

# Zero-shot and Few-shot Learning with Instruction-following LLMs for Claim Matching in Automated Fact-checking

Dina Pisarevskaya and Arkaitz Zubiaga

Queen Mary University of London, UK

{d.pisarevskaya, a.zubiaga}@qmul.ac.uk

## Abstract

The claim matching (CM) task can benefit an automated fact-checking pipeline by putting together claims that can be resolved with the same fact-check. In this work, we are the first to explore zero-shot and few-shot learning approaches to the task. We consider CM as a binary classification task and experiment with a set of instruction-following large language models (GPT-3.5-turbo, Gemini-1.5-flash, Mistral-7B-Instruct, and Llama-3-8B-Instruct), investigating prompt templates. We introduce a new CM dataset, ClaimMatch, which will be released upon acceptance. We put LLMs to the test in the CM task and find that it can be tackled by leveraging more mature yet similar tasks such as natural language inference or paraphrase detection. We also propose a pipeline for CM, which we evaluate on texts of different lengths.

## 1 Introduction

Claim matching (CM) is the task of determining if two claims can be verified using the same piece of evidence or fact-check, which can make the automated components of the fact-checking process more efficient (Zeng et al., 2021). Barrón-Cedeño et al. (2020) stated the problem, and Shaar et al. (2020) defined the task, introducing the first dataset and baselines. While some see CM as a ranking task that finds the claims most related to a given claim (Shaar et al., 2020, 2021b, 2022; Kazemi et al., 2021, 2022), others tackle it as a binary classification task determining if a pair of claims should be matched to each other or not (Kazemi et al., 2021, 2022; Choi and Ferrara, 2024a).

Large language models (LLMs) have become a powerful tool in Natural Language Processing (NLP) in recent years to solve a wide variety of tasks (OpenAI et al., 2024). The suitability of instruction-following LLMs to CM is yet to be studied. Finding out if LLMs can be used to determine

semantic similarity (SS) between statements has recently received a great deal of attention from the NLP community (Bubeck et al., 2023). However, two matching claims can have low SS. A complex input claim can contain multiple sub-claims (Shaar et al., 2020) or encompass reasoning between statements. LLMs have shown high results without task fine-tuning but based only on hand-crafted task descriptions with instructions, and zero or limited number of task examples (prompts) (Brown et al., 2020). Zero-shot and few-shot learning approaches can be particularly useful in the context of CM, given the expensive and time-consuming nature of manually labelling data where annotators need to go through long lists of claims to find matches. Hence, in our work, we explore the effectiveness of zero- and few-shot approaches to CM with instruction-following LLMs, where state-of-the-art methods (Kazemi et al., 2021, 2022) rely on larger training data for fine-tuning.

Given the newness of the CM task, little is known about how to approach the task: as a new task, or do approaches from close, more mature tasks suffice (e.g. natural language inference (NLI) or paraphrase detection (PD)). We investigate, for the first time, prompt-based, instruction-following LLMs for CM, in turn exploring the suitability of NLI, PD or bespoke CM prompts. We experiment on two datasets; (1) where there is a dearth of suitable datasets, we compile ClaimMatch, a new dataset extending that from CheckThat 2022 (Nakov et al., 2022), for claim matching with short texts, and (2) dataset based on Kazemi et al. (2022) for claim matching with long texts.

Claim matching is a specific and complicated fact-checking task, but, with contemporary instruction-following LLMs, it can be resolved as a natural language inference or paraphrase detection task. We find that LLMs yield superior results compared to the SS baselines, and state-of-the-art classification results with fine-tuned XLM-RoBERTa.

Providing few-shot examples helps LLMs better tackle the CM task. In zero-shot and few-shot setups, both PD and NLI templates are suitable for the CM task. Although different LLMs in the experiments require different prompt engineering approaches and differ if CM / PD / NLI setup is the best one for them, most of them are consistent in best prompts for CM, PD and NLI few-shot setups with a single instruction used. After running experiments with four models, we identify that Mistral with a PD template provides the best few-shot results (F1 95%), which with only 10 few-shot train examples gets close to the fine-tuned XLM-RoBERTa (F1 96.2%). With NLI and PD prompt templates combined, Gemini few-shot outperforms it (F1 97.2%), leading to the best score in the comparison.

Our main contributions are:

- we investigate PD, NLI and CM prompts to address our key research question to determine if CM is a unique task;
- we are the first to examine the CM task with instruction-following LLMs in zero-shot and few-shot learning scenario, comparing performance against baseline and state-of-the-art CM methods;
- we introduce and will release ClaimMatch, a new dataset for claim matching;
- we evaluate the CM pipeline on texts of different lengths.

