

Using Persuasive Writing Strategies to Explain and Detect Health Misinformation

Danial Kamali*, Joseph Romain*, Huiyi Liu†, Wei Peng‡, Jingbo Meng§, Parisa Kordjamshidi*

*Department of Computer Science and Engineering, Michigan State University,

†Department of Communication, Michigan State University,

‡Department of Media and Information, Michigan State University,

§School of Communication, Ohio State University

{kamalida,romainj2,liuhui5,pengwei}@msu.edu, meng.28@osu.edu, kordjams@msu.edu

Abstract

Nowadays, the spread of misinformation is a prominent problem in society. Our research focuses on aiding the automatic identification of misinformation by analyzing the persuasive strategies employed in textual documents. We introduce a novel annotation scheme encompassing common persuasive writing tactics to achieve our objective. Additionally, we provide a dataset on health misinformation, thoroughly annotated by experts utilizing our proposed scheme. Our contribution includes proposing a new task of annotating pieces of text with their persuasive writing strategy types. We evaluate fine-tuning and prompt-engineering techniques with pre-trained language models of the BERT family and the generative large language models of the GPT family using persuasive strategies as an additional source of information. We evaluate the effects of employing persuasive strategies as intermediate labels in the context of misinformation detection. Our results show that those strategies enhance accuracy and improve the explainability of misinformation detection models. The persuasive strategies can serve as valuable insights and explanations, enabling other models or even humans to make more informed decisions regarding the trustworthiness of the information.

Keywords: Misinformation Detection, Explainable AI, Persuasive Writing Strategy, Large Language Model

1. Introduction

The concept of misinformation has long been present, but in recent years, there has been an alarming surge in the dissemination of inaccurate information, particularly concerning health and the medical industry (Suarez-Lledo and Alvarez-Galvez, 2021). This influx of misinformation has had detrimental effects on society as a whole. People’s trust in news organizations has fallen drastically, resulting in a less informed populace (Islam et al., 2020). Many news resources purposefully make false news articles to spread misinformation and increase hostility between political groups. Most prior online misinformation research focuses on the political context (Lazer et al., 2018). However, health-related misinformation has grown exponentially during the COVID-19 pandemic, and these types of misinformation have a significant impact on individuals and society (Loomba et al., 2021). There are vast amounts of data for human fact-checkers, more than what they can go through manually, so professionals have turned to machine learning for automated misinformation identification for years, which is still a prominent challenge.

For people to trust a machine’s verdict in misinformation, the models must produce interpretable reasoning as to why an article is incorrect. This paper presents a novel annotation scheme containing a comprehensive hierarchy of persuasive writing strategies systematically identified from online mis-

information (Wei Peng and Meng, 2023), as shown in Figure 1. We train models to spot pieces of an article that can potentially make it untrustworthy. By understanding these tactics in the information source, users can have a deeper analysis and are less likely to be misled by false information. Our contributions include 1) A new corpus annotated with comprehensive hierarchical persuasive writing strategy labels. 2) A new challenging task of characterizing the persuasive writing strategies in text and developing baselines. 3) Demonstrating the integration of persuasive strategies enhances misinformation detection accuracy and offers insightful explanations that elucidate the underlying reasons for classifying particular information as false.

2. Related Work

Systematic research on rumors in mass media and deception in interpersonal communication dates back more than 70 years ago (Allport and Lepkin, 1945; Knapp, 1944). Nowadays, the significance of misinformation is increasingly evident as it effortlessly transcends temporal and geographical boundaries. Furthermore, the ease of searching, archiving, and accessing the information (Fernández-Luque and Bau, 2015) exacerbates the situation (Vosoughi et al., 2018; Shu et al., 2017).

In this study, we focus on health misinformation

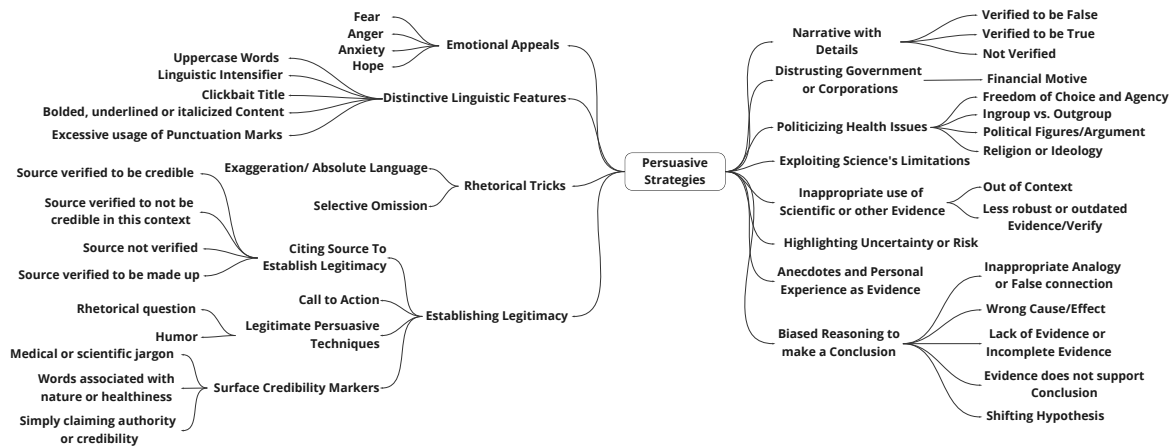


Figure 1: Persuasive Strategies annotation hierarchy scheme.

(Wei Peng and Meng, 2023; Swire-Thompson and Lazer, 2019). Before the COVID-19 pandemic, health misinformation had already attracted the attention of researchers due to the resurgence of childhood vaccination misinformation on social media (Wang et al., 2019). However, the burgeoning health misinformation during the COVID-19 pandemic (Kouzy et al., 2020; Cuan-Baltazar et al., 2020) and its associated negative impacts (e.g., vaccination hesitancy (Loomba et al., 2021; Roozenbeek et al., 2020) has brought heightened attention from the research community to this pressing issue.

With the vast daily news output, it is unrealistic for human fact-checkers to verify every detail. As a result, many have turned to machine learning models to aid the process of claim verification and misinformation detection in its diverse forms (Rani et al., 2022; Yuliani et al., 2019; Della Vedova et al., 2018; Burfoot and Baldwin, 2009; De Sarkar et al., 2018; Martino et al., 2020; Khanday et al., 2021). Here, we overview the datasets, models, and techniques that are specifically relevant to health misinformation.

Datasets. Many datasets are constructed for training and evaluating misinformation detection models. These datasets tend to vary significantly in attributes and content. All datasets contain claims and labels from fact-checking websites such as Politifact (Wang, 2017), Snopes (Popat et al., 2016), or both (Vo and Lee, 2019). In recent years, more datasets have been created by crawling various fact-checking websites to increase generalizability (Shahi and Nandini, 2020). Besides claims and labels, datasets often include evidence pages and metadata from sources such as Google (Augenstein et al., 2019) and Wikipedia (Aly et al., 2021; Yang et al., 2022). Recently, Yang et al. (2022) constructed two explainable fake news datasets to help develop explainable models. Our dataset uses a subset of claims from the MultiFC

dataset (Augenstein et al., 2019) that we augmented with supporting articles for our in-depth context analysis purposes.

Existing Models. Different approaches have been adopted to address automated misinformation detection. One research branch utilizes Natural Language Processing (NLP) techniques and large language models (LLMs) to detect misinformation within news texts (Manzoor et al., 2019; Zhang and Gao, 2023; Jiang et al., 2023). Recent studies highlight the beneficial impact of incorporating syntactic information into transformer models, as demonstrated by Kamali and Kordjamshidi (2023), and its specific advantages in detecting misinformation (Nagy and Kapusta, 2021). Another emerging research area focuses on identifying misinformation by examining specific textual features. For instance, Ghanem et al. (2019) highlighted that falsified information often carries unique emotional signatures. Building on this observation, Ghanem et al. (2021) introduced FakeFlow, a model that captures the flow of affective information and emotions within an article. This can assist in identifying articles that strategically manipulate reader emotions. Furthermore, the DEF (Luvembe et al., 2023) study considers both the emotion of the publisher and the social emotions it provokes, potentially broadening the context for detecting misinformation. Recent studies have increasingly underscored the utility of LLMs. Notably, Lee et al. (2020) employed LLMs such as GPT-2 (Radford et al., 2019) to assess claim factuality through evidence-conditioned perplexity. While there remains a scarcity of research on LLM misinformation detection (Zhang and Gao, 2023), emerging works (Press et al., 2023; Pelrine et al., 2023; Yao et al., 2023; Jiang et al., 2023) suggest that fusing LLM reasoning with external knowledge can substantially bolster performance in reasoning-intensive NLP tasks. Compared to approaches like FakeFlow, which centers on emotional manipula-

tion, our research adopts a more holistic perspective. We investigate a broader spectrum of factors through the lens of persuasive strategies. This approach not only enhances our ability to detect misinformation but also improves the interpretability framework. Consequently, this framework further aids non-communication experts in misinformation detection (Basol et al., 2021). In contrast to LLM-based approaches such as those by Jiang et al. (2023); Zhang and Gao (2023), our model labels misinformation without requiring access to any external knowledge base.

3. Persuasive Writing Strategies

Persuasion is an essential part of news misinformation. Misinformation often plays on people's suspicions, premises, and biases to create narratives that feel plausible and compelling (Vosoughi et al., 2018). By employing common persuasive techniques such as appeals to emotion, anecdotal evidence, and bandwagon appeals, misinformation can seem more credible than it actually is. Previous studies (Eisend and Tarrahi, 2022; Guess et al., 2020) have demonstrated that when individuals become aware of the persuasive tactics or have adequate media or information literacy to understand these persuasive attempts, they become less susceptible to persuasion attempts such as political propaganda and advertising. Such media literacy interventions coupled with fact-checking were found to be more effective than each component used alone (Hameleers, 2022).

