

PCoT: Persuasion-Augmented Chain of Thought for Detecting Fake News and Social Media Disinformation

Arkadiusz Modzelewski^{1,2}, Witold Sosnowski², Tiziano Labruna¹,
Adam Wierzbicki², Giovanni Da San Martino¹

¹University of Padua, Italy

²Polish-Japanese Academy of Information Technology, Poland

Correspondence: contact@amodzelewski.com

Abstract

Disinformation detection is a key aspect of media literacy. Psychological studies have shown that knowledge of persuasive fallacies helps individuals detect disinformation. Inspired by these findings, we experimented with large language models (LLMs) to test whether infusing persuasion knowledge enhances disinformation detection. As a result, we introduce the Persuasion-Augmented Chain of Thought (PCoT), a novel approach that leverages persuasion to improve disinformation detection in zero-shot classification. We extensively evaluate PCoT on online news and social media posts. Moreover, we publish two novel, up-to-date disinformation datasets: EUDisinfo and MultiDis. These datasets enable the evaluation of PCoT on content entirely unseen by the LLMs used in our experiments, as the content was published after the models' knowledge cut-offs. We show that, on average, PCoT outperforms competitive methods by 15% across five LLMs and five datasets. These findings highlight the value of persuasion in strengthening zero-shot disinformation detection.

1 Introduction

The spread of disinformation in digital communication poses significant risks to the state of democracy by shaping public opinion, reinforcing ideological divides, and fostering distrust in political institutions. (Jungherr and Rauchfleisch, 2024; Farhall et al., 2019; Brummette et al., 2018). The growing accessibility of digital media, coupled with reduced funds for traditional fact-checking efforts and the rise of alternatives like Birdwatch on Platform X (formerly Twitter), underscores the urgent need for complementary disinformation detection systems (Saeed et al., 2022; Allen et al., 2022). Traditional supervised detection methods, which rely on human-annotated data, face challenges in generalization and the scarcity of labeled data. This reinforces the need for zero-shot detection systems.

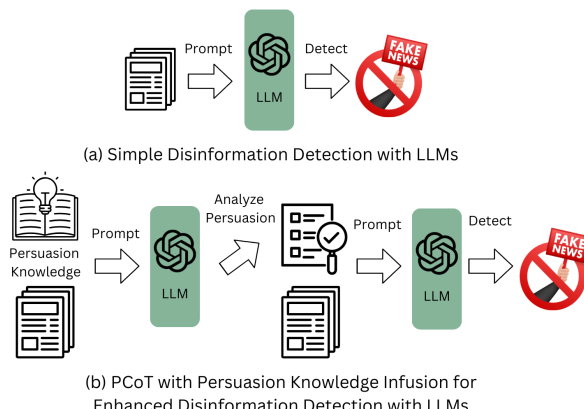


Figure 1: The comparison between detecting disinformation with LLMs in a simple zero shot setting and detecting with PCoT and infused knowledge about persuasion.

A critical aspect of disinformation is its coexistence with manipulation and persuasion to mislead audiences (Modzelewski et al., 2024; Chen et al., 2021). Psychological studies show that teaching individuals to recognize persuasive fallacies improves their ability to distinguish between real and fake news (Hruschka and Appel, 2023). Building on this, we explored whether infusing knowledge of persuasion into generative LLMs enhances disinformation detection.

As a result, we present **Persuasion-Augmented Chain of Thought (PCoT)**, a novel zero-shot method leveraging persuasion signals to improve disinformation detection that more effectively addresses generalization and annotated data scarcity challenges compared to supervised models. PCoT operates through a two-stage process where the LLM first identifies and analyzes persuasion within a given text, using infused knowledge. This analysis is then utilized in subsequent reasoning to determine the presence of disinformation. By augmenting models' decision-making process with persuasion knowledge, PCoT achieves significant gains

in detection performance across multiple datasets.

We conducted experiments on five datasets covering fake news and social media disinformation to evaluate our method rigorously. We evaluated PCoT on two novel datasets, MultiDis and EUDisinfo, which contain up-to-date articles from 2024 onwards, ensuring they were not part of the pre-training data for any tested LLMs. The **Multitopic Disinformation** is a high-quality dataset developed with fact-checking experts with prior experience in debunking organizations accredited by the International Fact-Checking Network¹. For a comprehensive evaluation, we also used three publicly available datasets containing texts before the knowledge cutoff of all tested models.

For evaluation, we selected three top-performing methods on zero-shot disinformation detection from Lucas et al. (2023) and adapted them to incorporate our PCoT approach. Using five different LLMs, we demonstrate that PCoT delivers significant performance improvements over chosen competitive methods.

Our main contributions are as follows:

- We introduce a novel Persuasion-Augmented Chain of Thought (PCoT) method that significantly enhances the ability of generative LLMs to detect disinformation in a zero-shot setting.
- We provide a thorough analysis of the effectiveness of our method across various datasets, including fake news and disinformative social media posts from X (formerly Twitter).
- We introduce two novel datasets, MultiDis and EUDisinfo, to further assess our method. These datasets consist of articles collected after the knowledge cutoff of the tested LLMs, allowing for a thorough and rigorous PCoT evaluation on texts that were entirely unseen by the models.
- We analyze how LLM-predicted persuasion affects disinformation detection effectiveness.

We release the final prompts and the codebase².

2 Datasets

To ensure robust performance of our method across diverse data conditions and inspired by the work of Lucas et al. (2023), we designed our evaluation to address potential dataset overlap with LLM pretraining. We tested our method on two dataset

¹The International Fact-Checking Network gives accreditation to debunking and fact-checking organizations that sign its code of principles. See <https://www.poynter.org/ifcn/>

²Repository with data, prompts and codebase: <https://github.com/ArkadiusDS/PCoT>

types: (i) prior-cutoff datasets, which may contain pretraining content, and (ii) two novel datasets of articles published after the models' knowledge cutoff. This setup enables a rigorous evaluation of our PCoT method on potential pretraining content and entirely new information. Moreover, we evaluated our method on social media posts versus longer articles, such as news.

2.1 Prior-Cutoff Datasets

The following datasets, published before January 1, 2024, may overlap with the models' training data.

- **CoAID** – A dataset for COVID-19 misinformation detection, comprising 4k+ news articles and 1k+ social posts, all annotated with ground-truth labels (Cui and Lee, 2020).
- **ISOT Fake News** – A dataset of 44k+ fake and truthful articles from reputable and unreliable sources, identified via Politifact³ (Ahmed et al., 2018, 2017).
- **ECTF** – An extended version of CTF (Paka et al., 2021) for detecting fake news on Platform X about COVID-19, with additional data to improve early-stage detection (Bansal et al., 2021).

2.2 Novel Post-Cutoff Datasets

2.2.1 MultiDis

The **Multitopic Disinformation Dataset** comprises nearly 2,000 English articles on European and global disinformation. It has been created by researchers from multiple European universities to support disinformation detection research.

Annotation Process The annotation process involved four key stages:

1. **Methodology and Data Preparation** – Researchers, fact-checking and debunking experts developed a robust methodology and guidelines before collecting a database of articles.
2. **In-Depth Training** – A three-day hybrid training led by the most experienced fact-checking expert aimed to deliver in-depth on-site training to all European teams while ensuring accessibility for remote annotators. Each team was assigned two supervisors, usually a disinformation researcher. The training concluded with an initial annotation round, reviewed and discussed by a fact-checking expert. These preliminary annotations were excluded from the final dataset to maintain high quality.

³PolitiFact is a nonprofit fact-checking project by the Poynter Institute.

3. **Article Annotation** – Independent annotation by a less experienced annotator and a supervisor.
4. **Final Evaluation** – The supervisor reviewed both annotations and resolved disagreements through discussion when necessary. A senior fact-checking expert contributed when needed. If consensus was unattainable, the article was labeled *Hard-to-say*.

Appendix A shows dataset and annotation details.

Data Sources We selected diverse sources to ensure access to both reliable and unreliable content, categorizing each as *Reliable*, *Unreliable*, or *Mixed*. A team of experts evaluated sources through consensus, thoroughly analyzing the source’s regularly published content and cross-checking with established tools (e.g., Media Bias/Fact Check). The assigned categories were not revealed to annotators to prevent biases toward sources.

The MultiDis dataset includes a variety of sources: global news agencies, regional publications, thematic platforms, fact-checking organizations, and independent media. All used 44 distinct sources are freely accessible. To ensure transparency, we make these sources publicly available.

Thematic Category Before detailed analysis, articles were manually assigned to one of eight thematic categories. The selection of these topics was informed by the EU DisinfoLab report⁴ (Sessa, 2023). The categories are: (i) *Anti-Europeanism and Anti-Atlanticism*; (ii) *Anti-migration and Xenophobia*; (iii) *Climate Change and the Energy Crisis*; (iv) *Health*; (v) *Institutional and Media Distrust*; (vi) *Gender Issues*; (vii) *Ukraine War and Refugees*; (viii) *LGBT+*. Table 1 shows the distribution of articles by thematic category in the MultiDis dataset.

Annotators, during the credibility evaluation, could label articles as *Inconsistent with the topic*, excluding them from further analysis to ensure high-quality topic assignments.

Credibility Analysis Annotators assessed each article using a debunking technique, auxiliary complemented by fact-checking, as defined by the NATO Strategic Communications Centre of Excellence (Pamment and Kimber, 2021).

Given an article, annotators analyze its content to determine whether it belongs to one of four categories. The main categories in our guidelines are:

⁴The EU DisinfoLab’s report, grounded in expert research from 20 countries across Europe, guarantees high quality and credibility.

Category	#DOC	#PERC
Anti-Europeanism & Anti-Atlanticism	219	11.4%
Anti-migration and Xenophobia	117	6.1%
Climate Change and the Energy Crisis	324	16.9%
Health	285	14.8%
Institutional and Media Distrust	317	16.5%
Gender Issues	97	5.0%
Ukraine War and Refugees	361	18.8%
LGBT+	202	10.5%

Table 1: Number of articles (#DOC) per thematic category and their (#PERC) percentage in MultiDis dataset.

Credible Information, *Disinformation*, with the latter following the European Commission’s High-Level Expert Group: *Disinformation is false, inaccurate or misleading information designed, presented, and promoted to intentionally cause public harm or for profit* (de Cock Buning, 2018). This definition has also been adopted in other disinformation studies (Modzelewski et al., 2024; Sosnowski et al., 2024). Two additional labels, *Hard-to-say* and *Inconsistent with the Topic*, were respectively assigned to articles where annotators did not reach a consensus or where the content did not match the assigned topic. Articles labeled with these or published before January 2024 were excluded from experiments.

Bias Prevention and Data Quality Our guidelines require each article to be annotated independently by two experts to minimize bias. Given the time demands of the annotation process, only two independent evaluations per article were guaranteed. Supervisors provided the third final annotation by reviewing the two previous annotations. However, supervisors were instructed to resolve uncertainties through discussions, and the lead fact-checking expert provided clarification when needed. These discussions helped ensure consistency among annotators and reduced human errors and bias. We achieved full agreement in the first two rounds for 86.78% of articles, with the remainder undergoing a more detailed third analysis.

Note: We publicly release the complete dataset, including annotations from all three rounds.

2.2.2 EUDisinfo

We introduce the EUDisinfo dataset, collected with usage of the EUvsDisinfo database⁵, which com-

⁵<https://EUvsDisinfo.eu/disinformation-cases/>

prises 18,464 disinformation cases⁶. EUvsDisinfo is an EU initiative dedicated to identifying, analyzing, and countering pro-Kremlin disinformation. Each entry concisely summarizes a disinformation case, along with links to the original misleading content and credible sources debunking the claims. Since EUvsDisinfo provides predefined evaluation for each disinformation case as either *credible* or *disinformation*, we did not conduct additional annotation. The EUvsDisinfo database comprises articles published in multiple languages, some previously analyzed in Leite et al. (2024a). However, as all articles in that study date before 2024, this dataset was unsuitable for our research. To address this limitation, we independently curated a collection of approximately 400 English articles published in 2024 or later.

Category	MultiDis	EUDisinfo
Credible Information	65.3%	67.1%
Disinformation	32.8%	32.9%

Table 2: Percentage of articles per main credibility category in MultiDis and EUDisinfo datasets.

To collect English news article content, we leveraged the *Trafilatura* tool (Barbaresi, 2021), which efficiently scrapes web content while preserving article structure. Additionally, we employed *Selenium* (Selenium, 2024) to navigate and extract HTML pages and *Beautiful Soup 4* (Richardson and Katz, 2024) to parse article content.