## 2 Related Work

**Claim matching.** The task can be formulated for different settings: at the short / claim level or at the longer / document level. The task can be addressed as a ranking or as a classification task. Regarding the matching of short claims with other short texts, sentence-based BERT-like embeddings cosine similarity methods for estimating textual similarity of claims at the claim level are usually used. [Shaar et al. \(2020\)](#) proposed to handle CM as a ranking task, using methods based on BM25 ranking and BERT-based cosine similarity. Textual similarity methods were addressed further for the CM task, using pre-trained LMs (Universal Sentence Encoder, RoBERTa, sentence-BERT, etc.), for English (i.e. [Barrón-Cedeño et al., 2020](#); [Mihaylova et al., 2021](#); [Pritzkau, 2021](#)).

SS detection methods can be considered as the state-of-the-art baseline. [Shaar et al. \(2021a\)](#) investigated the task from a document-level perspec-

tive: given an input document, how to identify all sentences that contain a claim that can be verified by previously fact-checked claims. State-of-the-art similarity measures and ranking measures between possible input-verified pairs were implemented: BM25, NLI Score ([Nie et al., 2020](#)), BERTScore ([Zhang et al., 2020](#)), Sentence-BERT (SBERT) cosine similarity scores ([Reimers and Gurevych, 2019](#)), and SimCSE ([Gao et al., 2021b](#)). For non-English data, [Kazemi et al. \(2021\)](#) trained Indian XLM-R model, performing SS-based ranking, and classification experiments. To the best of our knowledge, we are the first to explore CM in zero-shot and few-shot settings, using pre-trained instruction-following LLMs.

[Kazemi et al. \(2022\)](#) studied how to find matches from longer fact-checks for short claims made in tweets. The task was addressed as a ranking and as a binary classification task. Multilingual XLM-RoBERTa pre-trained LM was fine-tuned on the dataset (with different settings: all data together or data in different languages separately). [Choi and Ferrara \(2024a,b\)](#) utilised domain fine-tuning (based on a synthetic training set) of generative pre-trained LLMs (GPT-3.5-Turbo, Llama-2-13b-chat-hf, and Llama-2-7b-chat-hf) for CM, considering it as a RTE (recognising textual entailment) classification task with 3 classes: entailment (the truth of the first claim implies the truth of the second claim), neutral, contradiction between two claims (truth of the first claim implies that second claim is false). This task setup differs from our task setup where CM is examined as a binary classification task to check if two claims can be checked with the same piece of evidence. In our setup, matches are not necessarily related to entailment (e.g. “the sky is blue” and “the sky is red” are definite matches in our setup, while one doesn’t entail the other). Following the recent efforts, we also tackle CM as a binary classification task in hitherto unexplored zero-shot and few-shot settings, studying it for shorter and for longer texts.

**Zero-shot and few-shot learning** Let models perform on new tasks without fine-tuning. Zero-shot and few-shot setups assume addition of extra prompt text to each input example, so a model can learn the task based on such train examples and instructions ([Schick and Schütze, 2021b](#); [Le Scao and Rush, 2021](#)) and provide the answer that corresponds to a final label, filling the unfilled slots in a prompt. This method depends on manually

or automatically creating better label words and templates for prompts (prompt engineering). As some prompts work better than others, methods to automatically generate high-quality and diverse prompts have been suggested (Jiang et al., 2020; Perez et al., 2021; Lu et al., 2022; Lester et al., 2021; Qin and Eisner, 2021; Li and Liang, 2021). Schick and Schütze (2021a) studied how to detect appropriate label words automatically, but the hand-picked ones performed better, while Gao et al. (2021a) applied prompt generation with automated selection of both templates and label words, that achieved comparable results to hand-crafted prompts. Webson and Pavlick (2022) evaluated different manually written prompt templates for NLI task. Logan IV et al. (2022) showed that tuning only the prompt yields worse performance than with manually designed prompts, but suggested that it might depend on the models and settings. Promptsource (Bach et al., 2022) provides a variety of prompts for zero-shot and few-shot learning, including prompts for PD and NLI.