3.1. Persuasive Strategies Serve as Explanations

The deliverable of the current misinformation models usually indicates whether the information is true or false or evidence supporting or denying claims. A single label can only have a limited effect in diminishing the influence of misinformation (Margolin et al., 2018). One step forward from the simple fact-checking label is to provide an explainable justification for why specific claims are misleading (Kotonya and Toni, 2020; Atanasova et al., 2020). The research in communications has long discovered that providing justification and understanding the underlying techniques used in claims to influence the information receiver will be critical to helping individuals cope with persuasive attempts (Friestad and Wright, 1994). The current study attempts to automatically annotate persuasive features in misinformation, which can serve as intermediate features for classifying misinformation and as auxiliary information for future media literacy training tools to assist individuals in better misinformation detection.

3.2. Persuasive Strategies in Health Misinformation

The persuasive features implemented in the current study were based on a review that systematically screened 1,700 research articles related to online health misinformation to identify 12 thematic groups of persuasive strategies from 58 eligible articles (Wei Peng and Meng, 2023). For instance, for the theme of politicizing health issues, there are subgroups such as the trope of freedom and choice, the rhetoric of ingroup vs. outgroup, citing political figures or political arguments, and the use of religion and ideology. Although some similar approaches are available, such as bias detection based on political ideology (Baly et al., 2020), propaganda detection based on linguistic features (Rashkin et al., 2017) or using multiple categories of features (Da San Martino et al., 2019, 2020; Dimitrov et al., 2021), persuasive strategies based on a particular theory such as Aristotle's rhetoric theory (Chen et al., 2021) or Moral Foundation Theory (Lin et al., 2018), the persuasive writing techniques identified and annotated in our study have a competitive advantage because they were based on a systematic and comprehensive review of current literature to provide a full-scale framework.

4. Dataset Construction

4.1. Health-related Subset from MultiFC

To build our persuasive writing strategy dataset, we have chosen to use a subset of the MultiFC (Augenstein et al., 2019) dataset, consisting of claims labeled as health, medical, or food, containing 599 claims (7 duplicates). We refer to this as *Health subset*. Each sample of the MultiFC dataset includes a claim (the main point of a news article), a few sentences of supporting articles with their links, and a domain label providing the context for each claim. The misinformation labels for this subset are *mostly true*, *truth!*, *true*, *disputed!*, *in the green*, *in the red*, *mostly truth!*, *mostly false*, *fiction!*, *false*. The *unproven* (215 claims), *satirical* (22 claims), and non-health or non-medical related articles (33 claims) were excluded or filtered out through an expert review process, leaving a total of 329 entries.

4.2. Retrieval of Full Articles for Annotation

We manually retrieved the complete article of each claim from the MultiFC health subset by searching the claim on the Google search engine because the URLs provided in the MultiFC mostly turned invalid. A health communication master's student

used the claims or the keywords in the claim to conduct the search and browse the search results to identify the full article that matched the claim. Out of the 329 claims investigated, 86 original articles were either deleted from the internet or existed in a format other than text, resulting in a total of **242** articles containing **5666** sentences. Three expert annotators annotated the remaining articles with persuasive writing labels using the WebAnno annotation tool (Eckart de Castilho et al., 2016).

4.3. Annotation of Full Articles

4.3.1. Annotation Codebook Development

Our annotation scheme organizes the persuasive writing strategies into three hierarchy levels, as shown in Figure 1. We call the levels the parent tag level, the child tag level, and the sub-tag level. The parent tag level is the overarching group of strategies, corresponding to 12 thematic groups of persuasive strategies identified in Wei Peng and Meng (2023), such as *Politicizing Health Issues*. The child tag level is the subgroup of more specific strategies, e.g., *Politicizing Health Issues- Religion/Ideology*, *Emotional Appeals-Fear*. Although the original framework of Wei Peng and Meng (2023) had only two levels, in the codebook development process, based on health communication and persuasion theories, a third level, i.e., sub-tag level, was added by the health communication experts in our team to provide even more specificity when needed, e.g., *Establishing legitimacy- Surface Credibility Markers- Medical/Scientific jargon*. These additional subgroup child tags and the sub-tag level were added in our paper to the original framework in Wei Peng and Meng (2023).

In our labeling scheme, we impose a hierarchical structure; this means that text can be annotated with a child tag without being annotated with any of its sub-tags. However, if a text sequence is annotated with a child tag, it must be annotated with its respective parent tag. If a text sequence is annotated with a sub-tag, it must also be annotated with its respective child tag. Figure 2 shows an example of persuasive strategy labeling. Although the annotations can apply to any piece of text and any level of granularity, such as phrases or sentences, as we will see in the task specification, we assign the strategies to the sentences containing them.

Two communication professors and a master's student specializing in health and risk communication, all our team members, developed and refined an annotation codebook for these three-level labels. The master's student annotated all the articles. Note that the annotators only annotated the persuasive strategies. The claim of each article was assumed to be the central idea of the full article. The label or annotation of the claim was

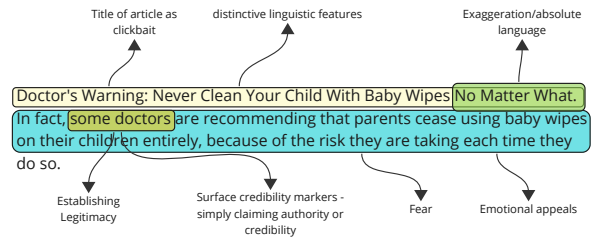


Figure 2: Persuasive strategy labeling example.

imported from MultiFC or their origin in MultiFC, which used fact-checking information to assign the label. In other words, we did not need to deal with the bias among the annotators regarding judgment on whether the claim or the whole article included false information.

4.3.2. Inter Annotator Agreement

A random sample of 25 articles (10 percent of the entire sample) was coded by the master's student and the two professors. The inter-annotator reliability was satisfactory with the following details: The average pairwise percentage agreement was 88.99%; pairwise agreement between the master's student and the two professors was 88.77% and 86.78%, respectively; and pairwise agreement between the two professors was 91.41%. The average pairwise Cohen's Kappa was 0.88; pairwise Cohen's Kappa between the master's student and the two professors were 0.88 and 0.86, respectively; and pairwise Cohen's Kappa between the two professors was 0.91. The master's student annotated all the articles and then double-checked by at least one of the professors to resolve any discrepancies. The detailed frequency of each persuasive strategy can be found in Appendix E.

5. Approach

We follow three learning tasks in this paper.

Task 1 (Misinformation Detection): Given a claim C and an article A , a model should classify the article with its trustworthiness label, Tw . In the simplest case, Tw can be a binary label indicating the document's truth. However, detecting the veracity of text is often more complex; thus, these label sets typically include more values, such as *Mixed*, *Unverifiable*, and *Outdated*.

Task 2 (Persuasive Strategy Labeling): Given an article A represented as a sequence of n sentences, $A = \{s_1, s_2, \dots, s_n\}$, the task is to assign each sentence s_i a set of persuasive strategy labels in $L = \{l_1, l_2, \dots, l_m\}$. A boolean value l_{ij} indicates whether sentence s_i follows persuasive strategy l_j . Each sentence s_i can carry multiple labels, indicating the presence of several strategies. In other words, sentences could be marked

with multiple labels in the hierarchy of persuasive writing strategies or none of them. In our setting, L follows a label hierarchy shown in Figure 1. In this task, the surrounding context can be used as a part of the input information for each sentence classification. The input sentence can be evaluated with three levels of context: *None*, *Low*, and *High*. The *Low* context includes two additional surrounding sentences, one to the left of the target sentence and one to the right. *High* context includes four additional sentences, two to the left of the target sentence and two to the right. While this task can be performed at the sentence or paragraph level, we found the sentence-level assignment of the labels to be more effective.

Task 3 (Combining Persuasive Strategies & Misinformation Detection): Since the persuasive strategies potentially can help both the detection and characterization of the news document, we define this third combined task. The input includes the claim C and article A . The output is the sentence-level persuasive strategy labels l_{ij} and the misinformation label Tw . This task can be modeled as a pipeline or a joint labeling setting.

5.1. Models

In recent years, foundation models have emerged as prominent tools, excelling in numerous NLP problems and establishing themselves as robust baselines (Bommasani et al., 2022). In our experiments, we employ two distinct types of foundation models. Firstly, we utilize RoBERTa (Liu et al., 2019), a pre-trained encoder-based model extensively adopted across the NLP community, trained using masked language modeling. Furthermore, we use GPT-3 (Brown et al., 2020) and the more recent GPT-4 (OpenAI, 2023), both of which have significantly larger parameter configurations, serving as our generative language models. For all tasks and experiments, we fine-tune RoBERTa-based models. We pass the RoBERTa model's CLS token result to a classification layer to perform classification tasks. GPT-3 and GPT-4 are used in the zero-shot setting due to their high training cost.

6. Experiments

The goal of our experiments and analysis is to answer the following research questions:

Q1. Do the foundation models provide reasonable accuracy on the misinformation detection task compared to earlier results? (Section 6.1)

Q2. How difficult is it to predict the various persuasive writing strategies in sentences after fine-tuning the baseline foundation models? (Section 6.2)

Q3. How well do the foundation models predict the article's veracity in *Health subset*? (Section 6.3)

Q4. How will including persuasive strategy labels affect a model's ability to detect misinformation in the articles? (Section 6.3)

Experimental details. We ran all experiments using PyTorch (Paszke et al., 2019). We used the pandas (pandas development team, 2020) library for data manipulation and the NLTK (Bird et al., 2009) and spacy (Honnibal et al., 2020) libraries for tokenization. Our reported results are the average of three runs with different random seeds on an NVIDIA A6000 GPU. The details of the hyperparameters of our experiments are in Appendix C. Our code, models, and data are publicly available¹.

6.1. Competency of Foundation Models

In this section, we demonstrate a small experiment to evaluate the effectiveness of our baseline foundation model in misinformation detection to answer Question 1. For a broader view of the evaluation of a variety of language models, we refer the reader to many recent research results that underscore the effectiveness of GPT-4 and GPT-3. A series of studies, including those by Lee et al. (2020), Press et al. (2023), Pelrine et al. (2023), Yao et al. (2023), and Jiang et al. (2023), highlight the effectiveness of GPT models in misinformation detection. Hoes et al. (2023), in particular, emphasized GPT-3's capability in identifying misinformation in news articles, suggesting that its capabilities extend beyond mere memorization.

