources in cluster  $A$  exhibit a strong bias, which is nearly absent in cluster  $B$ . We explore the particularities of these clusters next.

### Different Types of Conspiracy Theories

We observe two clusters with high density of unreliable sources (clusters  $A$  and  $B$ ). Both clusters include many unreliable news sources, and there exist qualitative differences between them, which we describe in this section.

To uncover qualitative differences between sources in clusters  $A$  and  $B$ , we measure the shift in context between






Figure 6: Kernel Density Estimation of the clusters

**Cl. (U) (P) Core Sources**

<b>A</b>	.70	.25	NewsWars, Vets. Today, The D.C. Clothesline
<b>B</b>	.84	.03	Mercola, Healthy Hol. Living, Vaccine React.
<b>C</b>	.00	.11	The Washington Post, Vox, NPR
<b>D</b>	.25	.00	The American Cons., Roll Call
<b>E</b>	.00	.20	Chicago Tribune
<b>F</b>	.12	.03	Wash. Monthly, FiveThirtyEight, Atlantic
<b>G</b>	.80	.00	Ice Age Now

Table 5: (U)nreliability score (proportion of unreliable sources), average (P)artisanship bias score, and core sources (nearest neighbors to the centroid) of the identified clusters.

these clusters and the mainstream cluster  $C$ . Specifically, we computed the semantic shift across clusters of sources by training Word2Vec models  $E_A$ ,  $E_B$ , and  $E_C$  using articles from the core sources of each cluster and using the same hyper-parameters as in the previous experiments. Then, we extract the words with the highest cosine distance between pairs  $(E_C, E_A)$  and  $(E_C, E_B)$  to find the terms that most contribute to the deviation in the news from sources in  $C$  to each of the unreliable clusters  $A$  and  $B$ .

Let  $S_A$  and  $S_B$  be the lists of the 100 words most shifted to  $C$ , from  $A$  and  $B$ , respectively. We find that there is only one word in common between the  $S_A$  and  $S_B$ : “natural”. To characterize the words in both lists, we identify words that refer to people, entities and places, political issues, and health and nutrition. Examples of these words are given below and listed on Table 6.

The largest group of words shifted in cluster  $A$  are related to individuals, entities, places (25%), and political topics (12%). Almost all individuals found are political figures (with a few exceptions). There are only 1.5% of terms related to health and nutrition. Many of these news outlets are conspiracy theory websites such as NewsWars, Veterans Today, and InfoWars. According to a Media Bias/Fact Check analysis<sup>4</sup>, these sites often publish hate-speech-filled content in addition to misleading or false information.

In contrast, the largest group of shifted words was detected in cluster  $B$  (21.5%), with only 2% people and 1% related to political topics. According to MBFC journalists<sup>5</sup>, these sources promote alternative health notions, sell questionable products and supplements, and promote antivaccination positions with pseudoscience-based arguments.

<sup>4</sup><https://mediabiasfactcheck.com/veterans-today>

<sup>5</sup><https://mediabiasfactcheck.com/mercola>

People and Places	Political Terms	Health
Kamala Harris	BLM (Black Lives Matter)	Coronavirus
Bernie Sanders	Patriot	Food
Nancy Pelosi	Voting	Vaccines
Mike Pence	Abortion	Doctors
Alex Jones	Partisan	Mask

Table 6: List of words from clusters  $A$  and  $B$  that are most shifted from the mainstream cluster  $C$ . People and Places, and Political Terms appear as the most shifted words in cluster  $A$ , suggesting that its sources push politically-oriented misinformation, while sources in cluster  $B$  focus more on alternative health solutions.

Based on this, we conclude that while cluster  $A$  is a cluster of mostly politically-unreliable news sources covering COVID-19 stories mixed with other political topics, cluster  $B$  is much more focused on covering alternative medicine-based misinformation with slight political leaning, presumably to appeal to individuals with different political opinions. On these sites, health-based information is often mixed with promotion and affiliate links to sites selling alternative medicine products and supplements. Our method is able to properly distinguish these different types of COVID-19 misinformation, without explicitly training on related features.

## Discussion

**SciLander is a method for embedding news sources.** The results of the experiments (§) show that the representations learned from SciLander outperform other state-of-the-art feature models. Despite the final representation being a set of autoencoded features (i.e., embeddings learned from a neural network), it is directly explained by the product of a combination of the aforementioned indicators.

The applications of these learned features are not restricted to classification tasks. They can be used in any scenario where similarities between news sources are needed, such as in clustering analysis (details in §) and recommendation systems.