Table 2 presents the percentage of articles per main credibility category in our two datasets. More detailed statistics are provided in Appendix B.1.

3 Proposed Method

In this section, we introduce the Persuasion-Augmented Chain of Thought (PCoT) method, which leverages persuasion to enhance zero-shot disinformation detection using generative LLMs.

3.1 Persuasion-Augmented Chain of Thought

Empirical studies have shown that persuasion is an integral part of disinformation (Chen et al., 2021). Insights from psychological research highlight the potential of leveraging persuasion knowledge to more effectively discern between fake and credible news (Hruschka and Appel, 2023). Inspired by this, we propose the Persuasion-Augmented Chain of Thought. The PCoT method employs a two-stage

reasoning process that improves LLM’s disinformation detection by persuasion knowledge infusion.

In the first stage, an LLM is prompted to perform multi-faceted reasoning by analyzing persuasion strategies (see Figure 2) within the text. The second stage performs the disinformation detection task, enriched by the previously generated analysis of persuasion strategies. Figure 1 presents a simplified comparison between traditional zero-shot disinformation detection using LLMs and our PCoT method. Final prompt templates for each stage of our PCoT method are available in Appendix F.

3.2 Persuasion Detection Step

In the first stage LLM performs multifaceted reasoning by tackling the multi-class, multi-label task of detecting persuasion strategies, along with contextual question answering by explaining persuasion usage within each text. The persuasion detection task can be formally represented as follows: The model M takes as input the text T , the impersonation I_P , the infused knowledge K_P and guidelines G_P . Here, I_P establishes the context and overrides alignment tuning, while K_P encapsulates knowledge about a predefined set of high-level persuasion strategies P , and guidelines G_P that determine the task and specify the structure of the expected response. This combined input is represented as $X = (T, I_P, K_P, G_P)$, where the set of persuasion strategies is given by $P = \{p_1, p_2, \dots, p_k\}$. For each text, the model generates an output in a structured textual format that can be decoded into a *JSON*-like dictionary. This output contains, for each persuasion strategy $p_i \in P$, two components: a binary label y_{p_i} (‘Yes’ or ‘No’) indicating the presence of p_i in the text, and an explanation E_{p_i} justifying the prediction. The output can be formally expressed as:

$$A_T = \{p_i : (y_{p_i}, E_{p_i}) \mid p_i \in P\}. \quad (1)$$

The model M generates the output A_T by leveraging the combined input X , capturing both the text and infused persuasion knowledge:

$$A_T \sim M(T, I_P, K_P, G_P). \quad (2)$$

This stage leverages the capabilities of generative LLMs to integrate knowledge about persuasion into the reasoning process. The rationale for our approach is based on the observations that explanations can enhance the robustness of the final prediction (He et al., 2024), and that previous works

⁶Database size recorded as of February 11, 2025.

have shown that incorporating explanations can improve zero-shot classification performance (Menon et al., 2022).

3.3 Disinformation Detection Step

In the final stage of the PCoT method, LLM performs zero-shot binary classification on each input text. Formally, the model M evaluates the input text T to detect disinformation. It processes the combined input $X = (T, I_D, A_T, G_D)$, where I_D defines the impersonation that establishes the context, A_T provides the persuasion analysis from the first stage of PCoT, and G_D defines the task and specifies the structure of the expected response. The model then generates the output, Y_T indicating whether T contains disinformation ('Yes' or 'No').

$$Y_T \sim M(T, I_D, A_T, G_D). \quad (3)$$

We explored the zero-shot setting as many studies have shown that zero-shot prompting of LLMs like GPT-4 can outperform supervised models like BERT in detecting disinformation (Pelrine et al., 2023; Bang et al., 2023; Hassan and Lee, 2020). In addition, Lucas et al. (2023) demonstrated that fine-tuning BERT on different datasets and testing on unseen data leads to worse performance than zero-shot with LLMs. We confirm these findings on our data, as outlined in the Appendix H.

4 Experiments and Evaluation

We created five test sets by randomly selecting texts from each dataset. To evaluate our PCoT method’s ability to detect disinformation in data unseen by LLMs, we used two novel test sets, MultiDis and EUDisinfo. These test sets contain only articles published from 2024 onward, ensuring that the content was not part of any LLM training data. Each test set contained 400-500 articles or posts. Appendix B.2 provides statistics on the composition of the test datasets used in our experiments.

We conducted all experiments on five different LLMs: *GPT 4o Mini*, *Llama 3.1 8B*, *Claude 3 Haiku*, *Llama 3.3 70B*, and *Gemini 1.5 Flash*. To ensure the most deterministic results possible, we set the hyperparameter temperature to 0 in each model. Appendix D includes more details about the models used, including knowledge cutoff dates and the rationale behind our choice.

PCoT was evaluated using the F_1 score. To assess the significance of its difference from competitive methods, we used McNemar’s test, which

suits binary tasks comparing two methods on the same dataset (Dror et al., 2018; Dietterich, 1998). This statistical test has been widely applied in NLP (Card et al., 2020; Blitzer et al., 2006).

Attack on reputation [AR]	- the argument does not address the topic itself but targets the participant (personality, experience, etc.) to question and/or undermine their credibility. The object of the argumentation can also refer to a group of individuals, an organization, an object, or an activity.
Justification [J]	- the argument is made of two parts, a statement and an explanation or appeal, where the latter is used to justify and/or to support the statement.
Simplification [S]	- the argument excessively simplifies a problem, usually regarding the cause, the consequence, or the existence of choices.
Distraction [D]	- the argument takes focus away from the main topic or argument to distract the reader.
Call [C]	- the text is not an argument, but an encouragement to act or to think in a particular way.
Manipulative wording [MW]	- the text is not an argument per se, but uses specific language, which contains words or phrases that are either non-neutral, confusing, exaggerating, loaded, etc., in order to impact the reader emotionally.

Figure 2: Persuasion strategies used in our experiments. A detailed description of the techniques associated with these strategies can be found in Appendix C.

4.1 Persuasion Detection Step

To enhance first stage of the PCoT method we created prompts with infused persuasion knowledge. The knowledge applied within the prompt is based on the taxonomy presented by Piskorski et al. (2023a,c). This taxonomy categorizes persuasion techniques into six strategies (shortcuts in brackets): *Attack on reputation [AR]*, *Justification [J]*, *Simplification [S]*, *Distraction [D]*, *Call [C]*, *Manipulative wording [MW]* (definitions in Figure 2). Using a well-established taxonomy enabled a thorough first-stage evaluation since datasets annotated with it are widely used (Dimitrov et al., 2024; Piskorski et al., 2023b).

Method	F_1 Micro
DMT	↑9% 0.722 ±0.035
DTAT	↑4% 0.689 ±0.042
Base MT	0.664±0.030

Table 3: Average F_1 micro ($\pm std$, over five LLMs) for three methods evaluated in the first stage of the PCoT method. Percentage changes are computed relative to the *Base MT* method. The *DMT* variant is selected as the final best-performing method for this stage.

To develop the most effective prompt for detecting persuasive strategies, we conducted extensive experiments on the SemEval 2023 dataset (Piskorski et al., 2023b), using 536 English news articles

with ground truth on persuasion strategies and five LLMs. We used F_1 micro as the evaluation metric for this stage, following its use in a closely related task at SemEval 2023 (Piskorski et al., 2023b).

We tested various prompts, including:

- Detailed Multitask (DMT) - a single prompt for detecting all strategies and their explanations. Prompt with infused knowledge about persuasion strategies and their definitions (see Figure 2), and the specific techniques with definitions (see Appendix C) that fall under each strategy. These techniques are categorized according to the taxonomy proposed by Piskorski et al. (2023a,c).
- Detailed One Task At a Time (DTAT) - individual prompts for binary detection and explanations per strategy, infusing the same knowledge as DMT but divided into six parts as there are six persuasion strategies.
- Base Multitask - our baseline single prompt for detecting all strategies. It does not incorporate persuasion knowledge but simply lists strategy names and prompts identification of those present in the text. This served as our starting point.

As shown in Table 3, the DMT method achieved the highest F_1 micro score, outperforming our baseline prompt by 9%. As a result, DMT was used in the first stage of our final PCoT method. Our experiments revealed important finding that:

Finding

Using a single prompt to identify all persuasion strategies was more effective than separate prompts for each strategy’s binary classification.

The persuasion detection step provides an analysis that includes binary labels and explanations. To assess the impact of these explanations, we also evaluated PCoT without them. Testing PCoT without explanations showed that including LLM-generated insights improved performance.

This step establishes the foundation for the second stage of PCoT by analyzing the persuasion signals present in the input text. Appendix E presents more details on tested prompts, evaluation results and rationale behind chosen taxonomy.

4.2 Disinformation Detection Step

For disinformation detection stage, we selected three top-performing competitive methods based on an extensive evaluation by Lucas et al. (2023), specifically those that excelled on human-annotated datasets (Cui and Lee, 2020; Shu et al., 2020). We outline the three methods below:

- *VaN* - A vanilla prompt serving as a fundamental baseline, offering concise instructions to LLMs (Lucas et al., 2023).
- *Z-CoT* - Extends *VaN* with a prompt encouraging step-by-step reasoning, inspired by Kojima et al. (2022)’s findings on zero-shot reasoning.
- *DeF-SpeC* - Emphasizes contextual, deductive, and abductive reasoning (Lucas et al., 2023), addressing LLM limitations in inductive and multi-step reasoning (Bang et al., 2023).

The chosen competitive methods served as baselines, allowing us to evaluate the effectiveness of PCoT. We then adapted these methods to our PCoT approach by modifying prompts to incorporate persuasion analysis from the first stage. This approach enabled us to determine whether the PCoT method is sensitive to prompt variations or exhibits consistent behavior. For a rigorous evaluation, we conducted experiments on five datasets covering various themes and genres, such as news and social media posts. The diverse selection of datasets allows us to assess PCoT’s generalizability.

Method	F_1 Score
PCoT	↑15% 0.815 ±0.027
PCoT Single Step	↑8% 0.765 ±0.072
Base	0.711±0.055

Table 4: Average F_1 ($\pm std$, over five LLMs) for *PCoT* (two-stage) and *PCoT Single Step*, which uses one prompt for simultaneous persuasion analysis and disinformation detection. Percentage changes are computed relative to the *Base* method.

To demonstrate the need for two-stage PCoT, we tested a more straightforward single-step approach, where LLMs analyzed persuasion and detected disinformation simultaneously. As shown in Table 4 single-step PCoT outperformed the baseline by 8%, while the two-stage method provided an additional significant 7% improvement.