To establish the basis of our comparisons in previous research results and set up a reasonable foundation model baseline, we train and test RoBERTa on the *pmot* domain subset of the Multi-FC's training set to demonstrate RoBERTa's competency in misinformation detection. This subset is the most frequent domain in the Multi-FC dataset, containing 15,390 examples. We then compare our foundation model to the top-performing model reported in the original MultiFC results on the similar subset (Augenstein et al., 2019).

In this experiment, the input comprises a claim and related evidence snippets retrieved from the Multi-FC dataset. We concatenated these evidence snippets with the claim by separator tokens and tokenizing the resultant string. The model produces an output in the form of a one-hot vector of labels, corresponding to the nine unique classes of the *pmot* domain.

As indicated in Table 1, the RoBERTa-based model demonstrates acceptable performance compared to the previously mentioned Multi-FC model.

¹<https://github.com/HLR/Misinformation-Detection>

This is evident from the higher Micro and Macro F1 scores achieved by RoBERTa, confirming its ability to yield satisfactory accuracy levels with our collected dataset and marking it as a suitable baseline for further analysis.

The *pomt* domain consists of 9 labels, some of which contain close to no instances. Due to this label imbalance, the Macro F1 score better represents the model performance. This result indicates that Multi-FC’s model is more biased towards classes with many instances, while RoBERTa does equally well on all classes. Note the Multi-FC test set was not made available; thus, the results are reported on the development set and tuned with 10% of the training set. Given this result, we use RoBERTa in the next experiments and merely focus on the subset that we annotated with persuasive writing strategies.

Model	F1 _{Micro}	F1 _{Macro}
(Augenstein et al., 2019)	0.321	0.276
RoBERTa	0.327	0.325

Table 1: Multi-FC and RoBERTa-based model results on the *pomt* domain.

6.2. Detection of Persuasive Strategies

We evaluate the prediction of persuasive writing strategies (Task 2) using four RoBERTa-based classifiers for the classification of the labels in various levels of the hierarchy shown in Figure 1 and explained in Section 4. The input to these models is a tokenized text sequence extracted from an article. This text sequence includes a target sentence and is accompanied by nearby sentences in some of our model variations. This format is used for classification models of all levels of the hierarchy. The level 1 model is a binary classification model designed to detect whether the input text follows a persuasive writing strategy. The output of level 1 is $\{NotPersuasive, Persuasive\}$. In levels 2-4, a piece of text can be annotated with no tags or multiple tags. The level 2 classifier includes 12 labels (e.g., *establishing legitimacy*). The level 3 classifier includes 30 labels (e.g., *surface credibility marker*). The level 4 classifier model includes nine labels at depth 4 of the label hierarchy (e.g., *medical or scientific jargon*). The results of our experiments are shown in Table 2. We used 193 articles, including 4,249 sentences for training, and the remaining 49 articles, including 1,417 sentences for testing. Also, we ignored labels with a frequency lower than 15 due to label sparsity. Our experiments show that extra contextual information generally helps with persuasive strategy labeling, except for level 4. This exception can be justified considering that level

4 strategies are not context-dependent and are usually expressed in one sentence, such as *claiming authority* or *words associated with health*. In contrast, some labels, such as *wrong cause/effect* or sentiments, which are more context-dependent, improve as more context is provided. As mentioned earlier, since our data is highly imbalanced, we consider the Macro-F1 score as the prominent evaluation metric. The detailed evaluation of precision, recall, and F1 per label with the number of training/test examples can be found in Appendix G.

Level 1		
Context	F1 _{Micro}	F1 _{Macro}
None	0.668	0.662
Low	0.677	0.676
High	0.699	0.692

Level 2		
Context	F1 _{Micro}	F1 _{Macro}
None	0.468	0.342
Low	0.501	0.372
High	0.497	0.383

Level 3		
Context	F1 _{Micro}	F1 _{Macro}
None	0.430	0.270
Low	0.464	0.283
High	0.463	0.288

Level 4		
Context	F1 _{Micro}	F1 _{Macro}
None	0.428	0.291
Low	0.434	0.265
High	0.441	0.256

Table 2: RoBERTa results for four levels of persuasive strategies, with three levels of context.

6.3. Misinformation Detection Using Persuasive Strategies

In this section, we first evaluate misinformation detection on our *Health subset* without using the strategy labels (Task 1) to answer Question 3. Then, we evaluate our misinformation detection model with persuasive strategies annotations coming from the prediction of annotation models, i.e., the output of Task 2. In fact, this is Task 3 with a pipeline configuration. In some additional experiments, we use ground-truth strategy labels to obtain the effect of gold annotations when no pipeline errors are propagated from the persuasive strategy labeling step. We specifically evaluate how different input combinations affect the performance, emphasizing whether the inclusion or exclusion of persuasive strategy labels from the articles has any significant impact on the target task. We fine-tune a

RoBERTa-based model on the *Health subset*. The input to this model is a set of tokens that can originate from articles, claims, or persuasive strategy labels, depending on the experiment. In the future, the hierarchical structure of the labels can be exploited as constraints in a more sophisticated approach for modeling this problem and joint training (Rajaby Faghihi et al., 2021, 2023).

Noisy and Imbalanced Labels. Our selected subset includes instances from various labeling domains, leading to labels with similar meanings like *Truth!* and *True*. Such ambiguous labeling can hinder effective training. To address this, we normalized the labels, reducing the initial ten labels to two labels, *True* and *False*. Detailed information about our label normalization approach is provided in Appendix F. Moreover, we excluded articles from the Multi-FC labeled as *Mixed*. Post-normalization, our subset comprises **170** instances labeled as *False* (indicating misinformation) and **71** instances labeled as *True* (indicating truthful information). The resulting data elements include a binary label and are used in a binary classification setting. Also, to manage the label imbalance, we integrated a weighted loss into our model.

Misinformation Detection without Persuasive Strategy Labels. We first trained and tested RoBERTa using the original article’s claim and article to obtain a baseline over our collected health subset. Table 3 shows the performance of our models is considerably higher than the majority baseline in terms of the F1-score.

Input Source	F1 _{Micro}	F1 _{Macro}
Majority Baseline	0.673	0.402
RoBERTa _{claim}	0.830	0.793
RoBERTa _{article}	0.810	0.765
RoBERTa _{claim+article}	0.823	0.799
GPT-4 _{claim}	0.837	0.791
GPT-4 _{article}	0.848	0.830
GPT-4 _{claim+article}	0.913	0.904

Table 3: Misinformation Detection results (Task 1). Model_{source} represents the foundation model classifier with *source* as the input source.

Misinformation Detection with Pipeline of Persuasive Strategy Labels. In this experiment, we assess the performance of our misinformation detection model when supplemented with persuasive strategy labels. This pipeline configuration begins with applying the persuasive strategy type classifiers (levels 2-4) to the article. Following this, the resulting strategy labels, alongside the articles, are passed to the misinformation detection model. In this configuration, the model outputs are represented as a token-separated string, and the article text was summarized using the Gensim summa-

rizer (Řehůřek and Sojka, 2010) so that the article text and all persuasive strategies could fit into the RoBERTa input layer with the 512 tokens limit. Alternatively, we could trim the documents and limit the input to the first 512 tokens. However, this setting was suboptimal as we would have lost a large portion of the potentially important information. We did not have such a limitation for the GPT experiments. We evaluated our system with different input source combinations. The results are shown in Table 4.

Input Source	F1 _{Micro}	F1 _{Macro}
GT	0.898	0.872
Pred _{none}	0.687	0.622
Pred _{low}	0.735	0.704
Pred _{high}	0.741	0.690
Claim+GT	0.912	0.891
Claim+Pred _{none}	0.694	0.646
Claim+Pred _{low}	0.728	0.683
Claim+Pred _{high}	0.735	0.672
Claim+Article+GT	0.871	0.845
Claim+Article+Pred _{none}	0.748	0.721
Claim+Article+Pred _{low}	0.762	0.722
Claim+Article+Pred _{high}	0.782	0.749
GPT-4 _{claim+Pred}	0.783	0.681
GPT-4 _{claim+GT}	0.878	0.850
GPT-4 _{claim+article+Pred}	0.913	0.901
GPT-4 _{claim+article+GT}	0.939	0.933
GPT-4 _{in-context learning}	0.932	0.922

Table 4: Misinformation Detection results with different sources (Task 3). Pred_{context} refers to persuasive strategy prediction labels concerning the context level, and GT refers to the persuasive strategy ground truth labels. Articles in the RoBERTa experiments are summarized to satisfy its 512 token limit. GPT-4_{in-context learning} result has persuasive strategies information in the prompt.

Zero-shot GPT-4. To explore the ability of generative large language models such as the GPT family on the misinformation detection task, we evaluated GPT-4 in a zero-shot setting using our test set. Similar to the RoBERTa experiment, we evaluate variations of input sources in this experiment by concatenating the information to the zero-shot. We structure our prompt in a way that asks the model to classify the claim based on auxiliary information such as *Article* or *Labels*. Figure 3 shows the structure of our prompt for this experiment. The results of this experiment are shown in Table 3 and 4. The extension of these results, including the GPT-3 evaluation, can be found in Appendix A.

Input Variation Observation. One experimental observation is that incorporating multiple input sources can boost the accuracy of misinformation detection. Our results demonstrate that,

```

Assuming you are a journalism and communication
expert. Is this claim correct?
claim: {claim}
We have found this article supporting the claim:
article: {article}
our communication expert reported that the article
supporting this claim follows the below persuasive
strategies
{labels}
You must answer to the best of your knowledge.
Give me one word answer "Yes" or "No"?
-----
[GPT]: "Yes"

```

Figure 3: An example of misinformation detection prompt for GPT model in a zero-shot setting. The prompt structure would vary based on the input sources. More prompt templates can be found in Appendix A.

in many cases, adding inputs such as *Article* or *Preds* to the model enhances its performance compared to using the *Claim* alone. While this is the case for GPT, sometimes, the additional information reduces the F1 score for the RoBERTa model. We associate this performance discrepancy with RoBERTa’s limited ability to process larger text. Additionally, there might be adverse effects due to the text summarization process, particularly when considering GPT models’ capacity to understand longer context and dependencies. This hypothesis becomes even more apparent when we compare the outcomes for *Claim + GT* with those for *Claim + Article + GT*, where the addition of the *Article* input led to a marked decrease in RoBERTa’s performance.