**LLMs applications.** Recent advances with LLMs (Dang et al., 2022) have led to their exploration across various tasks. Prompt-based approaches can be used for classification tasks in zero-shot and few-shot setups: NLI and SS as a classification task (Gu et al., 2022), topic classification (Hu et al., 2022), and paragraph classification (Deng et al., 2022; Wang et al., 2022). Wang et al. (2021) reformulated NLP tasks as the entailment ones, and then fine-tuned the model with as little as 8 examples. Reasoning abilities of recent LLMs were yielded in (Wei et al., 2022; Zhou et al., 2023; Press et al., 2023; Wei Jie et al., 2024), encouraging their use in tasks that require common-sense reasoning and complex language understanding, such as multiple fact-checking tasks: zero-shot scientific claims verification with GPT-3.5 and GPT-4 (Alvarez et al., 2024), generating fact-checking explanations with GPT-4, Claude 2, and PaLM 2 (Hsu et al., 2024), creating claim-focused summaries in zero-shot and few-shot setups with GPT-3.5 for claim verification (Chen et al., 2024), or exploring few-shot scenarios with GPT-3.5 to verify claims with a knowledge source (Li et al., 2024). In this work, we study the CM task, using instruction-following, prompt-based LLMs for it.

### 3 Datasets

We use two datasets in our research: ClaimMatch, which we create by extending and adapting data from Nakov et al. (2022) to experiment with short texts; and longer texts (LT) dataset based on Kazemi et al. (2022) to experiment with long texts.

**Short texts.** To create the ClaimMatch dataset, we rely on texts from the fully available public dataset (Nakov et al., 2022), which was in turn created by building on and extending Shaar et al. (2020). We use the multi-domain English subset of the data and get rid of the remainder data pertaining to the political domain only. It consists of tweets as input claims and previously verified claims from the corresponding articles from the Snopes website. It contains 1,398 claim pairs (training, development and dev-test sets combined), and 14,245 previously checked claims in general.

We extend this data as it originally contained only positive cases of claim matches, with no negative cases. For positive class examples, we chose 500 claim pairs from the dataset. For the negative class examples, we created them: for each of the selected input claims, we took another verified claim not from their pair (each verified claim from the dataset could be used only once). We used such new pairs as negative examples.<sup>1</sup> Table 4 in the appendix shows samples of positive and negative claim matches.

All texts were preprocessed: urls, retweet markers (“RT”) and emojis<sup>2</sup> were removed. Username (@) and hashtag (#) markers were removed while leaving the mentions, to keep the content of short texts. Text of a verified claim includes its title, subtitle and main text, taking into account the importance of the titles of articles about claims in previous research (Mansour et al., 2022).

Our resulting test set, used in the LLMs experiments, contains 1,000 examples: 500 for the positive class and 500 for the negative class. Both input claims and verified claims are short. For the input claims, the average length after preprocessing is 194 characters (39 tokens). For verified claims, average length after preprocessing is 303 characters (56 tokens). This new dataset is specified for CM as a binary classification task, and will be released.

**Longer texts.** We use the Kazemi

<sup>1</sup>We manually double-checked all generated pairs to prevent false positives or false negatives.

<sup>2</sup>Emojis were processed with the package <https://pypi.org/project/emoji/>

dataset (Kazemi et al., 2022)) for domain transfer evaluation on longer texts. It is based on fact-checks obtained from several sources, and contains 5,028 pairs of long texts & short claims, of which we firstly sample 1,000 due to computational restrictions. The preprocessing of texts from this dataset is the same as for the first one. In addition, we remove pairs with Levenshtein distance similarity ratio between a text and a claim more than 80%, to remove (near-)duplicates. The average lengths after preprocessing are: 209 characters (41 token) for claims, 3,776 characters (729 tokens) for articles. To create the LT dataset, we take a sample of 129 positive examples and create 129 negative examples with the same procedure as described for short texts above (258 test examples in general). The dataset (Kazemi et al., 2022) is public, but its usage required collecting twitter data to build connections between texts. As the dataset was not publicly available online, we assume that it was not used for pre-training LLMs, which is important while using public datasets for evaluation with LLMs. Hence, the LT dataset is suitable for the evaluation of the proposed CM approach with LLMs.

## 4 Experimental Setup

We use four popular instruction-following LLMs for prompt-based experiments, including 2 models with paid API, and 2 freely available models (accessed by request):

1. Gpt-3.5-turbo-0125 (**GPT-3.5**): the latest and cost effective model in the GPT-3.5 family allows instruction following and works for traditional completions non-chat tasks as well<sup>3</sup>. Tokens context length was limited by the model context length.
2. Gemini-1.5-flash (**Gemini**): the fastest and cheapest model from the Gemini family<sup>4</sup> allows instruction following. Default context length was used.
3. Mistral-7B-Instruct-v0.3 (**Mistral**): the latest instruct fine-tuned version of Mistral-7B-v0.3<sup>5</sup>, the popular example of open-source LLMs of such size. Maximal new tokens is

<sup>3</sup><https://platform.openai.com/docs/models/gpt-3-5>

<sup>4</sup><https://blog.google/technology/ai/google-gemini-update-flash-ai-assistant-io-2024/>

<sup>5</sup><https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3>

set to 400, as it places the full model’s answers.