5 Results and Discussion

General Overview The results of our experiments, presented in Table 5, compare the performance of our PCoT method with baseline approaches. PCoT significantly improves performance, achieving an average F_1 score of 0.815, about a 15% improvement over the baselines. McNemar’s test confirmed that, across all models and methods, PCoT consistently achieves significantly better results on overall data at the 0.01 significance

	Overall		Articles		Posts		Prior Cutoff		Post Cutoff	
	Base	PCoT	Base	PCoT	Base	PCoT	Base	PCoT	Base	PCoT
<i>GPT 4o Mini</i>										
VaN	0.759	0.845 ↑11%	0.788	0.885 ↑12%	0.700	0.762 ↑9%	0.742	0.830 ↑12%	0.790	0.874 ↑11%
Z-CoT	0.765	0.846 ↑11%	0.801	0.884 ↑10%	0.696	0.767 ↑10%	0.747	0.835 ↑12%	0.801	0.869 ↑8%
DeF-SpeC	0.772	0.834 ↑8%	0.816	0.867 ↑6%	0.690	0.766 ↑11%	0.742	0.813 ↑10%	0.832	0.875 ↑5%
<i>Gemini 1.5 Flash</i>										
VaN	0.681	0.810 ↑19%	0.673	0.843 ↑25%	0.695	0.748 ↑8%	0.683	0.778 ↑14%	0.679	0.875 ↑29%
Z-CoT	0.689	0.808 ↑17%	0.681	0.838 ↑23%	0.703	0.752 ↑7%	0.670	0.777 ↑16%	0.687	0.872 ↑27%
DeF-SpeC	0.744	0.834 ↑12%	0.764	0.876 ↑15%	0.708	0.754 ↑6%	0.721	0.810 ↑12%	0.790	0.884 ↑12%
<i>Claude 3 Haiku</i>										
VaN	0.710	0.797 ↑12%	0.714	0.820 ↑15%	0.702	0.747 ↑6%	0.728	0.797 ↑9%	0.677	0.796 ↑18%
Z-CoT	0.588	0.774 ↑32%	0.601	0.800 ↑33%	0.550	0.716 ↑30%	0.565	0.767 ↑36%	0.626	0.786 ↑26%
DeF-SpeC	0.780	0.795 ↑2%	0.806	0.810 ↑0%	0.727	0.763 ↑5%	0.809	0.812 ↑0%	0.727	0.766 ↑5%
<i>Llama 3.3 70B</i>										
VaN	0.740	0.845 ↑14%	0.747	0.881 ↑18%	0.727	0.768 ↑6%	0.733	0.839 ↑14%	0.752	0.856 ↑14%
Z-CoT	0.722	0.843 ↑17%	0.725	0.878 ↑21%	0.718	0.770 ↑7%	0.707	0.837 ↑18%	0.750	0.855 ↑14%
DeF-SpeC	0.732	0.832 ↑14%	0.740	0.863 ↑17%	0.717	0.768 ↑7%	0.719	0.806 ↑12%	0.755	0.880 ↑17%
<i>Llama 3.1 8B</i>										
VaN	0.627	0.792 ↑26%	0.565	0.802 ↑42%	0.736	0.773 ↑5%	0.649	0.788 ↑21%	0.585	0.801 ↑37%
Z-CoT	0.660	0.791 ↑20%	0.623	0.804 ↑29%	0.725	0.764 ↑5%	0.670	0.789 ↑18%	0.638	0.795 ↑25%
DeF-SpeC	0.697	0.773 ↑11%	0.688	0.784 ↑14%	0.712	0.752 ↑6%	0.683	0.767 ↑12%	0.724	0.785 ↑8%
Average	0.711	0.815 ↑15%	0.715	0.842 ↑18%	0.700	0.758 ↑8%	0.705	0.803 ↑14%	0.721	0.838 ↑16%

Table 5: Results with F_1 scores for five LLMs. The *Base* columns shows the competitive method results, while the *PCoT* columns presents results for prompts adapted to the PCoT method. McNemar’s test confirmed that, across all models and methods, PCoT achieves significantly better results on *Overall* data at the 0.01 significance level.

level (see Appendix G.1 for more details).

PCoT significantly improves disinformation detection across various scenarios, including news articles, social media posts, and novel post-cutoff datasets. It achieves the most substantial improvement in articles, with a 18% increase. Additionally, PCoT shows a 16% improvement for post-cutoff datasets, leading to our next key finding:

Finding

Infusing persuasion knowledge into prompts improves generative LLMs’ disinformation detection, especially for long texts and data not seen during pretraining.

Better performance on unseen data confirms superior effectiveness on longer articles, as these datasets consist exclusively of such texts. We attribute PCoT’s improved effectiveness on articles to the greater prevalence of persuasive strategies in longer texts, which complicate disinformation detection even for humans (Peng et al., 2024), underscoring the need for persuasion knowledge. Furthermore, PCoT deliver the largest average improvement, about 18%, for the smallest model.

Impact of Persuasion As Figure 3 shows, at least one persuasion strategy was found in 92% of disinformation and in 72% of credible texts. These results suggest that persuasion is more commonly used in disinformation than in credible information, though a significant proportion of credible content also contains persuasion. The strongest correlation



Figure 3: Average percentage of persuasion strategies predicted across 5 models for disinformation (*DIS*) and reliable information (*REL*). *ALL* represents the percentage of instances with at least one detected persuasion strategy. Other abbreviations are explained in Figure 2.

is observed between disinformation and the prediction of four specific strategies: *Attack on reputation*, *Simplification*, *Distraction*, and *Manipulative wording*. In contrast, the remaining two strategies, namely *Justification* and *Call*, occur with similar frequencies in both disinformation and credible information. More analysis confirming these findings can be found in Appendix G.2. The comparable presence of *Call* and *Justification* in both disinformation and credible content may be explained by the broad applicability of the persuasion techniques they encompass. For instance, *Call* techniques like *Slogans* such as "Make America Great Again!" are highly persuasive but not inherently misleading, making them familiar across various types of content. Similarly, *Conversation Killers* like "That’s just your opinion" appear in discussions to shut

down debate rather than mislead. Likewise, *Justification* includes techniques often found in credible information. For instance, *Appeal to Authority* is a standard persuasion technique in legitimate discourse, where expert opinions are cited to support claims. Similarly, *Appeal to Popularity*, justifying an argument based on widespread acceptance can be found in factual contexts.

Model	Persuasion		No Persuasion	
	PCoT	Base	PCoT	Base
GPT 4o Mini	0.872 ↑ ± 0,006	0.824 ± 0,008	0.342 ↑ ± 0,025	0.305 ± 0,009
Gemini 1.5 Flash	0.844 ↑ ± 0,014	0.738 ± 0,036	0.444 ↑ ± 0,013	0.430 ± 0,007
Claude 3 Haiku	0.831 ↑ ± 0,014	0.756 ± 0,101	0.177 ↓ ± 0,043	0.295 ± 0,084
Llama 3.3 70B	0.871 ↑ ± 0,007	0.781 ± 0,007	0.409 ↑ ± 0,010	0.343 ± 0,006
Llama 3.1 8B	0.812 ↑ ± 0,008	0.679 ± 0,050	0.536 ↑ ± 0,014	0.494 ± 0,059
Average	0.847 ↑	0.753	0.392 ↑	0.368

Table 6: Results comparison across two subsets: *Persuasion*, containing texts with at least one predicted persuasion strategy, and *No Persuasion* texts with no predicted persuasion. The table reports the average F₁ score and standard deviation for each model across three different prompting methods.

In addition, we evaluated how the PCoT method enhances disinformation detection across two subsets of used datasets: one where at least one persuasion strategy was predicted and another where none was detected. As shown in Table 6, PCoT improves detection by an average of 12% in the persuasion-present subset and about 7% in the persuasion-absent subset (detailed results for individual prompting methods are provided in Appendix G.3). Our findings highlight that:

Finding

Detecting disinformation is particularly challenging in texts where no persuasion strategy has been predicted.

Persuasive strategies may introduce emotionally charged language, making deception more apparent when these strategies are analyzed carefully. In contrast, when persuasion is absent, false statements alone may evade detection (Sosnowski et al., 2024). In this scenario, fact-checking techniques become more crucial, and semantic analysis of the language alone may be insufficient.

6 Further Evaluation and Ablation Study

We present additional experiments: a comparison with other prompting methods (section 6.1), a comparison of PCoT against cutting-edge reasoning

models (section 6.2), and an ablation study to assess the impact of the definitions of the persuasion strategies to the overall performance (section 6.3).

6.1 Prompting Methods Comparison

We compare PCoT with other recent prompting methods, including CoT (Wei et al., 2022) in a zero-shot version (Z-CoT) (Lucas et al., 2023), Chain-of-Verification (CoVe) (Dhuliawala et al., 2024) and Rephrase and Respond (RaR) (Deng et al., 2023). As shown in Table 7, PCoT consistently outperforms these methods.

Model	Z-CoT	RaR	CoVe	PCoT
GPT 4o Mini	0.765	0.698	0.790	0.846
Gemini 1.5 Flash	0.689	0.573	0.736	0.808
Claude 3 Haiku	0.588	0.768	0.441	0.774
Llama 3.3 70B	0.722	0.657	0.835	0.843
Llama 3.1 8B	0.660	0.566	0.764	0.791

Table 7: Overall F₁ scores of different prompting methods on five datasets.

6.2 Evaluation Against Reasoning Models

To further evaluate our approach, we compared PCoT-enhanced models to OpenAI’s advanced reasoning models, *o1-mini* and *o3-mini*. Specifically, we selected the best-performing (*GPT-4o Mini*) and worst-performing (*LLaMA 3.1 8B*) models from our zero-shot disinformation detection experiments using PCoT (see Table 5) and compared them against the reasoning models.

As shown in Table 8, even the weakest model, when used with PCoT, outperforms both *o1-mini* and *o3-mini* in zero-shot disinformation detection. This highlights PCoT’s ability to boost reasoning performance, even in smaller models.

Model	Overall
GPT 4o Mini + PCoT	0.846
Llama 3.1 8B + PCoT	0.791
o3-mini	0.770
o1-mini	0.634

Table 8: Overall F₁ scores for PCoT-enhanced models vs. OpenAI reasoning models on five datasets.

6.3 PCoT Base Version and Ablation Results

To better understand the contribution of explicit persuasion knowledge in PCoT, we conducted an ablation study using a simplified base version, which provides in the prompt a general definition of persuasion avoiding to mention persuasion strategies.

Remarkably, even without detailed knowledge, this simplified version yields notable performance gains over baseline prompting methods across five datasets (see Table 9). Although the original variant of PCoT remains stronger, these findings underscore the role of persuasion-augmented reasoning in zero-shot disinformation detection.

Model	PCoT BV	Base
GPT 4o Mini	0.814 \uparrow \pm 0.007	0.765 \pm 0.007
Gemini 1.5 Flash	0.790 \uparrow \pm 0.014	0.705 \pm 0.034
Claude 3 Haiku	0.736 \uparrow \pm 0.013	0.693 \pm 0.097
Llama 3.3 70B	0.831 \uparrow \pm 0.007	0.731 \pm 0.009
Llama 3.1 8B	0.785 \uparrow \pm 0.011	0.661 \pm 0.035

Table 9: Comparison of average F₁ scores and standard deviations between Base prompts and PCoT without persuasion strategy augmentation. Results are shown for VaN, Z-CoT, and DeF-SpeC (as *Base*), and their adaptations for PCoT’s base version (*PCoT BV*).

7 Related Works

Disinformation Detection. Disinformation detection has become a focal research focus due to its increasing impact on digital communication and societal trust (Flew, 2019; Olan et al., 2024; Iosifidis and Nicoli, 2020; Martens et al., 2018). Traditional approaches have used machine learning and deep learning to analyze lexical, semantic, and engagement-based features (Aslam et al., 2021; Ali et al., 2022; Nguyen et al., 2020). Given the high stakes, explainability is crucial. Hybrid frameworks combining deep learning with feature-specific explanations enhance transparency, trust, and understanding in NLP applications (Hashmi et al., 2024; Reis et al., 2019; Cartwright et al., 2022; Shu et al., 2019). Recent research has focused on detecting disinformation in human-annotated and LLM-generated data (Lucas et al., 2023; Chen and Shu).

Limited annotated data has driven zero and few-shot learning, highlighting the adaptability of pre-trained transformers across tasks without domain-specific training (Sivarajkumar and Wang, 2023; Rizinski et al., 2023; Kumar et al., 2023; Casola et al., 2023). Researchers have shown that zero-shot detection with LLMs like GPT-4 can outperform supervised models like BERT in detecting disinformation (Pelrine et al., 2023; Bang et al., 2023; Hassan and Lee, 2020).

Datasets. High-quality data is crucial for disinformation detection research (D’Ulizia et al., 2021), including datasets focused on COVID-19 disinformation (Bansal et al., 2021; Cui and Lee,

2020), persuasion techniques (Da San Martino et al., 2019), short statements (Wang, 2017) and fake news articles (Ahmed et al., 2018, 2017; Shu et al., 2020). To the best of our knowledge, previous datasets are released without revealing the intermediate labels. In contrast, our MultiDis dataset includes annotations from each of the three annotation steps.

Persuasion in Disinformation. Different studies have shown that disinformation often uses persuasion and manipulation to mislead audiences (Modzelewski et al., 2024; Peng et al., 2023; Chen et al., 2021; Musi and Reed, 2022; Ward et al., 2022). First attempts to use persuasion as intermediate labels in healthcare misinformation detection within a few-shot scenario have shown promising potential (Kamali et al., 2022). Nevertheless, no prior research has proposed a structured method that integrates persuasion knowledge and is applicable across diverse models and datasets to improve zero-shot disinformation detection.

8 Conclusions

In this study, we present **Persuasion-Augmented Chain of Thought (PCoT)**, a novel zero-shot approach that enhances disinformation detection by integrating persuasion knowledge into the LLM reasoning process. By leveraging persuasion knowledge and LLM-generated analysis, PCoT improves zero-shot classification, demonstrating the value of utilizing persuasion in disinformation detection.

Alongside PCoT, we present two novel disinformation datasets: MultiDis and EUDisinfo. EUDisinfo was collected using the database created by an EU initiative dedicated to identifying and analyzing pro-Kremlin disinformation. In contrast, MultiDis is a high-quality dataset created with debunking and fact-checking experts. These datasets enabled a robust evaluation of PCoT on texts beyond the knowledge cutoff of tested LLMs (2024 onward).