Moreover, the findings depicted in Table 4 demonstrate that simply integrating persuasive strategy labels into the prompt of a GPT model can enhance accuracy when operating with cutting-edge large language models. These outcomes indicate the benefits of incorporating persuasive strategies into language models across various sizes and training frameworks.

Using Ground Truth Strategy Labels. We observed that including ground truth persuasive strategies, or *GT*, markedly bolsters the accuracy of misinformation detection. This is clearly illustrated when comparing the results of *Claim + GT* and *Claim + Article + GT* in Table 4 with their counterpart experiments without *GT* in Table 3. For instance, $GPT-4_{claim+article+GT}$ outperforms $GPT-4_{claim+article}$ about 2.9% Macro F1-score. This substantial improvement emphasizes these annotations’ significance to the misinformation detection task. On the other hand, the results of the predicted persuasive strategy do not show the same level of performance, both independently

and when integrated with the claim or article. As shown in Table 4, persuasive strategy labeling quality directly correlates with misinformation detection quality. This implies that augmenting models with high-accuracy persuasive strategy labels can indeed help with misinformation detection, while incorporating noisy strategy labels could, in fact, be detrimental to the model’s performance. Further detailed results are demonstrated in Appendix D.

Two-step Misinformation Detection with predicted Persuasive Strategies in context. Here, our aim is to evaluate the real-world applicability of the persuasive strategy labels method by assessing the in-context learning capabilities using the GPT-4 model. To this end, we first select one example from each of our persuasive strategies and the corresponding article span, e.g., "No Matter What" \Rightarrow ["Exaggeration/absolute language,']. Given this context, we ask the GPT model to identify the persuasive strategies employed within a given article. After acquiring the labeled strategies, we ask the model to determine the veracity of the claim based on known facts and the marked persuasive writing strategies. As evident from Table 4, this step-by-step in-context learning approach enhances the misinformation detection performance compared to zero shot. Even with imperfect labeling of persuasive strategies, we observed a consistent improvement when leveraging them in the prompt. This result further indicates the effectiveness of the persuasive strategies in both fin-tuning and in-context learning with language models. The prompts we used for this experiment are shown in Appendix A. We evaluated the intermediate strategy labeling with GPT, focusing on the second level of the hierarchy that includes labels such as *Narrative with Details* and *Emotional Appeals*. Our results demonstrated an average F1-score of 0.429 in detecting persuasive writing techniques within an article. This result indicates the challenge of accurately detecting the strategies even though it was positively impacting our in-context learning. For further information regarding this experiment, please refer to Appendix B.

It is worth mentioning that generative models like GPT-4 and GPT-3 offer enhanced explainability. By employing the two-step and CoT prompting in the misinformation detection phase. The detected strategies can be used to explain the reasoning behind the decision that is tailored for the end-user. An example of this is shown in Figure 4.

Experiments with a More Recent Dataset. To further evaluate the effectiveness of the persuasive strategies when used as in-context learning, we employed the RAWFC dataset (Yang et al., 2022), which is widely used for in-context learning experiments. The RAWFC dataset is a fact-checking benchmark, providing claims categorized under a

User: Here we show example of persuasive strategy detection. Examples below show text spans with their corresponding persuasive strategy:

span₁ ⇒ ['persuasive strategy₁';persuasive strategy₂,...]
span₂ ⇒ ['persuasive strategy₁';persuasive strategy₂,...]
...
Mark the sentences in the text below with persuasive strategy labels: {article}

GPT-4: [Then after getting the labels from the GPT-4 model]

User: Given the labeled persuasive strategies in the above supporting article, Do you think this claim is correct?

Claim: The Government is banning organic Farming
Give me a Yes or No answer with your chain of thoughts.

GPT-4: No, the claim that the government is banning organic farming is not correct based on the labeled persuasive strategies in the supporting article. The article uses a lot of fear and uncertainty to make its point, often exaggerating the potential effects of the bill. It also shows a distrust of government and corporations, suggesting a bias in the information presented. Furthermore, the call to action to read the legislation for oneself implies that the author may not have provided a full or accurate summary of the bill. Therefore, the claim seems to be based more on speculation and fear-mongering than on concrete evidence.

Figure 4: The prompt structure for the explained misinformation detection using persuasive strategies with an example response.

three-class scheme: *True*, *False*, or *Half*. This dataset source is real-world news, making it an authentic representation of the claims circulating in today’s media landscape. It comprises 200 claims for each training, validation, and testing split. Of these, 67 claims are labeled as *True*, another 67 as *False*, and the remaining 66 are categorized as *Half* for each split. It also provides at least one article for each claim. For this experiment, first, we chose the two longest articles (if available) of the claim from the dataset. Second, we ask our model to identify persuasive strategies with in-context samples similar to the previous experiment. Subsequently, we ask the model to determine the veracity of the claim provided. The specific prompts utilized for these experiments can be found in Appendix A. As shown in Table 5, even using noisy persuasive strategies within the in-context learning setting not only outperforms the standard prompting and Chain-of-thought techniques but also exceeds the accuracy of methods that access online search engine (Zhang and Gao, 2023; Yao et al., 2023). This is a strong result indicating the advantage of using LLM’s capability for recognizing communication strategies in language and employing them to inform users about the potential risks of misinformation.

7. Conclusion

This study highlights the positive effect of persuasive strategy labeling in misinformation detection, which also helps the need for interpretable results in misinformation detection models to enhance

Model	P (%)	R (%)	F1 (%)
Standard	48.5	48.5	48.5
Standard _{witharticles}	48.5	47.5	47.9
CoT	42.4	46.6	44.4
ReAct [†]	51.2	48.5	49.8
HiSS [†]	53.4	54.4	53.9
ICL _{ours}	56.9	55.5	55.8

Table 5: Experimental results of claim verification on the RAWFC. Bold denotes the best performance. † denotes the access to search engines. The results on the RAWFC dataset are brought from the HiSS (Zhang and Gao, 2023). We used the davinci-003 to ensure a fair comparison.

transparency and public trust. We introduce a novel annotation scheme, which categorizes numerous persuasive writing strategies recognized by communication experts and presents a human-annotated dataset. Leveraging large language models, we evaluate various paradigms of fine-tuning and in-context learning to automate strategy detection while discussing associated challenges. We demonstrate a state-of-the-art misinformation detection model that provides explainability without requiring access to external knowledge. In addition, our dataset and the proposed challenge offer valuable resources for future research, aiming to foster critical thinking about online news articles by analyzing their employed writing strategies. By doing so, we strive to contribute towards a more informed populace.

8. Limitations

This study has notable limitations. Our primary aim was to analyze the influence of persuasive strategies on competitive baseline models. However, our dataset expansion has some challenges. Annotating persuasive writing strategies demands both expertise and substantial resources, which constrained the size of our dataset. Additionally, the costs associated with GPT's API calls, especially those involving long contexts, are significant and limit the dataset size used for the experiments. Consequently, we opted for RAWFC, a widely used benchmark with a manageable size. Furthermore, the inherent limitations of automated misinformation detection models should not be overlooked. While they offer valuable insights, and inform the user about the risks, they are not infallible. Their primary role is to supplement, not supplant, human judgment in addressing this critical societal issue. Lastly, the technical demands of large language models might restrict their applicability. Large pre-trained language models require significant computational power, potentially rendering them unsuitable for devices with limited resources.

9. Acknowledgement

We are grateful to our anonymous reviewers and Dr. Christos Christodoulopoulos for their invaluable suggestions and discussions. Support for this project was in part provided by the Trifecta Initiative Interdisciplinary Health Research at Michigan State University.

Credit Statement: **DK:** Formal Analysis, Methodology, Software, Experimentation, Writing-Original Draft, Writing-Review and Editing; **JR:** Formal Analysis, Methodology, Software, Experimentation, Writing-Original Draft; **HL:** Data Curation, Investigation; **WP:** Conceptualization, Data Curation, Funding Acquisition, Writing-Original Draft; **JM:** Conceptualization, Data Curation; **PK:** Conceptualization, Funding Acquisition, Experimental Design, Supervision, Writing-Review and Editing.

10. Bibliographical References

- Floyd H Allport and Milton Lepkin. 1945. Wartime rumors of waste and special privilege: Why some people believe them. *The Journal of Abnormal and Social Psychology*, 40(1):3.
- Rami Aly, Zhijiang Guo, Michael Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, Oana Cocarascu, and Arpit Mittal. 2021. [Feverous: Fact extraction and verification over unstructured and structured information](#).
- Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. Generating fact checking explanations. *arXiv preprint arXiv:2004.05773*.
- Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. [MultiFC: A real-world multi-domain dataset for evidence-based fact checking of claims](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4685–4697, Hong Kong, China. Association for Computational Linguistics.
- Ramy Baly, Giovanni Da San Martino, James Glass, and Preslav Nakov. 2020. [We can detect your bias: Predicting the political ideology of news articles](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4982–4991, Online. Association for Computational Linguistics.
- Melisa Basol, Jon Roozenbeek, Manon Berriche, Fatih Uenal, William P. McClanahan, and Sander van der Linden. 2021. [Towards psychological herd immunity: Cross-cultural evidence for two prebunking interventions against covid-19 misinformation](#). *Big Data & Society*, 8(1):20539517211013868.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. "O'Reilly Media, Inc."
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab,

- Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kudipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. 2022. [On the opportunities and risks of foundation models](#).
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *CoRR*, abs/2005.14165.
- Clint Burfoot and Timothy Baldwin. 2009. Automatic satire detection: Are you having a laugh? In *Proceedings of the ACL-IJCNLP 2009 conference short papers*, pages 161–164.
- Sijing Chen, Lu Xiao, and Jin Mao. 2021. Persuasion strategies of misinformation-containing posts in the social media. *Information Processing & Management*, 58(5):102665.
- Jose Yunam Cuan-Baltazar, Maria José Muñoz-Perez, Carolina Robledo-Vega, Maria Fernanda Pérez-Zepeda, and Elena Soto-Vega. 2020. Misinformation of covid-19 on the internet: infodemiology study. *JMIR public health and surveillance*, 6(2):e18444.
- Giovanni Da San Martino, Shaden Shaar, Yifan Zhang, Seunghak Yu, Alberto Barrón-Cedeño, and Preslav Nakov. 2020. [Prta: A system to support the analysis of propaganda techniques in the news](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 287–293, Online. Association for Computational Linguistics.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, Rostislav Petrov, and Preslav Nakov. 2019. [Fine-grained analysis of propaganda in news article](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5636–5646, Hong Kong, China. Association for Computational Linguistics.
- Sohan De Sarkar, Fan Yang, and Arjun Mukherjee. 2018. [Attending sentences to detect satirical fake news](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3371–3380, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Marco L. Della Vedova, Eugenio Tacchini, Stefano Moret, Gabriele Ballarin, Massimo DiPierro, and Luca de Alfaro. 2018. [Automatic online fake news detection combining content and social signals](#). In *2018 22nd Conference of Open Innovations Association (FRUCT)*, pages 272–279.
- Dimitar Dimitrov, Bishr Bin Ali, Shaden Shaar, Firoj Alam, Fabrizio Silvestri, Hamed Firooz, Preslav Nakov, and Giovanni Da San Martino. 2021. [SemEval-2021 task 6: Detection of persuasion techniques in texts and images](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 70–98, Online. Association for Computational Linguistics.
- Richard Eckart de Castilho, Éva Mújdricza-Maydt, Seid Muhie Yimam, Silvana Hartmann, Iryna Gurevych, Anette Frank, and Chris Biemann. 2016. [A web-based tool for the integrated annotation of semantic and syntactic structures](#). In *Proceedings of the Workshop on Language Technology Resources and Tools for Digital Humanities (LT4DH)*, pages 76–84, Osaka, Japan. The COLING 2016 Organizing Committee.
- Martin Eisend and Farid Tarrahi. 2022. Persuasion knowledge in the marketplace: A meta-analysis. *Journal of Consumer Psychology*, 32(1):3–22.
- Luis Fernández-Luque and Teresa Bau. 2015. Health and social media: perfect storm of information. *Healthcare informatics research*, 21(2):67–73.
- Marian Friestad and Peter Wright. 1994. The persuasion knowledge model: How people cope

- with persuasion attempts. *Journal of consumer research*, 21(1):1–31.
- Bilal Ghanem, Simone Paolo Ponzetto, Paolo Rosso, and Francisco M. Rangel Pardo. 2021. [Fakeflow: Fake news detection by modeling the flow of affective information](#). *CoRR*, abs/2101.09810.
- Bilal Ghanem, Paolo Rosso, and Francisco Rangel. 2019. [An emotional analysis of false information in social media and news articles](#).
- Andrew M Guess, Michael Lerner, Benjamin Lyons, Jacob M Montgomery, Brendan Nyhan, Jason Reifler, and Neelanjan Sircar. 2020. A digital media literacy intervention increases discernment between mainstream and false news in the united states and india. *Proceedings of the National Academy of Sciences*, 117(27):15536–15545.
- Michael Hameleers. 2022. Separating truth from lies: Comparing the effects of news media literacy interventions and fact-checkers in response to political misinformation in the us and netherlands. *Information, Communication & Society*, 25(1):110–126.
- Emma Hoes, Sacha Altay, and Juan Bermeo. 2023. Using chatgpt to fight misinformation: Chatgpt nails 72% of 12,000 verified claims.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Md Saiful Islam, Tonmoy Sarkar, Sazzad Hossain Khan, Abu-Hena Mostofa Kamal, S M Murshid Hasan, Alamgir Kabir, Dalia Yeasmin, Mohammad Ariful Islam, Kamal Ibne Amin Chowdhury, Kazi Selim Anwar, Abrar Ahmad Chughtai, and Holly Seale. 2020. COVID-19-related infodemic and its impact on public health: A global social media analysis. *Am. J. Trop. Med. Hyg.*, 103(4):1621–1629.
- Zhengbao Jiang, Frank F. Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [Active retrieval augmented generation](#).
- Danial Kamali and Parisa Kordjamshidi. 2023. Syntax-guided transformers: Elevating compositional generalization and grounding in multimodal environments. In *GenBench: The first workshop on generalisation (benchmarking) in NLP*, page 130.
- Akib Mohi Ud Din Khanday, Qamar Rayees Khan, and Syed Tanzeel Rabani. 2021. Identifying propaganda from online social networks during covid-19 using machine learning techniques. *International Journal of Information Technology*, 13(1):115–122.
- Robert H Knapp. 1944. A psychology of rumor. *Public opinion quarterly*, 8(1):22–37.
- Neema Kotonya and Francesca Toni. 2020. Explainable automated fact-checking for public health claims. *arXiv preprint arXiv:2010.09926*.
- 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).
- David M J Lazer, Matthew A Baum, Yochai Benkler, Adam J Berinsky, Kelly M Greenhill, Filippo Menczer, Miriam J Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A Sloman, Cass R Sunstein, Emily A Thorson, Duncan J Watts, and Jonathan L Zittrain. 2018. The science of fake news. *Science*, 359(6380):1094–1096.
- Nayeon Lee, Belinda Z. Li, Sinong Wang, Wen-tau Yih, Hao Ma, and Madian Khabsa. 2020. [Language models as fact checkers?](#) In *Proceedings of the Third Workshop on Fact Extraction and VERification (FEVER)*, pages 36–41, Online. Association for Computational Linguistics.
- Ying Lin, Joe Hoover, Gwenyth Portillo-Wightman, Christina Park, Morteza Dehghani, and Heng Ji. 2018. Acquiring background knowledge to improve moral value prediction. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 552–559. IEEE.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pre-training approach](#).
- Sahil Loomba, Alexandre de Figueiredo, Simon J Piatek, Kristen de Graaf, and Heidi J Larson. 2021. Measuring the impact of covid-19 vaccine misinformation on vaccination intent in the uk and usa. *Nature human behaviour*, 5(3):337–348.
- Alex Munyole Luvembe, Weimin Li, Shaohua Li, Fangfang Liu, and Guiqiong Xu. 2023. [Dual emotion based fake news detection: A deep attention-weight update approach](#). *Information Processing & Management*, 60(4):103354.

- Syed Ishfaq Manzoor, Jimmy Singla, and Nikita. 2019. Fake news detection using machine learning approaches: A systematic review. *2019 3rd International Conference on Trends in Electronics and Informatics (ICOEI)*, pages 230–234.
- Drew B Margolin, Aniko Hannak, and Ingmar Weber. 2018. Political fact-checking on twitter: When do corrections have an effect? *Political Communication*, 35(2):196–219.
- Giovanni Da San Martino, Stefano Cresci, Alberto Barron-Cedeno, Seunghak Yu, Roberto Di Pietro, and Preslav Nakov. 2020. [A survey on computational propaganda detection](#).
- Kitti Nagy and Jozef Kapusta. 2021. [Improving fake news classification using dependency grammar](#). *PLOS ONE*, 16(9):1–22.
- OpenAI. 2023. [Gpt-4 technical report](#).
- The pandas development team. 2020. [pandas-dev/pandas: Pandas](#).
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Kellin Pelrine, Meilina Reksoprodjo, Caleb Gupta, Joel Christoph, and Reihaneh Rabbany. 2023. [Towards reliable misinformation mitigation: Generalization, uncertainty, and gpt-4](#).
- Kashyap Papat, Subhabrata Mukherjee, Jannik Strötgen, and Gerhard Weikum. 2016. [Credibility assessment of textual claims on the web](#). In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management, CIKM '16*, page 2173–2178, New York, NY, USA. Association for Computing Machinery.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah A. Smith, and Mike Lewis. 2023. [Measuring and narrowing the compositionality gap in language models](#).
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Hossein Rajaby Faghihi, Quan Guo, Andrzej Uszok, Aliakbar Nafar, and Parisa Kordjamshidi. 2021. [DomiKnowS: A library for integration of symbolic domain knowledge in deep learning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 231–241, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Hossein Rajaby Faghihi, Aliakbar Nafar, Chen Zheng, Roshanak Mirzaee, Yue Zhang, Andrzej Uszok, Alexander Wan, Tanawan Premisri, Dan Roth, and Parisa Kordjamshidi. 2023. [Gluecons: A generic benchmark for learning under constraints](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(8):9552–9561.
- Neetu Rani, Prasenjit Das, and Amit Kumar Bhardwaj. 2022. [Rumor, misinformation among web: A contemporary review of rumor detection techniques during different web waves](#). *Concurrency and Computation: Practice and Experience*, 34(1):e6479.
- Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. [Truth of varying shades: Analyzing language in fake news and political fact-checking](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2931–2937, Copenhagen, Denmark. Association for Computational Linguistics.
- Jon Roozenbeek, Claudia R Schneider, Sarah Dryhurst, John Kerr, Alexandra LJ Freeman, Gabriel Recchia, Anne Marthe Van Der Bles, and Sander Van Der Linden. 2020. Susceptibility to misinformation about covid-19 around the world. *Royal Society open science*, 7(10):201199.
- Gautam Kishore Shahi and Durgesh Nandini. 2020. [Fakecovid - A multilingual cross-domain fact check news dataset for COVID-19](#). *CoRR*, abs/2006.11343.
- 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.
- Victor Suarez-Lledo and Javier Alvarez-Galvez. 2021. [Prevalence of health misinformation on social media: Systematic review](#). *J Med Internet Res*, 23(1):e17187.
- Briony Swire-Thompson and David Lazer. 2019. Public health and online misinformation: challenges and recommendations. *Annual review of public health*, 41:433–451.
- Nguyen Vo and Kyumin Lee. 2019. [Learning from fact-checkers: Analysis and generation of fact-checking language](#). In *Proceedings of the*

42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'19, page 335–344, New York, NY, USA. Association for Computing Machinery.

Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *science*, 359(6380):1146–1151.

William Yang Wang. 2017. "liar, liar pants on fire": A new benchmark dataset for fake news detection. *CoRR*, abs/1705.00648.

Yuxi Wang, Martin McKee, Aleksandra Torbica, and David Stuckler. 2019. Systematic literature review on the spread of health-related misinformation on social media. *Social science & medicine*, 240:112552.