4. Llama-3-8B-Instruct (**Llama**): the last open-access Llama model of such size<sup>6</sup>. The hyper-parameters were: temperature 0.6, top\_p 0.9, maximum of new tokens 400.

The latter two models were set up to fit the Colab Pro+ resources (A100 GPU with 40 GB RAM). For all four models, their zero-shot and few-shot setups were examined with default system instruction & specified user instruction (“single instruction”), or with specified system instruction & specified user instruction (“ensemble instruction”).

We list the selected CM, PD and NLI prompt template sets for zero-shot and few-shot learning in Table 1. CM templates were selected after empirical tests from a broader initial list of CM templates. PD and NLI template sets are based on Prompt-Source base (Bach et al., 2022). For Mistral and Llama, questions were moved to the beginning of the template to adhere to model requirements.

In zero-shot settings, only templates were given to models. For few-shot, we used the priming method (in-context learning): together with a template, labeled examples are included in an input sequence, so a model makes a prediction (Brown et al., 2020). We chose a sample of 5 positive examples from the data. As negative examples, we took 5 random input claims and selected for them 5 verified claims not from their pairs. All these claims were not used in the test set. So the set of train examples was balanced. Based on a prompt with these selected examples, the models were able to produce binary classification labels for the new examples. We used “yes”/“no” (templates excepting NLI 1-2) or “true”/“false” (NLI 1-2 templates) as label words for the task, for two corresponding classes: “match” (positive)/“not match” (negative). The final order of positive and negative examples was mixed.

In addition to the instruction-following LLMs, we include in the evaluation:

1. State-of-the-art model (SOTA): multilingual XLM-RoBERTa-base<sup>7</sup> (**XLM-R**) model that, after fine-tuning on data, showed the best results for CM as classification task in (Kazemi et al., 2022). In our study, it was fine-tuned for the CM task on the dataset of remaining

<sup>6</sup><https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>

<sup>7</sup>[https://huggingface.co/docs/transformers/en/model\\_doc/xlm-roberta](https://huggingface.co/docs/transformers/en/model_doc/xlm-roberta)

CM templates	
<b>CM-1</b>	A Matches to B. Correct? Answer: [yes/no]
<b>CM-2</b>	A Means that B. Correct? Answer: [yes/no]
PD templates	
<b>PD-1</b>	A. B. Question: Do A and B express the same meaning? Yes or no? Answer: [yes/no]
<b>PD-2</b>	A. B. Question: Do A and B express the same meaning? Answer: [yes/no]
<b>PD-3</b>	A. B. Question: Do A and B have similar meanings? Yes or no? Answer: [yes/no]
<b>PD-4</b>	A. B. Question: Are A and B saying the same thing? Yes or no? Answer: [yes/no]
<b>PD-5</b>	A. B. Question: Are A and B essentially the same? Yes or no? Answer: [yes/no]
<b>PD-6</b>	A. B. Question: Do A and B both refer to the same event? Yes or no? Answer: [yes/no]
NLI templates	
<b>NLI-1</b>	Suppose it's true that A. Then, is B. Question: Is true or false? Answer: [true/false]
<b>NLI-2</b>	Take the following as truth: A. Then B is true or false? Answer: [true/false]
<b>NLI-3</b>	A. Based on the previous statement, is it true that B? Yes or no? Answer: [yes/no]
<b>NLI-4</b>	Given A Is it guaranteed true that B? Yes or no? Answer: [yes/no]
<b>NLI-5</b>	Suppose A. Can we infer that B? Yes or no? Answer: [yes/no]

Table 1: Templates used for claim matching (CM), paraphrase detection (PD) and natural language inference (NLI).

1,790 claim pairs, not included in test set (10 epochs, learning rate 1e-5).

- Baselines: (1) SS-based model using All-MiniLM-L6-v2<sup>8</sup> (**All-MiniLM**) and (2) SS-based model using text-embedding-3-small<sup>9</sup> (**embedding3**).

As SS-based methods need a threshold to determine the class, a separate validation set of 1,000 claim pairs was used to calculate this threshold. Median SS score for the positive class was chosen as a threshold (0.64 for All-MiniLM and 0.63 for embedding3). In the test set, all claim pairs with SS  $\geq$  this threshold got positive class labels.