SciLander, like most other AI/ML methods, is heavily data-driven. It uses signals found in the text of news articles to infer the relationships between sources. The above can cause SciLander to make biased decisions, especially if the input data is biased towards/against societal groups, such as underrepresented minorities and other vulnerable groups. We argue that SciLander, when deciding what content to recommend or promote, can provide assistance in human decision-making but not replace human judgment.

**The indicators used by SciLander complement each other.** The experiments shown in § demonstrate that the embeddings performed better in classification tasks when all four indicators are combined (copy, shift, jargon, and stance). The latter suggests that the indicators worked in a complementary manner, where a mistake made by one indicator is corrected by the other indicators. Furthermore, as seen in Figure 3, some indicators were better suited to detect negative pairs. For example, the indicator *copy* had a

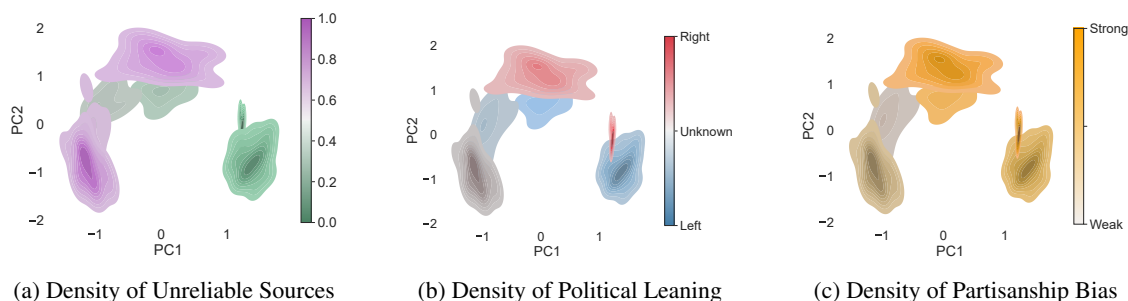


Figure 7: Density analysis of the clusters computed by SciLander. Components PC1 and PC2, obtained from Principal Component Analysis (PCA) on the source embeddings, are the components with the highest explained variance ratio.

*AUROC* score of 72.7% for positive samples and 51% for negative samples. This result can be explained by the fact that while the presence of copy behavior between sources  $x$  and  $y$  is an indicator of similarity between  $x$  and  $y$ , its absence does not necessarily imply that  $x$  and  $y$  are very distinct. In short, source  $x$  not copying from  $y$  does not imply that  $x$  is distant from  $y$ . Conversely, the *stance* indicator had a higher negative pair score (73.3%), suggesting that this indicator would perform better in finding negative samples.

**SciLander has the potential to be extended to general domains.** We based SciLander on features that work as indicators of source similarity or dissimilarity, motivated by previous research on language and misinformation (Chambers and Schilling 2018; Horne, Nørregaard, and Adali 2019; Smeros, Castillo, and Aberer 2019). The *shift* and *copy* indicators are agnostic to the news domain since they only require the presence of text. However, the *stance* and *jargon* indicators are closely tied to scientific news.

To extend the application of SciLander to other, including non-scientific, domains, a topic of choice must be specified prior to the application of the method. The chosen topic must include a set of entities referred to by news sources, such as political figures in the political news domain. In this case, the stance towards scientific references would be replaced by the stance towards such political figures, and the scientific jargon would be replaced by political jargon. Topics can be manually defined via a set of keywords and entities, or automatically defined, such as by applying topic modeling to extract the relevant keywords from the news documents.

**Limitations.** Our methodology was only applied and evaluated on an English dataset; extending to other languages would only require translation/adaptation of the domain-specific lexicon used to compute the jargon indicator or simply skipping this indicator and training using the other three. All the other indicators as well as the introduced embedding model, are based either on language-agnostic or already multilingual models.

Furthermore, our methodology supports only explicit citations, i.e., direct outgoing links to scientific papers, and not implicit mentions of science-related entities (e.g., universities) because the latter design choice introduces ambiguity and noisy source triplets.

Finally, we implicitly filter the scientific references re-

lated to COVID-19 as we filter the news corpus citing these references. Explicit filtering would require downloading and parsing the references from different formats, e.g., *pdf*, which is a demanding task not in the scope of this work.