Sue Lim Wei Peng and Jingbo Meng. 2023. [Persuasive strategies in online health misinformation: a systematic review](#). *Information, Communication & Society*, 26(11):2131–2148.

Zhiwei Yang, Jing Ma, Hechang Chen, Hongzhan Lin, Ziyang Luo, and Yi Chang. 2022. [A coarse-to-fine cascaded evidence-distillation neural network for explainable fake news detection](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2608–2621, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. [React: Synergizing reasoning and acting in language models](#). In *The Eleventh International Conference on Learning Representations*.

Sy Yuliani, Shahrin Sahib, and ZR MFBAF. 2019. Hoax news classification using machine learning algorithms. *International Journal of Engineering and Advanced Technology*, 9(2):3938–3944.

Xuan Zhang and Wei Gao. 2023. [Towards llm-based fact verification on news claims with a hierarchical step-by-step prompting method](#).

Radim Řehůřek and Petr Sojka. 2010. [Software framework for topic modelling with large corpora](#). In *Proceedings of LREC 2010 workshop New Challenges for NLP Frameworks*, pages 46–50, Valletta, Malta. University of Malta.

A. GPT Experiments

We evaluate the GPT-4² and GPT-3³ models in zero-shot and in-context learning setting on the test

²gpt-4 model version as of 9/20/2023

³text-davinci-003

set of our dataset. We used the prompts shown in Table 6 to evaluate models for the misinformation detection task. The labels are the names of the persuasive strategies concatenated with ". " (dot space). Table 7 analysis reveals GPT-4's superior performance over GPT-3 in misinformation detection, which could be due to its sophisticated architecture and updated training data. Remarkably, GPT-4 shows a significant improvement when leveraging *GT* ground truth data, which indicates capability in misinformation detection requiring nuanced context discernment. Its enhanced ability to handle *Claim + Article* input combinations is pivotal in detecting complex misinformation narratives. Importantly, the best performance for both models is achieved when combining *GT*, *Claim*, and *Article*, highlighting the crucial role of reliable ground truth in improving model accuracy. In conclusion, the observed results indicate a significant advancement in the capacity of language models, evidenced by the progression from GPT-3 to GPT-4. Moreover, the findings reinforce the assertion that the incorporation of persuasive strategies can contribute significantly to the efficiency of misinformation detection.

B. Persuasive Strategies Detection using GPT

In this study, we conduct a preliminary assessment of the capability of GPT-3.5 to identify persuasive writing strategies. Using the ChatGPT website interface for our experiment, we first provided the model with the descriptions of the persuasive writing strategies provided in the subsections of "Persuasive strategy themes" in [Wei Peng and Meng \(2023\)](#). Subsequently, we tasked GPT-3.5 with detecting the persuasive strategies used in a given text. An example of the interactions is shown in Figure 5. In this experiment, we evaluated the model's performance based solely on its ability to identify the existence of a persuasive writing strategy within the article without delving into specific sentences or segments. We repeated this experiment twice and reported the average accuracy of the GPT-3.5 model on this task in Table 8. According to the findings in Table 8, GPT-3.5 exhibits better capability in detecting certain strategies, such as *Distinctive Linguistic Features* and *Narrative with Details*, while it faces challenges with more nuanced tasks like identifying *Rhetorical Tricks*, which are inherently difficult to pinpoint. It is important to note that the results summarized in Table 8 encompass the entirety of the articles included in our dataset. These results further signify the difficulty of persuasive writing strategy detection in automated systems despite being very helpful for analyzing misinformation.

Inputs	Prompt
Zero-shot _{Claim}	Assuming you are a journalism and communication expert. Is this claim correct? claim: {claim} You must answer to the best of your knowledge. Give me one word answer "Yes" or "No"?
Zero-shot _{Labels}	Assuming you are a journalism and communication expert. Is a claim correct if our communication expert reported that an article supporting that claim follows the below persuasive strategies. {labels} You must answer to the best of your knowledge. Give me one word answer "Yes" or "No"?
Zero-shot _{Article}	Assuming you are a journalism and communication expert. Is this article correct? article: {article} You must answer to the best of your knowledge. Give me one word answer "Yes" or "No"?
Zero-shot _{Claim} +Labels	Assuming you are a journalism and communication expert. Is this claim correct? claim: {claim} our communication expert reported that the article supporting this claim follows the below persuasive strategies {labels} You must answer to the best of your knowledge. Give me one word answer "Yes" or "No"?
Zero-shot _{Claim} +Article+Labels	Assuming you are a journalism and communication expert. Is this claim correct? claim: {claim} We have found this article supporting the claim: article: {article} our communication expert reported that the article supporting this claim follows the below persuasive strategies {labels} You must answer to the best of your knowledge. Give me one word answer "Yes" or "No"?
GPT In-context Learning	Here we show examples of persuasive strategy detection. Examples below show text spans with their corresponding persuasive strategy: span ₁ ⇒ ['persuasive strategy ₁ ', 'persuasive strategy ₂ ,..'] span ₂ ⇒ ['persuasive strategy ₁ ', 'persuasive strategy ₂ ,..'] ... Mark the sentences in the text below with persuasive strategy labels: {article} Don't mark a sentence with one strategy more than once. [Then after getting the labels from the GPT model] Given the labeled persuasive strategies in the above-supporting article, is the following claim correct? Claim: {claim}. Give me a Yes or No answer with your chain of thoughts.

Table 6: GPT model prompts. *Labels* refers to either ground-truth persuasive strategy labels (*GT*) or prediction of our Task 2 models (*Pred_{context}*).

C. Models Hyperparameters

Models hyperparameters of each experiment are shown in Tables 9, 10, and 11.

D. Misinformation Detection Examples

Figure 6 is an example of how the correct persuasive strategy prediction helps with detection. In Figure 7, we can see that our persuasive strat-

egy labeler has detected more strategies than the ground truth, which has affected the model's decision to wrongfully mislabel as misinformation. Figure 8 portrays an example in which wrong strategy prediction did not help the model. However, we can see that if we use ground truth strategy labels, the model can correct its prediction and correctly label the claim.

Input Source	GPT3		GPT4	
	F1 _{Micro}	F1 _{Macro}	F1 _{Micro}	F1 _{Macro}
GT	0.653	0.612	0.696	0.510
Pred _{high}	0.592	0.372	0.653	0.395
Claim	0.837	0.779	0.837	0.791
Article	0.816	0.805	0.878	0.865
Claim + Article	0.857	0.851	0.898	0.891
Claim + Pred _{high}	0.755	0.686	0.776	0.688
Claim + GT	0.837	0.791	0.898	0.878
Claim + Article + Pred _{high}	0.653	0.648	0.918	0.910
Claim + Article + GT	0.857	0.779	0.939	0.933

Table 7: Misinformation Detection results (Task 3) of GPT models with *source* and prompts shown in Table 6. The results are an average of three runs with 0 temperature.

User: Please learn from the persuasive strategies in online health misinformation below. Persuasive strategies are used to create persuasive information or misinformation. [Persuasive strategies Descriptions].

GPT: [GPT Response]

User: Using the 12 strategies you learned above, identify the persuasive strategies in the following text: [A popular Chinese restaurant in Pretoria has been granted permission to sell Dog Meat to its customers. The restaurant successfully argued that the banning of consuming dog meat, which they pointed out, is in violation of their religious and cultural rights. The restaurant has been granted provisional permission to slaughter, sell and consume dog meat, while awaiting the Supreme Court’s final decision of their case. The restaurant is allowed to slaughter up to 35 dogs per week, under the conditions that the dogs are killed properly and humanely. One animal rights activist called the decision "shocking" and vowed to fight the restaurant in court to block the decision. "We can not believe that such actions would be allowed on South African soil. Dogs are precious and should not be used as food. We will not allow the courts to follow through with giving these savages permission to consume man’s best friend," said Jon Smit an animal rights activist.]

GPT: From the provided text, the following persuasive strategies are present: Highlighting Uncertainty and Risk: The text highlights potential ethical and health concerns by discussing the consumption of dog meat and referencing the need to kill dogs "properly and humanely." Exploiting Science’s Limitations: The text uses the term "properly and humanely" to imply that the practice might not be humane, questioning the treatment of animals. Inappropriate Use of Scientific Evidence: The claim about the restaurant being allowed to slaughter up to 35 dogs per week is presented without clear scientific evidence or context. Rhetorical Tricks: The text employs emotive language such as "shocking," "savages," and "man’s best friend" to evoke strong emotional reactions and sway opinion. Emotional Appeals: Emotional language, such as describing dogs as "precious" and "man’s best friend," aims to evoke strong emotional responses from the readers. Establishing Legitimacy: The text cites an animal rights activist named Jon Smit to strengthen the argument against the practice and lend credibility to the opposing viewpoint. It’s important to note that while these strategies are used in the text, the authenticity and accuracy of the information presented are uncertain. Always verify information from reliable sources before believing or sharing claims.

Figure 5: An example of persuasive strategy detection using GPT-3.5.

E. Dataset Statistics

Tables 12 and 13 present the distribution statistics of sentences marked with labels shown in Figure 1 considering their parent-children relationship. It should be mentioned that a parent is not necessarily tagged with any of its children’s tags, while a child tag must always be accompanied by its

parent tag.

F. Label Normalization

As mentioned in Section 5.1 our *Health subset* labels are highly imbalanced. Since our data is heterogeneous and is gathered from multiple domains, it contains various similar labels. We normalize our

Persuasive Writing Strategy	F-1 Score
Narrative with Details	0.534
Using Anecdotes and Personal Experience as Evidence	0.285
Distrusting Government or Corporations	0.407
Politicizing Health Issues	0.341
Highlighting Uncertainty and Risk	0.365
Exploiting Science's Limitations	0.067
Inappropriate Use of Scientific Evidence	0.477
Rhetorical Tricks	0.221
Biased Reasoning to Make a Conclusion	0.457
Emotional Appeals	0.528
Distinctive Linguistic Features	0.599
Establishing Legitimacy	0.869
Average	0.429

Table 8: GPT-3.5 performance on level-2 persuasive writing strategy detection. The results are at an article level instead of the sentence level compared to Task 2, which is sentence-based. The results are the average of two runs.