Our experiments using cutting-edge LLMs and five different datasets show that PCoT outperforms competitive methods, achieving an average 15% improvement. PCoT enhances disinformation detection, particularly for longer texts, such as news articles, where it achieves a 18% improvement while also providing a significant 8% increase in accuracy for social media posts. We also identified four persuasion strategies that most correlate with disinformation: *Attack on reputation*, *Simplification*, *Distraction*, and *Manipulative wording*.

Limitations

Datasets and Annotation Our annotation methodology for the MultiDis dataset categorized articles into one of eight thematic areas. Additionally, EUDisinfo focuses on disinformation with pro-Kremlin propaganda, while other datasets cover COVID-19 and political news disinformation. Despite this broad thematic coverage, we can not claim that the data fully represents all disinformation. Furthermore, our evaluation of the PCoT method was limited to English datasets. We leave multilingual analysis for future research.

Biases Human annotation can be prone to subjectivity. To minimize bias, annotations in the MultiDis dataset were conducted in cooperation with experienced debunkers and fact-checkers. We developed comprehensive annotation guidelines and provided thorough training. Two independent annotators annotated each article, followed by a final third review by a senior supervisor. The supervisor consulted the initial annotator and/or a lead fact-checking expert to ensure accuracy and consistency when necessary. However, experiments involving externally annotated datasets inevitably inherited any biases in those sources.

Method Our method heavily relies on the taxonomy and knowledge introduced by Piskorski et al. (2023a). Specifically, we incorporate a fixed set of high-level persuasive strategies in the first stage and integrate them into the prompt. Nevertheless, it is a taxonomy created by the Joint Research Centre, a scientific institution closely associated with the European Commission and used in many studies that confirm its high quality (Barrón-Cedeño et al., 2024; Dimitrov et al., 2024; Piskorski et al., 2023c; Szwoch et al., 2024; Leite et al., 2024b). Moreover, to the best of our knowledge, it is the only taxonomy extensively used for article annotation, which was crucial for our evaluation. A promising direction for future work is dynamically selecting persuasion techniques based on their relevance to disinformation detection. Nonetheless, we consider our current approach a crucial foundation for exploring dynamic selection, which we leave for future research.

Ethics

Datasets and Annotation The MultiDis and EU-Disinfo datasets consist entirely of publicly available data with no copyright restrictions. They do

not include any personally identifiable information and have been used exclusively for research purposes. The datasets will be released under the CC BY-NC-ND 4.0 license. Furthermore, our data collection protocol was reviewed and approved by an ethics board.

Crowdsourcing was not used at any stage of data collection or annotation. All annotators were employed by universities and fairly compensated. The annotation process remained entirely independent, free from political or commercial influence. Each team was overseen by two experienced supervisors and had direct access to a lead fact-checking expert for additional guidance.

Computational resources Leveraging large language models often requires substantial computational resources, which can contribute to environmental concerns (Strubell et al., 2020). However, our approach minimized computational demand as we relied on inference rather than training models from scratch. Most of our work conducted via usage of APIs, with no direct control over the computational resources involved. Additionally, we performed fine-tuning only on small BERT models. The computing infrastructure used for this research was acquired by the university specifically for research and educational purposes.

Acknowledgements

This research was supported by the projects: Infotester4Education (full title: Development and implementation of AI education methods and digital tools supporting tackling disinformation, number: 2023-2-PL01-KA220-HED-000180856) within the framework of Cooperation Partnership for Higher Education, ERASMUS+, and EUonAIR project (number 101177370, ERASMUS-EDU-2024-EUR-UNIV-1), within the framework of the EUonAIR Centre of Excellence in Responsible AI in Education, co-funded by the European Commission.

Giovanni Da San Martino would like to thank the Qatar National Research Fund, part of Qatar Research Development and Innovation Council (QRDI), for funding this work by grant NPRP14C0916-210015. He also would like to thank the European Union under the National Recovery and Resilience Plan (NRRP), Mission 4 Component 2 Investment 1.3 - Call for tender No. 341 of March 15, 2022 of Italian Ministry of University and Research – NextGenerationEU; Code

PE00000014, Concession Decree No. 1556 of October 11, 2022 CUP D43C22003050001, Progetto "SEcurity and RIghts in the CyberSpace (SERICS) - Spoke 2 Misinformation and Fakes - DEcision support systEm foR cybeR intelligENCE (Deterrence) for also funding this work.

References

- Hadeer Ahmed, Issa Traore, and Sherif Saad. 2017. Detection of online fake news using n-gram analysis and machine learning techniques. In *Intelligent, Secure, and Dependable Systems in Distributed and Cloud Environments: First International Conference, IS-DDC 2017, Vancouver, BC, Canada, October 26-28, 2017, Proceedings 1*, pages 127–138. Springer.
- Hadeer Ahmed, Issa Traore, and Sherif Saad. 2018. Detecting opinion spams and fake news using text classification. *Security and Privacy*, 1(1):e9.
- Abdullah Marish Ali, Fuad A Ghaleb, Bander Ali Saleh Al-Rimy, Fawaz Jaber Alsolami, and Asif Irshad Khan. 2022. Deep ensemble fake news detection model using sequential deep learning technique. *Sensors*, 22(18):6970.
- Jennifer Allen, Cameron Martel, and David G Rand. 2022. Birds of a feather don't fact-check each other: Partisanship and the evaluation of news in twitter's birdwatch crowdsourced fact-checking program. In *Proceedings of the 2022 CHI conference on human factors in computing systems*, pages 1–19.
- Nida Aslam, Irfan Ullah Khan, Farah Salem Alotaibi, Lama Abdulaziz Aldaej, and Asma Khaled Al-dubaikil. 2021. Fake detect: A deep learning ensemble model for fake news detection. *complexity*, 2021(1):5557784.
- Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023. A multi-task, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–718.
- Rachit Bansal, William Scott Paka, Nidhi, Shubhashis Sengupta, and Tanmoy Chakraborty. 2021. Combining exogenous and endogenous signals with a semi-supervised co-attention network for early detection of covid-19 fake tweets. In *Pacific-Asia conference on knowledge discovery and data mining*, pages 188–200. Springer.
- Adrien Barbaresi. 2021. [Trafalatura: A Web Scraping Library and Command-Line Tool for Text Discovery and Extraction](#). In *Proceedings of the Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 122–131. Association for Computational Linguistics.
- Alberto Barrón-Cedeño, Firoj Alam, Julia Maria Struß, Preslav Nakov, Tanmoy Chakraborty, Tamer Elsayed, Piotr Przybyła, Tommaso Caselli, Giovanni Da San Martino, Fatima Haouari, et al. 2024. Overview of the clef-2024 checkthat! lab: checkworthiness, subjectivity, persuasion, roles, authorities, and adversarial robustness. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 28–52. Springer.
- John Blitzer, Ryan McDonald, and Fernando Pereira. 2006. Domain adaptation with structural correspondence learning. In *Proceedings of the 2006 conference on empirical methods in natural language processing*, pages 120–128.
- John Brummette, Marcia DiStaso, Michail Vafeiadis, and Marcus Messner. 2018. Read all about it: The politicization of “fake news” on twitter. *Journalism & Mass Communication Quarterly*, 95(2):497–517.
- Dallas Card, Peter Henderson, Urvashi Khandelwal, Robin Jia, Kyle Mahowald, and Dan Jurafsky. 2020. With little power comes great responsibility. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9263–9274.
- Barry Cartwright, Richard Frank, George Weir, and Karmvir Padda. 2022. Detecting and responding to hostile disinformation activities on social media using machine learning and deep neural networks. *Neural Computing and Applications*, 34(18):15141–15163.
- Silvia Casola, Tiziano Labruna, Alberto Lavelli, Bernardo Magnini, et al. 2023. Testing chatgpt for stability and reasoning: A case study using italian medical specialty tests. In *CLiC-it*.
- Canyu Chen and Kai Shu. Can llm-generated misinformation be detected? In *The Twelfth International Conference on Learning Representations*.
- Sijing Chen, Lu Xiao, and Jin Mao. 2021. Persuasion strategies of misinformation-containing posts in the social media. *Information Processing & Management*, 58(5):102665.
- Limeng Cui and Dongwon Lee. 2020. Coaid: Covid-19 healthcare misinformation dataset. *arXiv preprint arXiv:2006.00885*.
- 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.
- Madeleine de Cock Buning. 2018. *A multi-dimensional approach to disinformation: Report of the independent high level group on fake news and online disinformation*. Publications Office of the European Union.
- Yihe Deng, Weitong Zhang, Zixiang Chen, and Qunqun Gu. 2023. [Rephrase and respond: Let large language models ask better questions for themselves](#). Preprint, arXiv:2311.04205.
- Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. 2024. Chain-of-verification reduces hallucination in large language models. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 3563–3578.
- Thomas G Dietterich. 1998. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural computation*, 10(7):1895–1923.
- Dimitar Dimitrov, Firoj Alam, Maram Hasanain, Abul Hasnat, Fabrizio Silvestri, Preslav Nakov, and Giovanni Da San Martino. 2024. Semeval-2024 task 4: Multilingual detection of persuasion techniques in memes. In *Proceedings of the 18th International Workshop on Semantic Evaluation (SemEval-2024)*, pages 2009–2026.
- Rotem Dror, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. The hitchhiker’s guide to testing statistical significance in natural language processing. In *Proceedings of the 56th annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 1383–1392.
- Arianna D’Ulizia, Maria Chiara Caschera, Fernando Ferri, and Patrizia Grifoni. 2021. Fake news detection: a survey of evaluation datasets. *PeerJ Computer Science*, 7:e518.
- European Commission. 2022a. Communication from the commission to the european parliament, the council, the european economic and social committee and the committee of the regions on the european democracy action plan. *Publications Office of the European Union*. Available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52020DC0790>.
- European Commission. 2022b. The strengthened code of practice on disinformation 2022. *Publications Office of the European Union*. Available at: <https://digital-strategy.ec.europa.eu/en/library/2022-strengthened-code-practice-disinformation>.
- Kate Farhall, Andrea Carson, Scott Wright, Andrew Gibbons, and William Lukamto. 2019. Political elites’ use of fake news discourse across communications platforms. *International Journal of Communication*, 13:23.
- Terry Flew. 2019. Digital communication, the crisis of trust, and the post-global. *Communication research and practice*, 5(1):4–22.
- Ehtesham Hashmi, Sule Yildirim Yayilgan, Muhammad Mudassar Yamin, Subhan Ali, and Mohamed Abomhara. 2024. Advancing fake news detection: hybrid deep learning with fasttext and explainable ai. *IEEE Access*.
- Fuad Mire Hassan and Mark Lee. 2020. Political fake statement detection via multistage feature-assisted neural modeling. In *2020 IEEE International Conference on Intelligence and Security Informatics (ISI)*, pages 1–6. IEEE.
- Xuanli He, Yuxiang Wu, Oana-Maria Camburu, Pasquale Minervini, and Pontus Stenetorp. 2024. Using natural language explanations to improve robustness of in-context learning. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13477–13499.
- Timon MJ Hruschka and Markus Appel. 2023. Learning about informal fallacies and the detection of fake news: An experimental intervention. *PLoS One*, 18(3):e0283238.
- Petros Iosifidis and Nicholas Nicoli. 2020. *Digital democracy, social media and disinformation*. Routledge.
- Andreas Jungherr and Adrian Rauchfleisch. 2024. Negative downstream effects of alarmist disinformation discourse: Evidence from the united states. *Political Behavior*, pages 1–21.
- Danial Kamali, Joseph Romain, Huiyi Liu, Wei Peng, Jingbo Meng, and Parisa Kordjamshidi. 2022. Using persuasive writing strategies to explain and detect health misinformation. *arXiv preprint arXiv:2211.05985*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Puneet Kumar, Kshitij Pathania, and Balasubramanian Raman. 2023. Zero-shot learning based cross-lingual sentiment analysis for sanskrit text with insufficient labeled data. *Applied Intelligence*, 53(9):10096–10113.
- João A Leite, Olesya Razuvayevskaya, Kalina Bontcheva, and Carolina Scarton. 2024a. Euvdisinfo: A dataset for multilingual detection of prokremlin disinformation in news articles. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pages 5380–5384.