## Conclusions

We have introduced SciLander, a method for learning a representation of news sources reporting science-related content. Our method uses a combination of signals to estimate the similarity between news sources. We have shown that these signals complement each other, capturing relationships between distinct sets of sources from a dataset of news articles related to COVID-19. Furthermore, the features learned by our model demonstrated superior performance to baselines for the task of source credibility detection, both in an offline and an online setting, requiring as little as three months of publication activity to accurately classify news sources. Lastly, we have shown that the learned source representations encode information of credibility and political leaning, forming clusters of sources that show similar reliability and political bias. In particular, we discovered two large clusters of unreliable sources to which different types of conspiracy news sources flock. One of them concentrates on alternative health misinformation, and the other promotes hyper-partisan political conspiracies.

**Reproducibility.** All the data, code, and models used for this paper are publicly available for research purposes in the following repository: <https://github.com/mgruppi/SciLander>.

## Ethics Statement

Our work aims at finding representations that capture the similarities and differences between news sources in their coverage of the COVID-19 pandemic. Our proposed method bases this representation on the language usage, content copy/sharing behavior, and their stance towards scientific references. We show that the representations learned from these signals are useful for several downstream tasks, including understanding the reliability of a source. This is accomplished by using proxies to trust scientific references, language, and content. One must be careful when applying this method to untested dimensions, such as the presence of language usage by minority groups. These groups may be underrepresented in the training data, which may cause the

model to make biased predictions about them. We propose that this method aids the decision-making process as a complement to human judgment rather than a replacement.

## Acknowledgements

This work was supported by the Rensselaer-IBM AI Research Collaboration (<http://airc.rpi.edu>), part of the IBM AI Horizons Network (<http://ibm.biz/AIHorizons>).

## References

- Anderson, R. N.; Warner, M.; Flagg, L. A.; and Ahmad, F. 2020. Guidance for Certifying Deaths Due to Coronavirus Disease 2019 (COVID-19). <https://stacks.cdc.gov/view/cdc/87029>. Accessed: 2021-03-15.
- Balntas, V.; Riba, E.; Ponsa, D.; and Mikolajczyk, K. 2016. Learning local feature descriptors with triplets and shallow convolutional neural networks. In *BMVC 2016, York, UK, September 19-22, 2016*. BMVA Press.
- Baly, R.; Karadzhov, G.; Alexandrov, D.; Glass, J. R.; and Nakov, P. 2018. Predicting Factuality of Reporting and Bias of News Media Sources. In *EMNLP, Brussels, Belgium, October 31 - November 4, 2018*, 3528–3539. ACL.
- Baly, R.; Karadzhov, G.; Saleh, A.; Glass, J. R.; and Nakov, P. 2019. Multi-Task Ordinal Regression for Jointly Predicting the Trustworthiness and the Leading Political Ideology of News Media. In *NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019*, 2109–2116. ACL.
- Beltagy, I.; Lo, K.; and Cohan, A. 2019. SciBERT: A Pre-trained Language Model for Scientific Text. In *EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, 3613–3618. ACL.
- Bourgeois, D.; Rappaz, J.; and Aberer, K. 2018. Selection Bias in News Coverage: Learning it, Fighting it. In *WWW 2018, Lyon, France, April 23-27, 2018*, 535–543. ACM.