<p>Claim: HIV infected needle at McDonald's in the ball pit</p> <p>Article: Email: (important) McDonalds Playground Some of you might not be parents, but you may have nieces, nephews, grandchildren or friends with children this will pertain to you too. As I read the following, my heart sank. I urge each and every one of you to pass this on to as many people as you can. I cannot stress how important this is! This is very disturbing news. In addition to the following true story, I will also add that my own sons were playing in the ball pit at Discovery Zone one day. One son lost his watch, and was very upset...</p> <p>GT: emotional appeals, call to action, establishing legitimacy, fear, using anecdotes and personal experience as evidence, narrative with details verified to be false, narrative with details, narrative with details verified to be false, narrative with details, highlighting uncertainty and risk, establishing legitimacy, legitimate persuasive techniques: rhetorical question, call to action, establishing legitimacy</p> <p>Pred: establishing legitimacy, call to action, emotional appeals, shifting hypothesis, using anecdotes and personal experience as evidence, narrative with details, narrative with details details not verified, using anecdotes and personal experience as evidence, using anecdotes and personal experience as evidence, narrative with details, using anecdotes and personal experience as evidence, narrative with details, narrative with details verified to be false, narrative with details, using anecdotes and personal experience as evidence, narrative with details verified to be false, narrative with details, establishing legitimacy, highlighting uncertainty and risk, narrative with details, narrative with details, narrative with details, establishing legitimacy, citing seemingly credible source, legitimate persuasive techniques: rhetorical question, establishing legitimacy, call to action</p> <p>Claim: No Misinformation</p> <p>Claim + Ground Truth: Misinformation</p> <p>Claim + Pred: Misinformation</p> <p>Actual Label: Misinformation</p>

Figure 6: Example of when Persuasive Strategies helped with misinformation Detection

labels using Table 14 to normalize similar labels such as *false* and *mostly false* into the same category.

G. Result of Persuasive Strategy Detection

The Precision, Recall, and F1-score of our best-performing model, along with the number of instances for each label, are presented for levels 1 to 4 in the following tables: level 1 in Tables 15, 16, and 17; level 2 in Tables 18, 19, and 20; level 3 in Tables 21, 22, and 23; and level 4 in Tables 24, 25,

Claim: In July 2017, it was announced that Finnish researchers would begin human trials of a Type 1 diabetes vaccine in 2018

Article: Finnish diabetes vaccine trials to start in 2018 Finnish researchers have developed an experimental vaccine against type 1 diabetes, and tested it successfully on mice. An enterovirus, which likely causes type 1 diabetes. Image: Lääketieteen ja biotieteiden tiedekunta / Tampereen yliopisto A vaccine for type 1 diabetes developed by Finnish researchers will be tested on mainly Finnish human subjects in late 2018, researchers announced on Tuesday...

GT: establishing legitimacy, citing source to establish legitimacy source verified to be credible, establishing legitimacy, surface credibility markers - simply claiming authority or credibility, establishing legitimacy, citing source to establish legitimacy source verified to be credible

Pred: surface credibility markers, exaggeration, establishing legitimacy, surface credibility markers, surface credibility markers - simply claiming authority or credibility, establishing legitimacy, less robust evidence or outdated evidence verified, establishing legitimacy, surface credibility markers - simply claiming authority or credibility, establishing legitimacy, ellipses, exaggerated/excessive usage of punctuation marks, citing source to establish legitimacy source not verified, inappropriate use of scientific evidence, lack of citation for evidence, highlighting uncertainty and risk, biased reasoning to make a conclusion, inappropriate analogy or false connection, narrative with details, narrative with details details not verified, establishing legitimacy, inappropriate use of scientific evidence, distinctive linguistic features, anxiety

Predictions:
Claim: Misinformation
Claim + Ground Truth: No Misinformation
Claim + Pred: Misinformation
Actual Label: No Misinformation

Figure 7: Example of when Persuasive Strategies Prediction did NOT help with the misinformation Detection

Parameter	Value
Batch Size	20
Learning Rate	5e-5
RoBERTa Max length	512
Epoch #: Level 1	40
Epoch #: Level 2-4	60
Linear Layer Size: Level 1	256
Linear Layer Size: Level 2	1024
Linear Layer Size: Level 3	1024
Linear Layer Size: Level 4	1024
Train/Test Split	80% / 20%
Dropout	50%

Table 9: Persuasive Strategy Models Hyperparameters

and 26, respectively.

Parameter	Value
Batch Size	10
Learning Rate	5e-5
RoBERTa Max length	512
Epoch #	12
Linear Layer Size	None
Train/Dev Split	90% / 10%
Dropout	50%

Table 10: RoBERTa model's hyperparameters for *pomt* baseline comparison

Parameter	Value
Batch Size	10
Learning Rate	5e-5
RoBERTa Max length	512
Epoch #	12
Linear Layer Size	None
Train/Test Split	80% / 20%
Dropout	50%
RoBERTa frozen layers #	181

Table 11: Misinformation Detection Model Hyperparameters

Claim: Doctors are now warning parents to never use baby wipes because they contain a "chemical" called methylisothiazolinone.

Article: Doctor's Warning: Never Clean Your Child With Baby Wipes No Matter What Many parents wouldn't think twice about what's in their child's baby wipes. More and more research is surfacing that baby wipes aren't the best thing for our children's skin. In fact, some doctors are recommending that parents cease using baby wipes on their children entirely because of the risk they are taking each time they do so. According to NBC News, a study has surfaced that confirms baby wipes are not safe to use on your children due to one key ingredient...

GT: title of article as clickbait, distinctive linguistic features, exaggeration/absolute language, rhetorical tricks, ellipses, exaggerated/excessive usage of punctuation marks, distinctive linguistic features, establishing legitimacy, legitimate persuasive techniques: rhetorical question, surface credibility markers - simply claiming authority or credibility, establishing legitimacy, fear, emotional appeals, surface credibility markers - simply claiming authority or credibility, establishing legitimacy, rhetorical tricks, exaggeration/absolute language, establishing legitimacy, citing source to establish legitimacy source not verified, highlighting uncertainty and risk, citing source to establish legitimacy source not verified, establishing legitimacy, inappropriate use of scientific and other evidence - out of context verified, inappropriate use of scientific and other evidence, surface credibility markers - words associated with nature or healthiness, establishing legitimacy, establishing legitimacy, citing source to establish legitimacy source not verified, rhetorical tricks, exaggeration/absolute language, highlighting uncertainty and risk, hope, emotional appeals

Pred: distinctive linguistic features, bolded words or underline, establishing legitimacy, citing seemingly credible source, establishing legitimacy, surface credibility markers, establishing legitimacy, surface credibility markers - simply claiming authority or credibility, establishing legitimacy, ellipses, exaggerated/excessive usage of punctuation marks, surface credibility markers, establishing legitimacy, ellipses, exaggerated/excessive usage of punctuation marks, citing source to establish legitimacy source not verified, narrative with details, narrative with details details not verified, narrative with details, narrative with details details verified to be true, narrative with details, narrative with details, narrative with details, establishing legitimacy, narrative with details details not verified, narrative with details, ellipses, exaggerated/excessive usage of punctuation marks, narrative with details, narrative with details details not verified, establishing legitimacy, ellipses, exaggerated/excessive usage of punctuation marks, citing source to establish legitimacy source not verified

Predictions:

Claim: Misinformation

Claim + Ground Truth: Misinformation

Claim + Pred: No Misinformation

Actual Label: Misinformation

Figure 8: Example of Persuasive Strategies Prediction when it worsened the Detection

Full Annotation Table		
Parent	Child	# of Instances
Narrative w/ Details	Total	555
	Verified to be False	253
	Verified to be True	182
	Not Verified	91
Anecdotes and Personal Experience as Evidence	Total	239
Distrusting Govt. or Corporations	Total	104
	Financial Motive	64
Politicizing Health Issues	Total	54
	Freedom of Choice and Agency	12
	Ingroup vs. Outgroup	16
	Political Figures/ Argument	13
	Religion or Ideology	10
Highlighting Uncertainty or Risk	Total	177
Exploiting Science's Limitations	Total	7
Inappropriate use of Scientific or other Evidence	Total	294
	Out of Context/ Verified	117
	Less robust or outdated Evidence/ Verify	51
Rhetorical Tricks	Total	60
	Exaggeration/ Absolute Language	50
	Selective Omission	0
Biased Reasoning to make a Conclusion	Total	382
	Inappropriate Analogy or False connection	43
	Wrong Cause/Effect	18
	Lack of Evidence or Incomplete Evidence	247
	Evidence does not support Conclusion	34
	Shifting Hypothesis	32
Emotional Appeals	Total	208
	Fear	98
	Anger	40
	Hope	19
	Anxiety	36
Distinctive Linguistic Features	Total	516
	Uppercase Words	233
	Linguistic Intensifier	1
	Clickbait Title	47
	Bolded, underlined, or italicized content	143
	Excessive usage of Punctuation Marks	131
Establishing Legitimacy	Total	1045
	Citing Source To Establish Legitimacy	399
	Legitimate Persuasive Techniques	191
	Surface Credibility Markers	346
	Call to Action	132

Table 12: Exhaustive Table of all Parent and Child Annotations, along with the number of examples annotated for each.

Sub Tag Table			
Parent	Child	Sub	# of Instances
Establishing Legitimacy	Citing Source To Establish Legitimacy	Total	399
		Verified to be Credible	86
		Verified to not be credible	29
		Not Verified	274
		Verified to be Made Up	8
	Legitimate Persuasive Techniques	Total	191
		Rhetorical Question	181
		Metaphor	0
		Humor	8
	Surface Credibility Markers	Total	346
		Medical/Scientific Jargon	76
		Words Associated w/ Health	44
		Words Associated w/ Uncertainty	0
		Simply Claiming Authority or Credibility	234
Call to Action	Total	132	

Table 13: Exhaustive Table of all Parent, Child, and Sub Annotations, along with the number of annotated examples for each.

Norm. Label	Original Label
True	mostly true, truth!, true, in-the-green, mostly truth!
False	mostly false, fiction!, false

Table 14: Table of all labels and our normalized versions of them

No Context				
	Precision	Recall	F1	# of Instances
No Persuasive Writing Strategy	0.647	0.586	0.615	633
Has Persuasive Writing Strategy	0.690	0.742	0.715	784

Table 15: Precision, Recall, F1 Score, and Support for the Test Set Level 1 of the persuasive strategy label's hierarchy with no context.