- João A Leite, Olesya Razuvayevskaya, Carolina Scarton, and Kalina Bontcheva. 2024b. A cross-domain study of the use of persuasion techniques in online disinformation. *arXiv preprint arXiv:2412.15098*.
- Jason Lucas, Adaku Uchendu, Michiharu Yamashita, Jooyoung Lee, Shaurya Rohatgi, and Dongwon Lee. 2023. Fighting fire with fire: The dual role of llms in crafting and detecting elusive disinformation. In *2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023*, pages 14279–14305. Association for Computational Linguistics (ACL).
- Bertin Martens, Luis Aguiar, Estrella Gomez-Herrera, and Frank Mueller-Langer. 2018. The digital transformation of news media and the rise of disinformation and fake news.
- Rakesh R Menon, Sayan Ghosh, and Shashank Srivastava. 2022. Clues: A benchmark for learning classifiers using natural language explanations. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6523–6546.
- Arkadiusz Modzelewski, Giovanni Da San Martino, Pavel Savov, Magdalena Wilczyńska, and Adam Wierzbicki. 2024. Mipd: Exploring manipulation and intention in a novel corpus of polish disinformation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 19769–19785.
- Elena Musi and Chris Reed. 2022. From fallacies to semi-fake news: Improving the identification of misinformation triggers across digital media. *Discourse & Society*, 33(3):349–370.
- Van-Hoang Nguyen, Kazunari Sugiyama, Preslav Nakov, and Min-Yen Kan. 2020. Fang: Leveraging social context for fake news detection using graph representation. In *Proceedings of the 29th ACM international conference on information & knowledge management*, pages 1165–1174.
- Femi Olan, Uchitha Jayawickrama, Emmanuel Ogiemwonyi Arakpogun, Jana Suklan, and Shaofeng Liu. 2024. Fake news on social media: the impact on society. *Information Systems Frontiers*, 26(2):443–458.
- William Scott Paka, Rachit Bansal, Abhay Kaushik, Shubhashis Sengupta, and Tanmoy Chakraborty. 2021. Cross-sean: A cross-stitch semi-supervised neural attention model for covid-19 fake news detection. *Applied Soft Computing*, 107:107393.
- James Pamment and Anneli Lindvall Kimber. 2021. *Fact-checking and debunking: a best practice guide to dealing with disinformation*. NATO Strategic Communication Centre of Excellence.
- Kellin Pelrine, Anne Imouza, Camille Thibault, Meilina Reksoprodjo, Caleb Gupta, Joel Christoph, Jean-François Godbout, and Reihaneh Rabbany. 2023. Towards reliable misinformation mitigation: Generalization, uncertainty, and gpt-4. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6399–6429.
- Wei Peng, Sue Lim, and Jingbo Meng. 2023. Persuasive strategies in online health misinformation: a systematic review. *Information, Communication & Society*, 26(11):2131–2148.
- Wei Peng, Jingbo Meng, and Barikisu Issaka. 2024. Navigating persuasive strategies in online health misinformation: an interview study with older adults on misinformation management. *Plos one*, 19(7):e0307771.
- Jakub Piskorski, Nicolas Stefanovitch, Valerie-Anne Bausier, Nicolo Faggiani, Jens Linge, Sopho Kharazi, Nikolaos Nikolaidis, Giulia Teodori, Bertrand De Longueville, Brian Doherty, et al. 2023a. News categorization, framing and persuasion techniques: Annotation guidelines. *European Commission, Ispra, JRC132862*.
- Jakub Piskorski, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2023b. Semeval-2023 task 3: Detecting the category, the framing, and the persuasion techniques in online news in a multilingual setup. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2343–2361.
- Jakub Piskorski, Nicolas Stefanovitch, Nikolaos Nikolaidis, Giovanni Da San Martino, and Preslav Nakov. 2023c. Multilingual multifaceted understanding of online news in terms of genre, framing, and persuasion techniques. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3001–3022.
- Julio CS Reis, André Correia, Fabrício Murai, Adriano Veloso, and Fabrício Benevenuto. 2019. Explainable machine learning for fake news detection. In *Proceedings of the 10th ACM conference on web science*, pages 17–26.
- Leonard Richardson and Jeremy Katz. 2024. **Beautiful soup 4**. A library for scraping information from web pages.
- Maryan Rizinski, Andrej Jankov, Vignesh Sankaradas, Eugene Pinsky, Igor Miskovski, and Dimitar Trajanov. 2023. Company classification using zero-shot learning. *arXiv preprint arXiv:2305.01028*.
- Mohammed Saeed, Nicolas Traub, Maelle Nicolas, Gianluca Demartini, and Paolo Papotti. 2022. Crowdsourced fact-checking at twitter: How does the crowd compare with experts? In *Proceedings of the 31st ACM international conference on information & knowledge management*, pages 1736–1746.
- Selenium. 2024. *Selenium Browser Automation*. Selenium. Accessed: 2024-09-11.
- M.G Sessa. 2023. Connecting the disinformation dots: insights, lessons, and guidance from 20 eu member states. <https://www.disinfo.eu/publications/connecting-the-disinformation-dots/>.

Kai Shu, Limeng Cui, Suhang Wang, Dongwon Lee, and Huan Liu. 2019. defend: Explainable fake news detection. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 395–405.

Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 2020. Fakenewsnet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big data*, 8(3):171–188.

Sonish Sivarajkumar and Yanshan Wang. 2023. Health-prompt: a zero-shot learning paradigm for clinical natural language processing. In *AMIA Annual Symposium Proceedings*, volume 2022, page 972.

Witold Sosnowski, Arkadiusz Modzelewski, Kinga Skorpowska, Jahna Otterbacher, and Adam Wierzbicki. 2024. Eu disinfotest: a benchmark for evaluating language models’ ability to detect disinformation narratives. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14702–14723.

Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2020. Energy and policy considerations for modern deep learning research. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 13693–13696.

Joanna Szwoch, Mateusz Staszko, Rafal Rzepka, and Kenji Araki. 2024. Limitations of large language models in propaganda detection task. *Applied Sciences*, 14(10):4330.

William Yang Wang. 2017. “liar, liar pants on fire”: A new benchmark dataset for fake news detection. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 422–426.

Katrina J Ward, Hamilton Link, Kiril Avramov, and Jean Goodwin. 2022. Identifying disinformation using rhetorical devices in natural language models. Technical report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States).

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

A Dataset and Annotation Details

A.1 Annotation Methodology and Guidelines

Our methodology and annotation guidelines were designed to standardize the assessment of articles for disinformation content, aiming to reduce subjectivity and enable comprehensive analysis. Utilizing these annotation guidelines, we analyzed numerous articles to identify disinformation. The methodology was developed in cooperation with analysts

(fact-checking and debunking experts) employed in the project based on their experience as experts, scientific knowledge available on the subject, and the experience of other institutions and organizations involved in research and detection of disinformation. The methodology improved throughout the project and subsequent testing to best reflect the disinformation environment. All authors of this methodology have at least three years of experience working for fact-checking or debunking organizations accredited by the International Fact-Checking Network. Moreover, our methodology and annotation guidelines draw on similar work on the annotation of disinformation, such as the guidelines presented by [Modzelewski et al. \(2024\)](#).

Main Assumptions of the Methodology Creating a uniform methodology and guidelines aims to guarantee the quality of the assessments made by annotators and minimize their subjectivity.

The analysis of articles is carried out mainly via the debunking technique, with the auxiliary use of the fact-checking technique. These terms for this methodology are defined in a manner analogous to the methodology developed for the NATO Strategic Communication Centre of Excellence ([Pamment and Kimber, 2021](#)). Fact-checking is the long-standing process of checking that all facts in a piece of writing, news article, or speech are correct. Debunking refers to exposing falseness or manipulating systematically and strategically (based on a chosen topic, classifications of selected techniques, narrative).

Preparation of Articles for Evaluation The first step is to select web portals from which articles on particular topics will be taken. Among them are both mainstream media and those presenting the alternative current. This is to ensure access to enough reliable as well as unreliable content. Each portal will be assigned to one of three categories, determining its credibility. This will be done by a team of experts by consensus. Assessing the credibility of a website requires an in-depth analysis of the content posted on it regularly, as well as checking it in reliable sources, including via the Media Bias/Fact Check search engine. The source’s rating will not be visible to annotators. The analysis consists in selecting the category that best suits a given domain:

- **Reliable** — sources that are reliable/publishing reliable content on a specific topic, in particular traditional news portals.

- **Unreliable** — sources publishing unreliable content, typically disinformation, e.g., all domains financed by the Kremlin, sites containing conspiracy theories, etc.
- **Mixed/Biased** — partially or potentially biased websites that may present false information on specific issues, e.g., typically political websites, and blog collections.

Thematic Category Before the analysis begins, articles will be assigned to eight topics. This will be done manually with the help of keywords through searches on selected web portals. Thematic categories were pre-defined. The selection of topics was based on EU DisinfoLab’s cross-cutting report on disinformation in Europe (Sessa, 2023). It is based on expert studies from 20 countries.

- Anti-Europeanism and anti-Atlanticism (anti-EU, anti-NATO)
- Anti-migration and xenophobia
- Climate change and the energy crisis
- Health (including COVID-19 and vaccines)
- Institutional and media distrust (public institutions)
- Gender-based disinformation
- Ukraine war and refugees
- Disinformation about LGBTQIA+

Content Analysis The next step requires analyzing the entire article’s content and recognizing whether the information is accurate or disinformative. If the article provides only factual information, it is marked as “credible information.” Selecting this category ends the assessment of the article. When information in the article is unreliable and misleads the recipients, content is considered disinformative. The unintentional dissemination of false information is known as misinformation. However, even unintentional dissemination of false information without the goal of manipulating recipients can fuel disinformation. Disinformation is particularly difficult to detect as the author’s intention is usually unspecified, and in most cases, it can only be presumed. Therefore, for this study, we assume that any form of false or manipulative information is considered disinformation.

For these guidelines, the definition of disinformation provided by the European Commission High-Level Group of Experts on False News and Disinformation on the Internet (HELG) will be used, as it covers all four aspects and does not exclude potentially harmful content presented in the form of political advertising or satire, as presented in the

EU Code of Practice. The definition is as follows (de Cock Buning, 2018):

“ All forms of false, inaccurate, or misleading information designed, presented, and promoted to intentionally cause public harm or for profit.”

However, a necessary supplement to this definition is taking into account the European Union Code of Practice on Disinformation, according to which disinformation is defined as: “verifiable false or misleading information which, cumulatively, (a) is created, presented and disseminated for economic gain or to intentionally deceive the public; and (b) may cause public harm, intended as threats to democratic political and policymaking processes as well as public goods such as the protection of EU citizens’ health, the environment, or security”. (European Commission, 2022b). The detected information must be verifiable, which means that it can be proved untrue, and, therefore, it cannot be, for example, a yet unproven theory or opinion, as long as it is not intended to mislead the recipients. In summary, disinformation is intentionally misleading by providing misleading or false information (European Commission, 2022a). Unlike disinformation, misinformation is *misleading information shared by people who do not recognize it as such* (de Cock Buning, 2018). However, as noted earlier, misinformation and disinformation are treated as a single category under “disinformation.”

When a given content is not verifiable (reliable/disinformative/misinformative), it is marked as the “Hard to say” category. Indicating this category ends the assessment the same as “Inconsistent with the topic”. Below, we present the main categories:

- Credible information
- Disinformation
- Hard to say
- Inconsistent with the topic

B Dataset Basic Statistics

This appendix provides additional details on the statistics of the two datasets created for this study. The second section presents statistics on the balance of the classes in the five test datasets. All datasets, including the test datasets used in our experiments, are available in our public repository⁷.

⁷Repository with data, prompts and codebase: <https://github.com/ArkadiusDS/PCoT>

B.1 MultiDis and EUDisinfo Statistics

Table 10 reports the number of articles within the three credibility categories in MultiDis dataset: *Credible Information*, *Disinformation*, and *Hard-to-say*. Additionally, Table 10 includes articles labeled as *Inconsistent with the Topic*. Articles in this category were excluded from further analysis. Similar statistics for the EUDisinfo dataset are presented in Table 11, which includes only two categories: *Credible Information* and *Disinformation*.

Category	#DOC	#PERC
Credible Information	1256	65.3%
Disinformation	630	32.8%
Hard-to-say	18	0.95%
Inconsistent with the Topic	18	0.95%

Table 10: Number of articles (#DOC) per each credibility evaluation category and their (#PERC) percentage in MultiDis dataset.