- Buchanan, M. 2020. Managing the infodemic. *Nature Publishing Group*, 16(9): 894–894.
- Castillo, C. 2016. *Big crisis data: social media in disasters and time-critical situations*. Cambridge University Press.
- Chambers, J. K.; and Schilling, N. 2018. *The handbook of language variation and change*. John Wiley & Sons.
- Chen, X.; Qin, Z.; Zhang, Y.; and Xu, T. 2016. Learning to Rank Features for Recommendation over Multiple Categories. In Perego, R.; Sebastiani, F.; Aslam, J. A.; Ruthven, I.; and Zobel, J., eds., *SIGIR 2016, Pisa, Italy, July 17-21, 2016*, 305–314. ACM.
- Chung, C. J.; Nam, Y.; and Stefanone, M. A. 2012. Exploring Online News Credibility: The Relative Influence of Traditional and Technological Factors. *Journal of Computer-Mediated Communication*, 17(2): 171–186.
- Cruse, D. A.; Cruse, D. A.; Cruse, D. A.; and Cruse, D. A. 1986. *Lexical semantics*. Cambridge university press.
- Czeisler, M. É.; Lane, R. I.; Petrosky, E.; Wiley, J. F.; Christensen, A.; Njai, R.; Weaver, M. D.; Robbins, R.; FacerChilds, E. R.; Barger, L. K.; et al. 2020. Mental health, substance use, and suicidal ideation during the COVID-19 pandemic—United States, June 24–30, 2020. *Morbidity and Mortality Weekly Report*, 69(32): 1049.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1*, 4171–4186. ACL.
- Grover, A.; and Leskovec, J. 2016. node2vec: Scalable Feature Learning for Networks. In *SIGKDD, San Francisco, CA, USA, August 13-17, 2016*, 855–864. ACM.
- Gruppi, M.; Chen, P.; and Adali, S. 2021. Fake it Till You Make it: Self-Supervised Semantic Shifts for Monolingual Word Embedding Tasks. In *AAAI 2021, Virtual Event, February 2-9, 2021*, 12893–12901. AAAI Press.
- Gruppi, M.; Horne, B. D.; and Adali, S. 2021. NELA-GT-2020: A Large Multi-Labelled News Dataset for The Study of Misinformation in News Articles. *CoRR*, abs/2102.04567.
- Gruppi, M.; Horne, B. D.; and Adali, S. 2021. Tell Me Who Your Friends Are: Using Content Sharing Behavior for News Source Veracity Detection. *CoRR*, abs/2101.10973.
- Hamilton, W. L.; Leskovec, J.; and Jurafsky, D. 2016. Cultural Shift or Linguistic Drift? Comparing Two Computational Measures of Semantic Change. In *EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, 2116–2121. ACL.
- Hansen, C.; Hansen, C.; Alstrup, S.; Simonsen, J. G.; and Lioma, C. 2019. Neural Check-Worthiness Ranking with Weak Supervision: Finding Sentences for Fact-Checking. In *WWW '19, San Francisco, CA, USA, May 13-17, 2019*, 994–1000. ACM.
- Hoffer, E.; and Ailon, N. 2015. Deep Metric Learning Using Triplet Network. In Feragen, A.; Pelillo, M.; and Loog, M., eds., *SIMBAD 2015, Copenhagen, Denmark, October 12-14, 2015, Proceedings*, volume 9370 of *LNCS*, 84–92. Springer.
- Horne, B. D.; and Adali, S. 2017. This Just In: Fake News Packs a Lot in Title, Uses Simpler, Repetitive Content in Text Body, More Similar to Satire than Real News. *CoRR*, abs/1703.09398.
- Horne, B. D.; Nørregaard, J.; and Adali, S. 2019. Different Spirals of Sameness: A Study of Content Sharing in Mainstream and Alternative Media. In *ICWSM 2019, Munich, Germany, June 11-14, 2019*, 257–266. AAAI Press.
- Jiang, S.; Baumgartner, S.; Ittycheriah, A.; and Yu, C. 2020. Factoring Fact-Checks: Structured Information Extraction from Fact-Checking Articles. In *WWW '20, Taipei, Taiwan, April 20-24, 2020*, 1592–1603. ACM / IW3C2.
- Joulin, A.; Bojanowski, P.; Mikolov, T.; Jégou, H.; and Grave, E. 2018. Loss in translation: Learning bilingual word mapping with a retrieval criterion. *arXiv preprint arXiv:1804.07745*.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *ACL 2020, Online, July 5-10, 2020*, 7871–7880. ACL.