We used the standard evaluation metrics for classification: F1 score (weighted) and accuracy (in Error analysis, precision (weighted) and recall (weighted) are added). For all models, less than 5% examples did not obtain a clear classification label (“partial match” in a longer model output), so they were included in the negative class.

## 5 Zero-shot and few-shot experiments

To identify the most effective user instructions for a prompt, we investigate different templates on ClaimMatch and compare results in Table 2.

### Few-shot results with a “single instruction”.

Overall, our results show that PD and NLI templates outperform CM templates, suggesting that reformulating the CM task as a PD and NLI helps. The improvement is primarily noticeable with PD, suggesting that, with the datasets at hand, matching

claims can be handled as paraphrases. This is however not consistent; certain templates such as PD-6 and NLI-5 can outperform CM templates, but not all PD and NLI templates do.

Different LLMs require different prompt engineering approaches and differ if CM / PD / NLI setup is the best one for them, but most of them are consistent in best prompts within each of the CM, PD and NLI setups: for GPT-3.5, Gemini and Mistral, templates CM-1, PD-6 and NLI-5 are the most effective for their sets. Only for Llama, templates CM-2, PD-4 and NLI-4 are better. For GPT-3.5 and Mistral, PD-6 template performs better, for Gemini - NLI-5, for Llama - NLI-4.

Prompt templates in the same set can have similar meanings yet provide distinct results: for example, PD-1 and PD-2 differ slightly, but all the model results yield significant differences. The results of the models also vary in different ways in the template sets and among the models and the templates. Wording-dependent prompt engineering separately for each model is required here. Mistral yields the highest overall metrics in all template sets. For this model, PD templates show highest scores, and the CM templates work better than the NLI templates. Llama has the lowest scores, and its results with the NLI-2 template demonstrate that the model understood the task wrongly (as a claim verification rather than as a CM task). GPT-3.5 and Gemini are not so consistent along results for template sets. Although Llama is very good with NLI-4, it performs much worse with NLI-5 where explicit inference understanding is needed. Mistral shows high consistency and impressive scores along PD templates results, but for NLI they vary much. These examples demonstrate that there are

<sup>8</sup>sentence similarity model with the highest number of downloads on huggingface: <https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

<sup>9</sup><https://platform.openai.com/docs/guides/embeddings/embedding-models>

FEW-SHOT						ZERO-SHOT		
Model	F1, %	Acc., %	Model	F1, %	Acc., %	Model	F1, %	Acc., %
<b>PD templates</b>			<b>NLI templates</b>					
<b>PD-1</b>			<b>NLI-1</b>			<b>PD-6</b>		
GPT-3.5	80.3	81.0	GPT-3.5	70.6	70.9	GPT-3.5	<b>89.2</b>	<b>89.3</b>
Gemini	80.5	81.2	Gemini	66.2	69.3	Gemini	<b>90.7</b>	<b>90.7</b>
Mistral	93.3	93.3	Mistral	85.5	85.7	Mistral	85.6	85.9
Llama	52.4	59.7	Llama	54.9	58.1	Llama	<b>88.0</b>	<b>88.1</b>
<b>PD-2</b>			<b>NLI-2</b>			<b>NLI-5</b>		
GPT-3.5	74.2	75.7	GPT-3.5	67.5	67.8	GPT-3.5	70.6	72.7
Gemini	77.2	78.3	Gemini	61.4	65.3	Gemini	71.0	73.0
Mistral	92.7	92.7	Mistral	84.1	84.2	Mistral	79.5	80.3
Llama	64.9	67.2	Llama	36.2	50.2	Llama	74.3	75.7
<b>PD-3</b>			<b>NLI-3</b>			<b>CM-1</b>		
GPT-3.5	74.2	75.7	GPT-3.5	69.8	70.8	GPT-3.5	69.9	71.7
Gemini	85.3	85.6	Gemini	74.9	76.1	Gemini	62.8	63.3
Mistral	93.7	93.7	Mistral	78.2	78.4	Mistral	<b>88.9</b>	<b>89.0</b>
Llama	72.9	72.9	Llama	86.4	86.5	Llama	70.9	72.9
<b>PD-4</b>			<b>NLI-4</b>					
GPT-3.5	78.4	79.3	GPT-3.5	63.8	64.2			
Gemini	80.2	80.9	Gemini	77.9	78.7			
Mistral	93.5	93.5	Mistral	79.1	79.2			
Llama	<b>78.4</b>	<b>78.4</b>	Llama	<b>91.8</b>	<b>91.8</b>			
<b>PD-5</b>			<b>NLI-5</b>			<b>MANY-SHOT</b>		
GPT-3.5	78.5	79.4	GPT-3.5	<b>75.7</b>	<b>76.3</b>	SOTA model		
Gemini	85.5	85.8	Gemini	<b>93.4</b>	<b>93.4</b>	XLM-R	<b>96.2</b>	<b>96.2</b>
Mistral	94.2	94.2	Mistral	<b>88.3</b>	<b>88.3</b>	Baselines		
Llama	66.2	67.7	Llama	52.8	59.7	All-MiniLM	76.6	77.8
<b>PD-6</b>			<b>CM templates</b>			embedding3		
GPT-3.5	<b>84.8</b>	<b>85.1</b>	<b>CM-1</b>			<b>77.2</b>		
Gemini	<b>90.3</b>	<b>90.4</b>	GPT-3.5	<b>75.8</b>	<b>76.9</b>			
Mistral	<b>95.0</b>	<b>95.0</b>	Gemini	<b>72.2</b>	<b>73.8</b>			
Llama	60.0	64.5	Mistral	<b>90.6</b>	<b>90.6</b>			
			Llama fs	77.6	78.3			
			<b>CM-2</b>					
			GPT-3.5	71.1	72.7			
			Gemini	67.2	69.8			
			Mistral	89.9	89.9			
			Llama	<b>81.7</b>	<b>81.8</b>			