Category	#DOC	#PERC
Credible Information	241	67.1%
Disinformation	118	32.9%

Table 11: Number of articles (#DOC) per main credibility evaluation category and their (#PERC) percentage in the EUDisinfo test dataset.

B.2 Test Datasets Statistics

Table 12 presents the class distribution of *Disinformation* and *Credible Information* across the five test datasets. Table 13 shows the same distribution across different content categories and time-based splits, indicating that social media posts and prior cutoff texts contain a higher proportion of disinformation.

Dataset	Disinformation	Credible Information
CoAID	21%	79%
ECTF	41%	59%
EUDisinfo	33%	67%
ISOT Fake News	55%	45%
MultiDis	26%	74%

Table 12: Class distribution across evaluation datasets. The proportions reflect the nature of each dataset and its composition regarding disinformation and credible content.

Category	Disinformation	Credible Information
All Texts	35%	65%
Articles	33%	67%
Social Media Posts	41%	59%
Prior Cutoff	39%	61%
Post Cutoff	29%	71%

Table 13: Class distribution by text type and time period. Social media and pre-cutoff texts show a higher share of disinformation compared to articles and post-cutoff samples.

C Persuasion Strategies and Techniques

The six general persuasion strategies in our study are linked to specific persuasion techniques, as identified by Piskorski et al. (2023c,a). Definitions of these techniques are provided in the final prompt created for the first stage of PCoT method.

C.1 Attack on Reputation

- **Name Calling or Labelling:** a form of argument in which loaded labels are directed at an individual, group, object or activity, typically in an insulting or demeaning way, but also using labels the target audience finds desirable.
- **Guilt by Association:** attacking the opponent or an activity by associating it with another group, activity or concept that has sharp negative connotations for the target audience.
- **Casting Doubt:** questioning the character or personal attributes of someone or something in order to question their general credibility or quality.
- **Appeal to Hypocrisy:** the target of the technique is attacked on its reputation by charging them with hypocrisy/inconsistency.
- **Questioning the Reputation:** the target is attacked by making strong negative claims about it, focusing specially on undermining its character and moral stature rather than relying on an argument about the topic.

C.2 Justification

- **Flag Waving:** justifying an idea by exhaling the pride of a group or highlighting the benefits for that specific group.
- **Appeal to Authority:** a weight is given to an argument, an idea or information by simply stating that a particular entity considered as an authority is the source of the information.
- **Appeal to Popularity:** a weight is given to an argument or idea by justifying it on the basis that allegedly "everybody" (or the large majority) agrees with it or "nobody" disagrees with it.

- **Appeal to Values:** a weight is given to an idea by linking it to values seen by the target audience as positive.
- **Appeal to Fear, Prejudice:** promotes or rejects an idea through the repulsion or fear of the audience towards this idea.

C.3 Distraction

- **Strawman:** consists in making an impression of refuting an argument of the opponent's proposition, whereas the real subject of the argument was not addressed or refuted, but instead replaced with a false one.
- **Red Herring:** consists in diverting the attention of the audience from the main topic being discussed, by introducing another topic, which is irrelevant.
- **Whataboutism:** a technique that attempts to discredit an opponent's position by charging them with hypocrisy without directly disproving their argument.

C.4 Simplification

- **Causal Oversimplification:** assuming a single cause or reason when there are actually multiple causes for an issue.
- **False Dilemma or No Choice:** a logical fallacy that presents only two options or sides when there are many options or sides. In extreme, the author tells the audience exactly what actions to take, eliminating any other possible choices.
- **Consequential Oversimplification:** is an assertion one is making of some "first" event/action leading to a domino-like chain of events that have some significant negative (positive) effects and consequences that appear to be ludicrous or unwarranted or with each step in the chain more and more improbable.

C.5 Call

- **Slogans:** a brief and striking phrase, often acting like emotional appeals, that may include labeling and stereotyping.
- **Conversation Killer:** words or phrases that discourage critical thought and meaningful discussion about a given topic.
- **Appeal to Time:** the argument is centred around the idea that time has come for a particular action.

C.6 Manipulative Wording

- **Loaded Language:** use of specific words and phrases with strong emotional implications (ei-

ther positive or negative) to influence and convince the audience that an argument is valid.

- **Obfuscation, Intentional Vagueness, Confusion:** use of words that are deliberately not clear, vague or ambiguous so that the audience may have its own interpretations.
- **Exaggeration or Minimisation:** consists of either representing something in an excessive manner or making something seem less important or smaller than it really is.
- **Repetition:** the speaker uses the same phrase repeatedly with the hopes that the repetition will lead to persuade the audience.

D LLMs Explored in Experiments

In our experiments, we used five different cutting-edge LLMs: *GPT-4o-mini*, *Meta-Llama 3.1* (8B-Instruct), *Claude 3 Haiku*, *Llama 3.3* (70B-Instruct-Turbo), and *Gemini 1.5 Flash*. We aimed to include widely recognized, state-of-the-art models from the largest available while ensuring they remain affordable. We also selected two open-weight models to demonstrate that our method can be applied without access to closed models through APIs. Additionally, we chose the smaller Llama 3.1 with 8B parameters to ensure that our method could be applied to models that do not require costly infrastructure.

Table 14 lists the Large Language Models used in our experiments, detailing their knowledge cutoff dates, access methods, licenses, and sizes. The knowledge cutoff dates confirm that our datasets, *MultiDis* and *EUDisinfo*, which contain articles from 2024 onward, were not part of the models' pretraining.

E Persuasion Detection Step Evaluation and Impact on PCoT

We conducted a series of extensive experiments to optimize the first stage of our PCoT method. Our experiments extensively relied on the taxonomy developed by Piskorski et al. (2023a,c), particularly for crafting prompts to detect persuasive strategies. The taxonomy was developed by researchers at the Joint Research Centre (JRC)⁸. This choice was driven by the fact that the taxonomy and the annotated data are publicly available. These datasets

⁸JRC is a scientific institution closely associated with the European Commission. The center provides independent, evidence-based insights and research to support the EU for societal benefit.

API Model Name	Knowledge Cutoff Date	Access Details	License	Model Size
gpt-4o-mini	October 2023	OpenAI API 02.2025	Commercial	Not Disclosed
gemin-1.5-flash	November 2023	Google API 02.2025	Commercial	Not Disclosed
claude-3-haiku-20240307	August 2023	Anthropic API 02.2025	Commercial	Not Disclosed
meta-llama/LLama-3.3-70B-Instruct-Turbo	December 2023	DeepInfra API 02.2025	Meta Llama 3 Community	70B
meta-llama/Meta-LLama-3.1-8B-Instruct	December 2023	DeepInfra API 02.2025	Meta Llama 3 Community	8B

Table 14: Large Language Models used in our experiments.

were used in the International Workshop on Semantic Evaluation, focusing on persuasion detection (Piskorski et al., 2023b; Dimitrov et al., 2024). By leveraging this resource, we could evaluate our taxonomy using ground truth data. Additionally, to our knowledge, this is the only high-quality taxonomy applied to longer-form news articles. News articles are a key component of our extensive experiments.

Below, we provide a description of the five key approaches we tested with shortcut in brackets:

1. **Multitask [MT]** - In this approach, we used a single prompt that included the names and definitions of persuasion strategies, as outlined by Piskorski et al. (2023c) and showed in Figure 2. This zero-shot prompting method guided the LLM in classifying persuasion strategies across multiple labels and categories. Furthermore, we instructed the model to provide explanations for each classification decision.
2. **Detailed Multitask [DMT]** - In this approach, we used a single, comprehensive prompt that provided additional context for each persuasion strategy to improve the performance of the classification tasks. Along with the definitions of the persuasion strategies, we included various techniques related to each strategy, with their definitions outlined by Piskorski et al. (2023c). Specifically, this method incorporated the names and definitions of strategies listed in Figure 2 and the names and definitions of techniques from Section C. Furthermore, we requested that the model explain each decision made in classifying persuasion strategies.
3. **One Task at a Time [TAT]** - In this approach, we used a separate prompt for each persuasion strategy, treating each as a binary classification task. This approach resulted in six distinct prompts, each focusing on a specific persuasion strategy from Piskorski et al. (2023c). Each prompt included only the name and definition of a single strategy, as listed in Figure 2. Additionally, we asked the model to explain each classification decision related to the corresponding persuasion strategy.
4. **One Detailed Task at a Time [DTAT]** - This approach is similar to the *One Task at a Time* method but with more detailed information to aid in the binary classification of each persuasion strategy. For each strategy, we used a separate prompt that not only included the name and definition of the strategy, as listed in Figure 2, but also provided the names and definitions of the persuasion techniques associated with that strategy, as outlined in Section C. As with the other approaches, we followed the taxonomy from Piskorski et al. (2023c) to structure the prompts.
5. **One Task at a Time with Broad Knowledge [TATB]** - This approach is similar to the *One Task at a Time* method but with a broader scope. Instead of providing knowledge about a single persuasion strategy per prompt, we used six distinct prompts, each containing knowledge about all the persuasion strategies. However, the LLM was still tasked with detecting and analyzing only one specific strategy within each prompt, treating it as a binary classification task.

Prompting Method	F ₁ Score
PCoT DMT	0.815 ↑ 1.6 p.p.
PCoT No Exp	0.799 ↑ 3.4 p.p.
PCoT Single Step	0.765 ↑ 5.4 p.p.
Base	0.711

Table 15: Impact of different prompting methods on final PCoT method performance

Table 17 presents the average results for each of the described methods. It includes the overall average performance in detecting persuasion and the results for each persuasion strategy. The *Detailed Multitask [DMT]* method outperformed the others in detecting persuasion. As a result, we selected *DMT* for the final version of the first stage of our PCoT method.

The results in Table 15 underscore the impact of different prompting strategies within the PCoT method for disinformation detection. The best-performing variant, *PCoT DMT*, achieves an F₁ score of 0.815, surpassing the baseline by 10.4

model	Explanation	F ₁ Score
GPT 4o mini	Yes	0.841
	No	0.830
Gemini 1.5 Flash	Yes	0.817
	No	0.798
Claude 3 Haiku	Yes	0.789
	No	0.771
Llama 3.3 70B	Yes	0.844
	No	0.842
Llama 3.1 8B	Yes	0.785
	No	0.756
Average	Yes	0.815
	No	0.799

Table 16: Results for PCoT with usage of explanation for each persuasion strategy and without explanation.

percentage points. Excluding explanations of the persuasion strategy in the first stage (*PCoT No Exp*) reduces the performance to 0.799, while adopting a single-stage approach (*PCoT Single Step*) further reduces it to 0.765. These findings emphasize the critical role of a two-stage reasoning process and persuasion strategy analysis in enhancing disinformation detection.

Table 16 presents a comparative analysis of PCoT with and without persuasion strategy explanations across various models. The impact of explanations varies, with the most significant improvement observed in the smallest open-weight model, Llama 3.1 8B, while Llama 3.3 70B shows minimal change. We observe a consistent average improvement when using explanations. Since inference is conducted with a temperature of 0, making the results more stable and reproducible, this further reinforces the importance of explanations. Notably, the benefits are most pronounced for smaller models, underscoring the value of explanations in enhancing their disinformation detection performance.

F Prompts

In this section, we provide an overview of the prompts used in our study and present prompt templates for each step of the PCoT method. Given the large number and substantial length of the prompts, we do not include them in full in the paper. Instead, the complete set of prompts is available in our online repository.

F.1 Baselines

Figure 10 illustrates the baseline prompt template used for zero-shot disinformation detection, specifically for the *VaN*, *Z-CoT*, and *DeF-SpeC* methods

introduced by Lucas et al. (2023). These methods were selected because Lucas et al. (2023) comprehensively evaluated various approaches using disinformation datasets, testing prompts on human-annotated and LLM-generated data. Since our study focuses exclusively on human-annotated data, we chose three of the best-performing methods on human-annotated data.

F.2 Persuasion Detection Step

Figure 11 presents the final template of the best-performing prompt used in the first stage of the PCoT method, designed specifically for detecting persuasion strategies and generating corresponding explanations. This prompt was meticulously crafted following a comprehensive evaluation of various approaches applied to data with ground truth labels for persuasion strategies. In particular, we tested multiple methods on the dataset from the International Workshop on Semantic Evaluation 2023 (SemEval 2023) shared task on persuasion (Piskorski et al., 2023b). The final prompt incorporates the names and definitions of persuasive strategies and the associated techniques outlined in Piskorski et al. (2023c,a). Figure 11 offers a detailed view of the prompt used in our study. Additionally, we make the final prompts publicly available.