- Li, C.; and Goldwasser, D. 2019. Encoding Social Information with Graph Convolutional Networks for Political Perspective Detection in News Media. In *ACL 2019, Florence, Italy, July 28- August 2, 2019*, 2594–2604. ACL.
- McKay, S.; and Tenove, C. 2021. Disinformation as a Threat to Deliberative Democracy. *Political Research Quarterly*, 74(3): 703–717.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient Estimation of Word Representations in Vector Space. In *ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013*.
- Rappaz, J.; Bourgeois, D.; and Aberer, K. 2019. A Dynamic Embedding Model of the Media Landscape. In *WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, 1544–1554. ACM.
- Reis, J. C. S.; Correia, A.; Murai, F.; Veloso, A.; Benvenuto, F.; and Cambria, E. 2019. Supervised Learning for Fake News Detection. *IEEE Intell. Syst.*, 34(2): 76–81.
- Ribeiro, F. N.; Lima, L. H. C.; Benvenuto, F.; Chakraborty, A.; Kulshrestha, J.; Babaei, M.; and Gummadi, K. P. 2018. Media Bias Monitor: Quantifying Biases of Social Media News Outlets at Large-Scale. In *ICWSM 2018, Stanford, California, USA, June 25-28, 2018*, 290–299. AAAI Press.
- Schlechtweg, D.; Hätyy, A.; Tredici, M. D.; and im Walde, S. S. 2019. A Wind of Change: Detecting and Evaluating Lexical Semantic Change across Times and Domains. In *ACL 2019, Florence, Italy, July 28- August 2, 2019*, 732–746. ACL.
- Shaar, S.; Babulkov, N.; Martino, G. D. S.; and Nakov, P. 2020. That is a Known Lie: Detecting Previously Fact-Checked Claims. In *ACL 2020, Online, July 5-10, 2020*, 3607–3618. ACL.
- Shu, K.; Wang, S.; and Liu, H. 2019. Beyond News Contents: The Role of Social Context for Fake News Detection. In *WSDM 2019, Melbourne, VIC, Australia, February 11-15, 2019*, 312–320. ACM.
- Shugars, S.; Gitomer, A.; McCabe, S.; Gallagher, R. J.; Joseph, K.; Grinberg, N.; Doroshenko, L.; Welles, B. F.; and Lazer, D. 2021. Pandemics, Protests, and Publics: Demographic Activity and Engagement on Twitter in 2020. *Journal of Quantitative Description: Digital Media*, 1.
- Smeros, P.; Castillo, C.; and Aberer, K. 2019. SciLens: Evaluating the Quality of Scientific News Articles Using Social Media and Scientific Literature Indicators. In *WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, 1747–1758. ACM.
- Smeros, P.; Castillo, C.; and Aberer, K. 2021. SciClops: Detecting and Contextualizing Scientific Claims for Assisting Manual Fact-Checking. In *CIKM '21, Queensland, Australia, November 1 - 5, 2021*, 1692–1702. ACM.
- Van Bavel, J. J.; Baicker, K.; Boggio, P. S.; Capraro, V.; Cichocka, A.; Cikara, M.; Crockett, M. J.; Crum, A. J.; Douglas, K. M.; Druckman, J. N.; et al. 2020. Using social and behavioural science to support COVID-19 pandemic response. *Nature human behaviour*, 4(5): 460–471.
- Vishwakarma, D. K.; Varshney, D.; and Yadav, A. 2019. Detection and veracity analysis of fake news via scrapping and authenticating the web search. *Cogn. Syst. Res.*, 58: 217–229.
- Wang, L. L.; Lo, K.; Chandrasekhar, Y.; Reas, R.; Yang, J.; Eide, D.; Funk, K.; Kinney, R.; Liu, Z.; Merrill, W.; Mooney, P.; Murdick, D. A.; Rishi, D.; Sheehan, J.; Shen, Z.; Stilson, B.; Wade, A. D.; Wang, K.; Wilhelm, C.; Xie, B.; Raymond, D.; Weld, D. S.; Etzioni, O.; and Kohlmeier, S. 2020. COVID-19: The Covid-19 Open Research Dataset. *CoRR*, abs/2004.10706.
- Wang, R.; Shivanna, R.; Cheng, D. Z.; Jain, S.; Lin, D.; Hong, L.; and Chi, E. H. 2021. DCN V2: Improved Deep & Cross Network and Practical Lessons for Web-scale Learning to Rank Systems. In Leskovec, J.; Grobelnik, M.; Najork, M.; Tang, J.; and Zia, L., eds., *WWW '21, Ljubljana, Slovenia, April 19-23, 2021*, 1785–1797. ACM / IW3C2.
- Weinberger, K. Q.; and Saul, L. K. 2009. Distance metric learning for large margin nearest neighbor classification. *Journal of machine learning research*, 10(2).
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; and Brew, J. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *CoRR*, abs/1910.03771.
- Yang, S.; Shu, K.; Wang, S.; Gu, R.; Wu, F.; and Liu, H. 2019. Unsupervised Fake News Detection on Social Media: A Generative Approach. In *AAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, 5644–5651. AAAI Press.
- Yin, Z.; Sachidananda, V.; and Prabhakar, B. 2018. The Global Anchor Method for Quantifying Linguistic Shifts and Domain Adaptation. In *NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, 9434–9445.
- Zarocostas, J. 2020. How to fight an infodemic. *The Lancet*, 395(10225): 676.
- Zhang, A. X.; Ranganathan, A.; Metz, S. E.; Appling, S.; Sehat, C. M.; Gilmore, N.; Adams, N. B.; Vincent, E.; Lee, J.; Robbins, M.; Bice, E.; Hawke, S.; Karger, D. R.; and Mina, A. X. 2018. A Structured Response to Misinformation: Defining and Annotating Credibility Indicators in News Articles. In *WWW 2018, Lyon, France, April 23-27, 2018*, 603–612. ACM.
- Zubiaga, A.; Kochkina, E.; Liakata, M.; Procter, R.; Lukasik, M.; Bontcheva, K.; Cohn, T.; and Augenstein, I. 2018. Discourse-aware rumour stance classification in social media using sequential classifiers. *Inf. Process. Manage.*, 54(2): 273–290.