Table 2: Few-shot (left) and zero-shot (top right) performance with PD, NLI and CM templates. Many-shot SOTA and baselines (bottom right).

some variation across LLMs, and one should be careful choosing the template type depending on the model.

As for SS baselines, Mistral and Llama outperform All-MiniLM and embedding3 in both CM templates, although GPT-3.5 and Gemini results are below the baselines. With most of the PD templates, GPT-3.5, Gemini and Mistral yield higher results than both baselines, for Llama only one template works better than them. GPT-3.5 does not outperform the baselines with all NLI templates, but two NLI templates work for Gemini and Llama (with NLI-4 for both of them) better than the baselines. All Mistral NLI templates outperform the baselines. The best few-shot results of PD-6 template with Mistral (F1 95.0%) and NLI-5 template with Gemini (F1 93.4%) do not outperform state-of-the-art classification results with fine-tuned XLM-

R (F1 96.2%), but can be compatible with them in real-word scenarios, where only little training data are available, not sufficient for models fine-tuning.

**Zero-shot results with a “single instruction”.** For comparison, we present zero-shot results for three templates, that showed the best results for most of the models: CM-1, PD-6, and NLI-5. As expected, for most of the settings, few-shot provides higher scores than zero-shot, and we need some examples to let LLMs better tackle the task. However, With PD-6 template, for GPT-3.5 zero-shot is better, and for Gemini - slightly better than few-shot. GPT-3.5 provides F1 89.2%, that is still below the XLM-R F1 96.2%. For Llama, zero-shot provides significantly higher scores with PD-6 and NLI-5 templates. It can be explained that this model gets distracted by the content of few-shot examples specifically in the given templates. Zero-

FEW-SHOT			ZERO-SHOT		
Model	F1, %	Acc., %	Model	F1, %	Acc., %
<b>CM-1 &amp; PD-6 template</b>			<b>CM-1 &amp; PD-6 template</b>		
GPT-3.5 fs	40.8	69.0	GPT-3.5 zs	95.7	95.7
Gemini fs	91.1	91.2	Gemini zs	<b>97.1</b>	<b>97.1</b>
Mistral fs	92.3	92.3	Mistral zs	<b>96.2</b>	<b>96.2</b>
Llama fs	66.0	69.2	Llama zs	91.9	91.9
<b>PD-6 &amp; PD-6 template</b>			<b>PD-6 &amp; PD-6 template</b>		
GPT-3.5 fs	87.8	87.9	GPT-3.5 zs	95.6	95.6
Gemini fs	95.1	95.1	Gemini zs	96.0	96.0
Mistral fs	94.9	94.9	Mistral zs	95.5	95.5
Llama fs	60.7	65.2	Llama zs	95.6	95.6
<b>NLI-5 &amp; PD-6 template</b>			<b>NLI-5 &amp; PD-6 template</b>		
GPT-3.5 fs	59.5	59.9	GPT-3.5 zs	<b>96.1</b>	<b>96.1</b>
Gemini fs	<b>97.2</b>	<b>97.2</b>	Gemini zs	97.1	97.1
Mistral fs	89.0	89.0	Mistral zs	92.5	92.5
Llama fs	48.0	57.4	Llama zs	94.7	94.7

Table 3: Few-shot (left) and zero-shot (right) performance with ensemble instructions.

shot results show the potential of tackling the CM task with PD templates, and it also shows that using just a few shots can lead to a significant gain generally, so having a few samples to train the model can improve performance.