F.3 Disinformation Detection Step

Figure 12 illustrates the final prompt template used in the second stage of the PCoT method, which focuses on disinformation detection. This prompt incorporates the persuasion analysis generated in the first stage of PCoT. For each test set, we experimented with three different disinformation prompts. We adjusted three methods *VaN*, *Z-CoT*, and *DeF-SpeC* (Lucas et al., 2023) to our PCoT method. This approach enabled us to compare the performance of the adapted methods against the baselines, where we applied the original methods from Lucas et al. (2023).

G Detailed Analysis

G.1 McNemar’s Test for PCoT Performance

To evaluate the statistical significance of PCoT, we conducted McNemar’s test comparing each prompting method to its PCoT-adjusted counterpart across various language models. The results, presented in Table 18, show that PCoT consistently improves performance at the 0.01 significance level

Approach	Metric	Attack on Reputation	Justification	Simplification	Distraction	Call	Manipulative Wording	Average
MT	F1 Micro	0.6407	0.6616	0.6198	0.7537	0.6366	0.7813	0.6823
		± 0.130	± 0.031	± 0.022	± 0.090	± 0.046	± 0.093	± 0.069
DMT	F1 Micro	0.7368	0.6710	0.6290	0.7082	0.6326	0.8440	0.7036
		± 0.103	± 0.044	± 0.028	± 0.104	± 0.046	± 0.058	± 0.081
TAT	F1 Micro	0.6522	0.6489	0.5940	0.5153	0.6541	0.7276	0.6320
		± 0.143	± 0.041	± 0.026	± 0.188	± 0.011	± 0.217	± 0.065
DTAT	F1 Micro	0.6963	0.6896	0.5985	0.4810	0.6407	0.8045	0.6518
		± 0.086	± 0.013	± 0.028	± 0.147	± 0.023	± 0.100	± 0.109
TATB	F1 Micro	0.6455	0.6437	0.5851	0.5269	0.6299	0.6858	0.6195
		± 0.133	± 0.045	± 0.047	± 0.248	± 0.058	± 0.225	± 0.056

Table 17: The table presents F₁ scores for each persuasion strategy and approach. Standard deviations, calculated from the results across five different LLMs, are provided below their corresponding scores. Detailed explanations of each approach can be found in Section E. The final approach used in the PCoT method is the best-performing *DMT*.

across all models and methods in overall evaluation. However, certain cases, such as experiments on posts for Llama 3.1 8B and experiments on articles for Claude 3 Haiku in DeF-Spec, exhibit non-significant differences.

G.2 Persuasion Strategy Correlations with Disinformation Detection

The results presented in Tables 19 and 20 provide key insights into the relationship between persuasion strategies and disinformation detection across different models and prompting methods. Table 19 presents the Matthews correlation coefficient (MCC) between various persuasion strategies and ground truth disinformation labels. The results reinforce previous findings, showing that across all models, *Attack on Reputation*, *Simplification*, *Distraction*, and *Manipulative Wording* exhibit positive correlations with disinformation, indicating that these strategies are strong signals of misleading content. In contrast, *Justification* and *Call* show in general a negligible correlation, suggesting that these strategies may be equally characteristic of credible and disinformation content.

Table 20 extends this analysis by evaluating the correlation between persuasion strategies and final disinformation predictions under different PCoT-adapted methods (*VaN*, *Z-CoT*, and *DeF-Spec*). The results demonstrate consistent patterns across all configurations, suggesting that PCoT’s effectiveness is not highly prompt-sensitive and remains stable across different prompting approaches.

It is important to note that we could not assess the impact of individual persuasive strategies in complete isolation, as all strategies were detected simultaneously. However, this analysis still provides valuable insight into which persuasive strategies are more characteristic of disinformation versus credible information.

G.3 PCoT Analysis on Predicted Persuasive and Non-persuasive Content

The results presented in Tables 21, 22, and 23 underscore the effectiveness of the proposed Persuasion-Augmented Chain-of-Thought approach in enhancing disinformation detection across various models and prompting methods. This improvement is evident in detecting disinformation in texts with predicted persuasive strategies (*Persuasion* subset) and those without (*No Persuasion* subset). PCoT consistently outperforms the baseline prompting methods (*VaN*, *Z-CoT*, *DeF-Spec*) in the *Persuasion* subset, where at least one persuasive strategy is identified. While PCoT also shows improvements in the *No Persuasion* subset, the gains are lower, highlighting the challenge of detecting misleading content without persuasive cues.

G.4 Persuasion Strategy Prediction in Disinformation and Reliable Content

In addition to Figure 3, we provide a heatmap in Figure 4 showing the distribution of predicted persuasion strategies within the final-stage predictions of the PCoT method. Figure 4 shows that the LLM-predicted distribution of persuasion strategies for predicted disinformation and reliable information closely matches the results in Figure 3.

Figures 5 to 9 show heatmaps depicting the distribution of persuasion strategies across all tested models. For each model, we illustrate the distribution of predicted persuasion strategies within ground truth disinformation and reliable information. In every case, disinformation exhibits a significantly higher percentage of texts containing at least one persuasive strategy. While credible content also employs persuasion, the distribution of specific strategies differs. *Justification* and *Call* are more characteristic of credible content, whereas other

strategies are more commonly associated with disinformation.

H Comparing BERT and LLMs on Unseen Data

Our experiments aim to validate the findings of (Lucas et al., 2023), which suggest that LLMs generalize more effectively and outperform BERT models in disinformation detection on unseen datasets. Furthermore, confirming these results strengthens the significance of our approach in advancing zero-shot classification.

H.1 Experimental Setup

Datasets We first selected three datasets: (i) CoAID, (ii) ISOT Fake News, and (iii) ECTF, to construct our training and validation sets. The validation set contained around 6,000 texts, while the training set included approximately 40,000. For testing, we used the same subsets as in the primary PCoT evaluation experiments, enabling a direct comparison with zero-shot classification results from baseline methods and our PCoT approach. Furthermore, no articles from EUDisinfo or MultiDis were included in the training or validation sets, ensuring they remained entirely unseen by BERT.

Model and Optimization We fine-tuned widely used pre-trained BERT model. The Hugging Face model name is as follows: google-bert/bert-large-uncased⁹. This model was also used by Lucas et al. (2023). For our computations, including hyperparameter optimization and final fine-tuning, we utilized an NVIDIA L40 GPU. Since these experiments were not the primary focus of our study, our hyperparameter exploration was limited in scope. However, we systematically varied two key hyperparameters: learning rate and weight decay. Specifically, we experimented with learning rates ranging from 5e-6 to 5e-5 and weight decay values between 0.005 and 0.03. The final selected values were a learning rate 1e-5 and a weight decay of 0.03. Other training hyperparameters were kept constant, including a batch size of 16 for both training and evaluation, three training epochs, and a warm-up phase covering approximately 8% of the total training steps.

⁹Hugging Face link to the BERT model and its details: [google-bert/bert-large-uncased](https://huggingface.co/google-bert/bert-large-uncased)

H.2 Results and Discussion

Tables 24 and 25 present the results of our experiments comparing the baseline method and the PCoT method across various models with result on BERT model. These tables present performance of each model in detecting disinformation on unseen data, so not available during pretraining and fine-tuning of any of models. BERT performs worse than all other models, with an F₁ score of 0.485.




Figure 4: Averaged percentage of persuasion strategies predicted across 5 models in predicted disinformation (*DIS*) and predicted reliable information (*REL*). *ALL* represents the percentage of instances with at least one detected persuasion strategy. Other abbreviations are explained in Figure 2.




Figure 5: Percentage of persuasion strategies predicted by GPT 4o mini for disinformation (*DIS*) and reliable information (*REL*). *ALL* represents the percentage of instances with at least one detected persuasion strategy. Other abbreviations are explained in Figure 2.




Figure 6: Percentage of persuasion strategies predicted by Gemini 1.5 Flash for disinformation (*DIS*) and reliable information (*REL*). *ALL* represents the percentage of instances with at least one detected persuasion strategy. Other abbreviations are explained in Figure 2.

Method	Data	Gemini 1.5 Flash	Claude 3 Haiku	GPT 4o mini	Llama 3.3 70B	Llama 3.1 8B
VaN	overall	0.01	0.01	0.01	0.01	0.01
VaN	articles	0.01	0.01	0.01	0.01	0.01
VaN	posts	0.01	0.01	0.01	0.01	Non-Significant
VaN	prior	0.01	0.01	0.01	0.01	0.01
VaN	post	0.01	0.01	0.01	0.01	0.01
Z-CoT	overall	0.01	0.01	0.01	0.01	0.01
Z-CoT	articles	0.01	0.01	0.01	0.01	0.01
Z-CoT	posts	0.01	0.01	0.01	0.01	Non-Significant
Z-CoT	prior	0.01	0.01	0.01	0.01	0.01
Z-CoT	post	0.01	0.01	0.01	0.01	0.01
DeF-Spec	overall	0.01	0.01	0.01	0.01	0.01
DeF-Spec	articles	0.01	Non-Significant	0.01	0.01	0.01
DeF-Spec	posts	0.01	0.01	0.01	0.01	0.01
DeF-Spec	prior	0.01	Non-Significant	0.01	0.01	0.01
DeF-Spec	post	0.01	0.01	0.01	0.01	0.05

Table 18: Results of McNemar’s test, comparing each prompting method (*VaN*, *Z-CoT*, and *DeF-Spec*) against its PCoT-adjusted counterpart across various language models. The values represent significance levels for different evaluation metrics, with *Non-Significant* indicating no statistically significant difference at the 0.05 threshold.

	persuasion	Attack on reputation	Justification	Simplification	Distraction	Call	Manipulative wording
GPT 4o mini	0.228	0.528	-0.160	0.611	0.230	0.008	0.507
Gemini 1.5 Flash	0.173	0.476	-0.219	0.511	0.203	-0.000	0.627
Claude 3 Haiku	0.220	0.378	-0.054	0.354	0.201	-0.029	0.628
Llama 3.3 70B	0.328	0.546	0.152	0.536	0.347	0.118	0.591
Llama 3.1 8B	0.178	0.484	-0.054	0.301	0.303	0.064	0.474

Table 19: The Matthews correlation coefficient between persuasion strategies and ground truth disinformation label. Table presents coefficients for each persuasion strategy. In addition, *persuasion* column shows correlation with predicted at least one persuasion strategy.




Figure 7: Percentage of persuasion strategies predicted by Claude 3 Haiku for disinformation (*DIS*) and reliable information (*REL*). *ALL* represents the percentage of instances with at least one detected persuasion strategy. Other abbreviations are explained in Figure 2.




Figure 9: Percentage of persuasion strategies predicted by Llama 3.1 8B for disinformation (*DIS*) and reliable information (*REL*). *ALL* represents the percentage of instances with at least one detected persuasion strategy. Other abbreviations are explained in Figure 2.



Figure 8: Percentage of persuasion strategies predicted by Llama 3.3 70B for disinformation (*DIS*) and reliable information (*REL*). *ALL* represents the percentage of instances with at least one detected persuasion strategy. Other abbreviations are explained in Figure 2.

	persuasion	Attack on reputation	Justification	Simplification	Distraction	Call	Manipulative wording
<i>VaN with PCoT</i>							
GPT 4o mini	0.307	0.601	-0.187	0.720	0.279	0.027	0.592
Gemini 1.5 Flash	0.228	0.495	-0.222	0.638	0.268	0.036	0.670
Claude 3 Haiku	0.353	0.491	-0.003	0.466	0.273	0.011	0.788
Llama 3.3 70B	0.422	0.648	0.176	0.622	0.385	0.158	0.700
Llama 3.1 8B	0.151	0.479	-0.155	0.362	0.333	0.027	0.481
<i>Z-CoT with PCoT</i>							
GPT 4o mini	0.308	0.597	-0.183	0.720	0.273	0.023	0.585
Gemini 1.5 Flash	0.227	0.495	-0.212	0.640	0.267	0.034	0.669
Claude 3 Haiku	0.334	0.504	0.018	0.419	0.257	0.012	0.766
Llama 3.3 70B	0.419	0.642	0.184	0.625	0.385	0.154	0.693
Llama 3.1 8B	0.166	0.484	-0.134	0.356	0.334	0.026	0.504
<i>DeF-SpeC with PCoT</i>							
GPT 4o mini	0.276	0.558	-0.203	0.720	0.277	0.025	0.557
Gemini 1.5 Flash	0.250	0.522	-0.225	0.638	0.260	0.032	0.709
Claude 3 Haiku	0.346	0.478	-0.003	0.443	0.255	0.010	0.782
Llama 3.3 70B	0.395	0.613	0.159	0.655	0.408	0.145	0.667
Llama 3.1 8B	0.163	0.455	-0.161	0.373	0.340	0.046	0.477

Table 20: The Matthews correlation coefficient between persuasion strategies and the final disinformation prediction. Table shows results for each base prompting method adopted to PCoT usage. Table presents coefficients for each persuasion strategy. In addition, *persuasion* column shows correlation with predicted at least one persuasion strategy.