Low Context				
	Precision	Recall	F1	# of Instances
No Persuasive Writing Strategy	0.657	0.670	0.664	633
Has Persuasive Writing Strategy	0.729	0.718	0.724	784

Table 16: Precision, Recall, F1 Score, and Support for the Test Set Level 1 of the persuasive strategy label's hierarchy with low context.

High Context				
	Precision	Recall	F1	# of Instances
No Persuasive Writing Strategy	0.679	0.660	0.669	633
Has Persuasive Writing Strategy	0.732	0.747	0.739	784

Table 17: Precision, Recall, F1 Score, and Support for the Test Set Level 1 of the persuasive strategy label's hierarchy with high context.

No Context				
	Precision	Recall	F1	# of Instances
Narrative w/ Details	0.629	0.629	0.629	105
Emotional Appeals	0.352	0.292	0.319	65
Distinctive Linguistic Features	0.537	0.797	0.642	182
Establishing Legitimacy	0.585	0.414	0.485	266
Using Anecdotes and Personal Exp. as Evidence	0.795	0.318	0.455	110
Distrusting Govt. or Corps.	0.583	0.356	0.442	59
Politicizing Health Issues	0.000	0.000	0.000	17
Highlighting Uncertainty and Risk	0.300	0.240	0.267	25
Inappropriate Use of Scientific Evidence	0.316	0.154	0.207	78
Rhetorical Tricks	0.273	0.240	0.255	25
Biased Reasoning to Make Conclusion	0.164	0.123	0.141	73

Table 18: Precision, Recall, F1 Score, and Support for the Test Set Level 2 of the persuasive strategy label's hierarchy with no context.

Low Context				
	Precision	Recall	F1	# of Instances
Narrative w/ Details	0.607	0.705	0.652	105
Emotional Appeals	0.250	0.231	0.240	65
Distinctive Linguistic Features	0.621	0.720	0.667	182
Establishing Legitimacy	0.572	0.523	0.546	266
Using Anecdotes and Personal Exp. as Evidence	0.745	0.373	0.497	110
Distrusting Govt. or Corps.	0.400	0.373	0.386	59
Politicizing Health Issues	0.154	0.118	0.133	17
Highlighting Uncertainty and Risk	0.474	0.360	0.409	25
Inappropriate Use of Scientific Evidence	0.338	0.333	0.335	78
Rhetorical Tricks	0.119	0.200	0.149	25
Biased Reasoning to Make Conclusion	0.260	0.178	0.211	73

Table 19: Precision, Recall, F1 Score, and Support for the Test Set Level 2 of the persuasive strategy label's hierarchy with low context.

High Context				
	Precision	Recall	F1	# of Instances
Narrative w/ Details	0.664	0.676	0.670	105
Emotional Appeals	0.261	0.185	0.216	65
Distinctive Linguistic Features	0.586	0.863	0.698	182
Establishing Legitimacy	0.554	0.560	0.557	266
Using Anecdotes and Personal Exp. as Evidence	0.733	0.300	0.426	110
Distrusting Govt. or Corps.	0.309	0.356	0.331	59
Politicizing Health Issues	0.364	0.235	0.286	17
Highlighting Uncertainty and Risk	0.526	0.400	0.455	25
Inappropriate Use of Scientific Evidence	0.429	0.308	0.358	78
Rhetorical Tricks	0.125	0.040	0.061	25
Biased Reasoning to Make Conclusion	0.260	0.178	0.211	73

Table 20: Precision, Recall, F1 Score, and Support for the Test Set Level 2 of the persuasive strategy label's hierarchy with high context.

No Context				
	Precision	Recall	F1	# of Instances
Fear	0.091	0.042	0.057	24
Anger	0.000	0.000	0.000	25
Hope	0.000	0.000	0.000	4
Anxiety	0.000	0.000	0.000	7
Uppercase Words	0.824	0.864	0.843	81
Title of Article as Clickbait	0.286	0.750	0.414	8
Bolded, Underline or Italicized	0.480	0.758	0.587	62
Exaggerated Usage of Punctuation Marks	0.581	0.783	0.667	46
Citing Source to Establish Legitimacy	0.459	0.487	0.473	80
Legitimate Persuasive Techniques	0.618	0.596	0.607	57
Surface Credibility Markers	0.519	0.456	0.485	90
Call to Action	0.486	0.375	0.424	48
Verified to be False	0.240	0.222	0.231	27
Details Verified to be True	0.607	0.270	0.374	63
Details Not Verified	0.200	0.250	0.222	16
Financial Motivation	0.400	0.205	0.271	39
Out of Context_Verified	0.000	0.000	0.000	34
Verified	0.286	0.167	0.211	12
Exaggeration/Absolute Language	0.500	0.105	0.174	19
Inappropriate Analogy or False Connection	0.167	0.100	0.125	10
Claims Without Evidence	0.233	0.194	0.212	36
Evidence Does Not Support Conclusion	0.000	0.000	0.000	10
Shifting Hypothesis	0.000	0.000	0.000	14

Table 21: Precision, Recall, F1 Score, and Support for the Test Set Level 3 of the persuasive strategy label's hierarchy with no context.

Low Context				
	Precision	Recall	F1	# of Instances
Fear	0.167	0.042	0.067	24
Anger	0.000	0.000	0.000	25
Hope	0.000	0.000	0.000	4
Anxiety	1.000	0.143	0.250	7
Uppercase Words	0.697	0.852	0.767	81
Title of Article as Clickbait	0.286	0.750	0.414	8
Bolded, Underline or Italicized	0.586	0.823	0.685	62
Exaggerated Usage of Punctuation Marks	0.635	0.717	0.673	46
Citing Source to Establish Legitimacy	0.427	0.512	0.466	80
Legitimate Persuasive Techniques	0.551	0.667	0.603	57
Surface Credibility Markers	0.528	0.311	0.392	90
Call to Action	0.655	0.396	0.494	48
Verified to be False	0.550	0.407	0.468	27
Details Verified to be True	0.852	0.365	0.511	63
Details Not Verified	0.222	0.125	0.160	16
Financial Motivation	0.400	0.103	0.163	39
Out of Context_Verified	0.200	0.029	0.051	34
Verified	0.000	0.000	0.000	12
Exaggeration/Absolute Language	0.000	0.000	0.000	19
Inappropriate Analogy or False Connection	0.000	0.000	0.000	10
Claims Without Evidence	0.346	0.250	0.290	36
Evidence Does Not Support Conclusion	0.000	0.000	0.000	10
Shifting Hypothesis	0.125	0.071	0.091	14

Table 22: Precision, Recall, F1 Score, and Support for the Test Set Level 3 of the persuasive strategy label's hierarchy with low context.

High Context				
	Precision	Recall	F1	# of Instances
Fear	0.211	0.167	0.186	24
Anger	0.000	0.000	0.000	25
Hope	0.000	0.000	0.000	4
Anxiety	0.167	0.143	0.154	7
Uppercase Words	0.724	0.877	0.793	81
Title of Article as Clickbait	0.222	0.500	0.308	8
Bolded, Underline or Italicized	0.621	0.871	0.725	62
Exaggerated Usage of Punctuation Marks	0.621	0.783	0.692	46
Citing Source to Establish Legitimacy	0.519	0.500	0.510	80
Legitimate Persuasive Techniques	0.603	0.667	0.633	57
Surface Credibility Markers	0.527	0.433	0.476	90
Call to Action	0.485	0.333	0.395	48
Verified to be False	0.237	0.667	0.350	27
Details Verified to be True	0.704	0.302	0.422	63
Details Not Verified	0.636	0.438	0.519	16
Financial Motivation	0.526	0.256	0.345	39
Out of Context_Verified	0.000	0.000	0.000	34
Verified	0.000	0.000	0.000	12
Exaggeration/Absolute Language	0.125	0.053	0.074	19
Inappropriate Analogy or False Connection	0.000	0.000	0.000	10
Claims Without Evidence	0.417	0.278	0.333	36
Evidence Does Not Support Conclusion	0.000	0.000	0.000	10
Shifting Hypothesis	0.000	0.000	0.000	14

Table 23: Precision, Recall, F1 Score, and Support for the Test Set Level 3 of the persuasive strategy label's hierarchy with high context.

No Context				
	Precision	Recall	F1	# of Instances
source verified to be credible	0.250	0.067	0.105	30
source verified to not be credible	0.000	0.000	0.000	6
source not verified	0.389	0.477	0.429	44
rhetorical question	0.674	0.574	0.620	54
medical or scientific jargon	0.400	0.167	0.235	12
words associated with nature or healthiness	0.400	0.167	0.235	12
simply claiming authority or credibility	0.596	0.456	0.517	68

Table 24: Precision, Recall, F1 Score, and Support for the Test Set- Level 4 of the persuasive strategy label's hierarchy with no context.

Low Context				
	Precision	Recall	F1	# of Instances
source verified to be credible	0.375	0.100	0.158	30
source verified to not be credible	0.000	0.000	0.000	6
source not verified	0.412	0.477	0.442	44
rhetorical question	0.544	0.685	0.607	54
medical or scientific jargon	1.000	0.250	0.400	12
words associated with nature or healthiness	0.000	0.000	0.000	12
simply claiming authority or credibility	0.647	0.324	0.431	68

Table 25: Precision, Recall, F1 Score, and Support for the Test Set- Level 4 of the persuasive strategy label's hierarchy with low context.

High Context				
	Precision	Recall	F1	# of Instances
source verified to be credible	0.000	0.000	0.000	30
source verified to not be credible	0.000	0.000	0.000	6
source not verified	0.383	0.409	0.396	44
rhetorical question	0.597	0.796	0.683	54
medical or scientific jargon	1.000	0.167	0.286	12
words associated with nature or healthiness	0.000	0.000	0.000	12
simply claiming authority or credibility	0.667	0.324	0.436	68

Table 26: Precision, Recall, F1 Score, and Support for the Test Set- Level 4 of the persuasive strategy label's hierarchy with high context.