**“Ensemble” instructions.** As almost all models have high performance with PD-6 template, we added it specifically to their system instructions, so each model setting has 2 instructions combined: PD-6 template in system instruction and another template in user instruction. Results are presented in Table 3 for CM-1, PD-6 and NLI-5 templates. In all settings, excluding Gemini NLI-5, zero-shot worked much better than few-shot. Only for Llama, “double” PD-6 template with zero-shot is better than CM-1 and NLI-5 templates. “Ensemble” instructions let Gemini improve results - both for zero-shot and few-shot setups. Compared to best few-shot results with default system instructions, zero-shot results for all models are significantly better, and more similar along all models and templates. It is unexpected, because with a “single” instruction few-shot was in general better than zero-shot. It can be explained that models can be distracted by few-shot examples provided in user instruction only for one of two templates. In zero-shot, Mistral with CM-1 & PD-6 template reaches state-of-the-art scores (F1 96.2%), and Gemini outperforms them (F1 97.1%). While CM-1 & PD-6 template is the best zero-shot one for Mistral and NLI-5 & PD-6 template is the best zero-shot one for GPT-3.5, Gemini, however, provides the best overall scores (F1 97.2%) in the comparison with its few-shot NLI-5 & PD-6 template (outperforming the state-of-the-art scores with both zero-shot and few-shot NLI-5). So few-shot training data, and templates

from close tasks, such as NLI and PD, are still helpful to achieve better results. Combination of templates leads to improved and more consistent results, which make it more reliable. While PD templates were very good alone, it shows that they can be further improved by using other templates, including a CM template. For most of the cases, combinations of templates yields better scores than these two templates separately, both for few-shot and zero-shot.

Both PD and NLI templates, as templates from close NLP tasks, seem to be appropriate for the CM task in zero-shot and few-shot setting with instruction-following LLMs. But they should be carefully examined specifically for each model, as they are more or less suitable for different models, and even similar templates let get different classification scores. Combinations of templates and choice of a zero-shot/few-shot setup should be carefully investigated, to yield better results.

**Longer texts.** We next created a pipeline for few-shot experiments for different domains and tested it on longer texts devoted to various news topics from LT dataset. Two setups were used in these experiments: Mistral few-shot PD-6 (as the best one with “single instructions”) and Gemini few-shot NLI-5 & PD-6 (as the best one with “ensemble” instructions). For both models, default context length was used. Despite Kazemi et al. (2022), longer texts were processed on the entire text level and not on paragraph level (due to LLMs context capabilities). We experimented with two options: 1. using train examples, the same as in experiments on short texts; 2. using a sample of new 10 domain train examples from this dataset (Kazemi et al., 2022). Mistral setup resulted in 83.7% (both F1 and accu-

racy scores) with short train examples, but domain train examples led to only F1 67.7% and accuracy 70.5%. Gemini setup yielded 98.4% (both F1 and accuracy scores) with short examples, and 95.7% (both F1 and accuracy scores) with domain examples. As the performance for both Mistral and Gemini models was better with domain-independent train examples (short texts) than with domain train examples, we suggest using data-independent train examples in the pipeline. Hence, in the proposed pipeline, firstly data-independent train examples should be chosen (manually or in automated way). The prompt template should be selected. Then, for each new test example, few-shot learning technique should be applied: 10 train examples should be given to a LLM, followed by the test example. Such domain transfer methods would be helpful for benchmarking and evaluation of the few-shot learning with instruction-based LLMs for CM task, on different domains and on texts of various lengths.

## 6 Error Analysis

We studied all model outputs, to deeper understand performance and why different models perform better with different prompts. For all templates except for NLI 3-5, in most of the cases the answers were accompanied with model explanations. Gemini few-shots provided more structured answers (starting with the class label, e.g. “yes”/“no”), with less explanations than other models.

**Fact-checking instead of CM.** “Single instruction” CM templates work for Mistral much better than for other models, and it understood the task very well (due to the model’s argumentation). But NLI 1-2 templates were problematic: the model did not understand the task and incorrectly addressed it as a fact-checking task, due to “true”/“false” in instruction question. For other models, such cases were also present for these tasks, so “true”/“false” questions should be excluded for CM. NLI 3-4 templates did not cause as many fact-checking answers: they contain “true”, but focus on two claims inference is expressed more clearly.