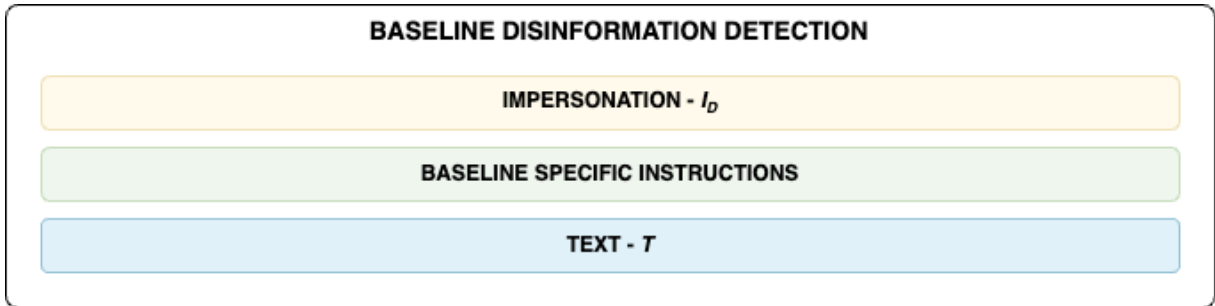


Figure 10: The prompt template for each baseline method in disinformation detection, namely, *VaN*, *Z-CoT*, and *DeF-SpeC*. The component I_D establishes context while overriding alignment tuning. Each baseline method differs in the *Baseline Specific Instructions* block. Generally, it provides method-specific guidelines defining the task and requests for structured output. Finally, the text T represents the content passed for disinformation evaluation.

Model	Persuasion		No Persuasion	
	PCoT	Base	PCoT	Base
GPT-4o-mini	0.876	0.815	0.315	0.303
Gemini 1.5 Flash	0.837	0.713	0.438	0.424
Claude 3 Haiku	0.840	0.787	0.128	0.304
Llama 3.3 70B	0.876	0.789	0.407	0.346
Llama 3.1 8B	0.816	0.631	0.551	0.561

Table 21: Performance comparison based on F_1 scores across two subsets: *Persuasion*, containing texts with at least one predicted persuasion strategy, and *No Persuasion*, containing texts with no predicted persuasion strategies. The table reports the F_1 score for *VaN* prompting method as *Base* and for our adaptation to *PCoT*.

Model	Persuasion		No Persuasion	
	PCoT	Base	PCoT	Base
GPT-4o-mini	0.876	0.827	0.348	0.297
Gemini 1.5 Flash	0.836	0.723	0.434	0.429
Claude 3 Haiku	0.815	0.644	0.196	0.206
Llama 3.3 70B	0.875	0.775	0.404	0.331
Llama 3.1 8B	0.818	0.676	0.535	0.473

Table 22: Performance comparison based on F_1 scores across two subsets: *Persuasion*, containing texts with at least one predicted persuasion strategy, and *No Persuasion*, containing texts with no predicted persuasion strategies. The table reports the F_1 score for *Z-CoT* prompting method as *Base* and for our adaptation to *PCoT*.

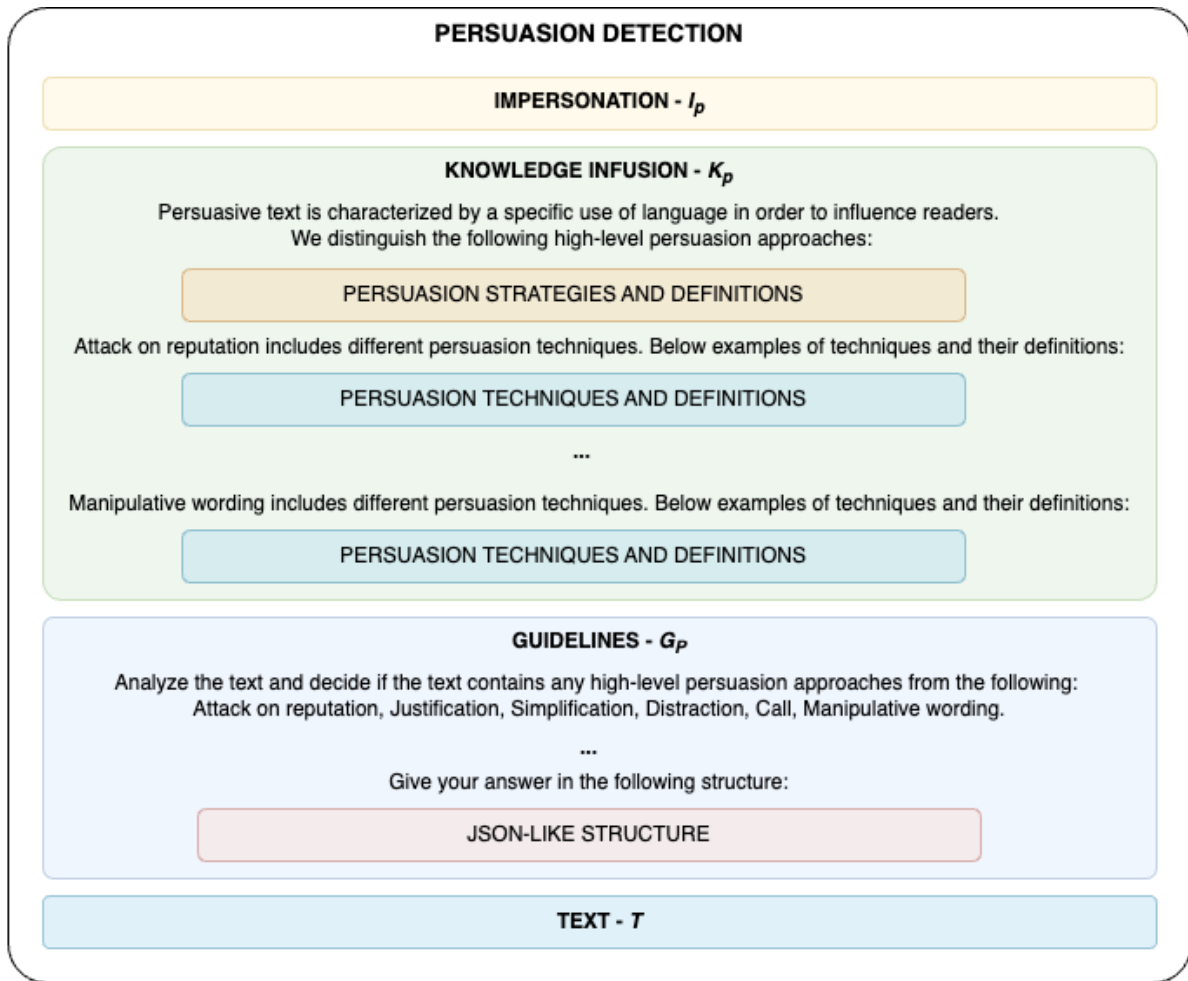


Figure 11: The prompt template for first stage of PCoT method, namely for persuasion detection step. The component I_P establishes the context and overrides alignment tuning, while K_P encapsulates knowledge about a predefined set of high-level persuasion strategies, and guidelines G_P determine the task and specify the structure of the expected response. The *Persuasion Strategies and Definitions* block includes names of persuasion strategies and definitions presented in Figure 2, while *Persuasion Techniques and Definitions* blocks includes names and definitions of techniques described in Appendix C. Finally, the text T represents the content passed for persuasion analysis.

Model	Persuasion		No Persuasion	
	PCoT	Base	PCoT	Base
GPT-4o-mini	0.865	0.829	0.364	0.315
Gemini 1.5 Flash	0.861	0.779	0.459	0.437
Claude 3 Haiku	0.837	0.838	0.208	0.374
Llama 3.3 70B	0.863	0.780	0.415	0.351
Llama 3.1 8B	0.803	0.730	0.523	0.448

Table 23: Performance comparison based on F_1 scores across two subsets: *Persuasion*, containing texts with at least one predicted persuasion strategy, and *No Persuasion*, containing texts with no predicted persuasion strategies. The table reports the F_1 score for *DeF-SpeC* prompting method as *Base* and for our adaptation to *PCoT*.

model	F_1 Score
BERT	0.485
GPT 4o mini	0.808
Gemini 1.5 Flash	0.719
Claude 3 Haiku	0.677
Llama 3.3 70B	0.752
Llama 3.1 8B	0.649

Table 24: Comparison between BERT performance and LLMs used with baseline methods (result averaged over 3 base methods) on post-cutoff datasets.

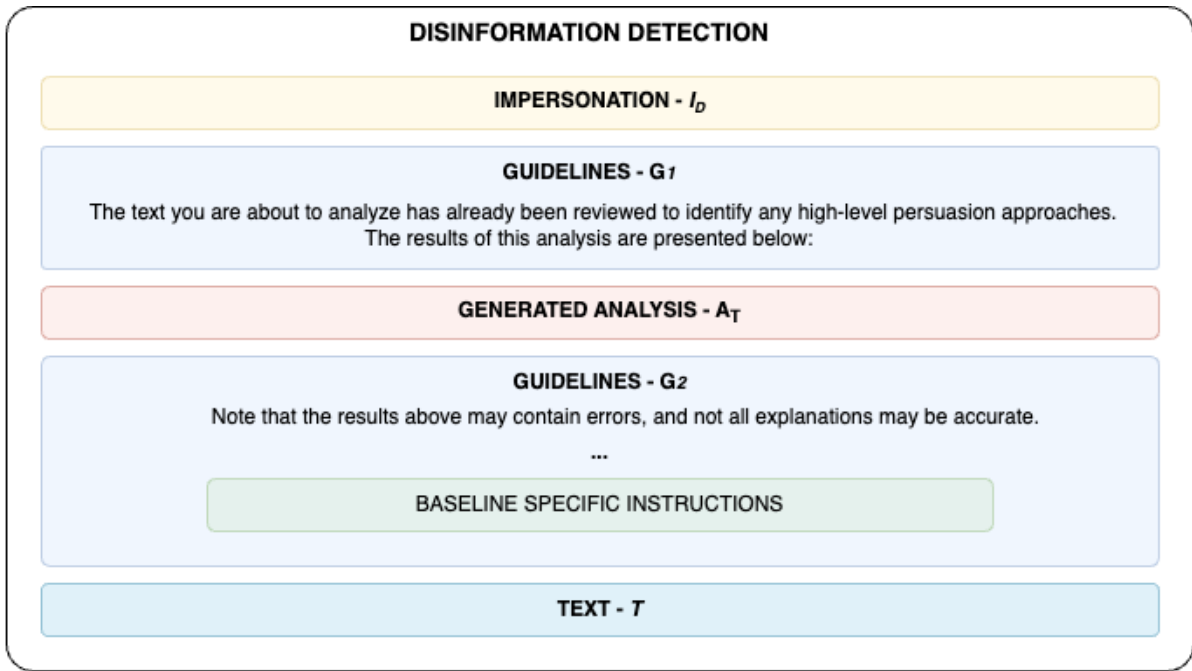


Figure 12: The prompt template for second final stage of PCoT method, namely for disinformation detection step. The component I_D establishes the context and overrides alignment tuning, while guidelines $G_D = \{G_1, G_2\}$ determine the task and specify the structure of the expected response. Next component is the generated analysis A_T from the output of first stage of PCoT and finally, the text T represents the content passed for disinformation evaluation. The *Baseline Specific Instructions* block is a part of guidelines and includes different instructions depending on which baseline method was adapted to PCoT method, namely it can be instruction from *VaN*, *Z-CoT*, or *DeF-SpeC*

model	F ₁ Score
BERT	0.485
GPT 4o mini	0.873
Gemini 1.5 Flash	0.877
Claude 3 Haiku	0.783
Llama 3.3 70B	0.864
Llama 3.1 8B	0.794

Table 25: Comparison between BERT performance and LLMs used with PCoT method (result averaged over 3 PCoT runs) on post-cutoff datasets.