**Focus on train examples** can also lead to errors. For Llama, “single instruction” templates performed worse than for other models (excluding NLI-4). In few-shot, it could be confused with train examples and start comparing two test claims not to each other, but to train claims (e.g. in NLI-5 few-shot recall for negative class was 21.0%). But its zero-shot with “ensemble” instructions got results

compatible to other models. This model requires a specific prompt engineering (e.g. more explicit numeration in train and test examples).

**PD and NLI templates require some clarification for CM task.** Even PD-6, one of the best templates, can still be too strict for CM. In false negative answers of all models, two claims were considered by models as referring to different events, if they vary in some details. But such claims can still match each other according to the gold label.

This error type is more specific for zero-shot shots. All four models can detect if two statements in a claim pair vary in some non-substantial or significant details. But the level of this difference, and the final class label verdict of a model are handled individually for each concrete example, model and template setup. For example, Mistral’s output with CM-1 template included 36 gold label positive class examples where the model highlighted that claims in a pair are about similar, but not the same events (e.g. “While Statement 1 does suggest that Joe Biden owns the largest mansion in his state, it does not definitively establish this fact. Therefore, it does not match Statement 2, which claims that the information is definitively known”). The model provided the positive class label only for 11 examples. But if we consider all 36 examples as positives, the final classification performance is improved: both F1 score and accuracy are 91.5% (instead of 88.9% and 89.0% respectively). Hence, templates can be modified for the task. Separate model explanations should be a step of the pipeline. If a model answers that two claims are about the same topic and event, but vary in not substantial details, the pair should be processed as positive during post-processing.

**Claim specifics** can also cause errors. The models in different setups classified one example as false negative: two claims mentioned different numbers. E.g. Llama’s output with PD-6 template, after providing the negative class label: “Statement 1 mentions that cameras on the M1 and M25 go live, and anyone going over 70 mph gets an instant ticket. This suggests that the cameras were activated with a speed limit of 70 mph. Statement 2 mentions that all speed cameras on the M1 and M25 were activated in January 2019 with a uniform speed limit of 72 miles per hour. The speed limit is different in the two statements, which indicates that they do not refer to the same event”. The discrepancies, highlighted by the model, really take place in the texts, so it can be explained why two



claims can be considered as not fully matching and not verifiable by the same piece of evidence.

**Complex claims and nested sub-claims.** Addressing this challenging CM issue, all models demonstrated understanding of sub-claims. However, with PD templates 1-5 models can answer that statements A and B are about similar, but not the same event, as B includes another sub-claim - information that is not present in A; or the same persons and locations are mentioned, but the context is different. This issue can be solved by exploring prompt templates and decomposing the CM task.

We conducted a fine-grained study of 10 claim pairs, where at least one of the texts is a complex claim with several sub-claims in it, e.g.: “1) In the 1640’s the Dutch inhabitants of New Amsterdam built a 12’ wall to keep the bad hombres out. In 1664 the British ignored the wall and took New Amsterdam by sea. It’s now called New York. They took down the wall and built a street. It’s called Wall Street. — Joe Delmonaco (JoeDelmonaco) January 7, 2019 2) Was Wall Street Originally the Site of a “Border Wall” Meant to Protect New Amsterdam? Social media memes compared a defensive wall built along the northern border of New Amsterdam during the 17th century to President Donald Trump’s border wall. In the 17th century, New Amsterdam built a protective wall along its northern perimeter (analogous to Trump’s border wall) to keep “bad hombres” out, but it failed to achieve its stated purpose in that the British successfully invaded the city by sea.”

We did not find a consistently best PD template, but found some patterns. With PD-1 and PD-2 templates, Llama and GPT-3.5 have difficulties with the provided example. Llama with PD-2 template correctly noticed that there are several sub-claims in Statement 2. It justified the negative class label for this claim pair: “No, Statement 1 and Statement 2 do not express the same meaning. Statement 1 discusses the construction of a wall in New Amsterdam in the 1640s and its failure to prevent the British from taking the city by sea. Statement 2 compares the wall built by New Amsterdam to President Trump’s border wall and states that it failed to achieve its purpose, citing the British invasion of the city.” In general, CM for complex claims is well solved by the models. Compared to few-shot results with a “single instruction”, all four models with “ensemble instructions” perform better. It shows that different instructions help a model understand the task from different perspectives.

## 7 Discussion & Conclusion

We propose a novel direction to tackle claim matching as a classification task: prompt-based zero-shot and few-shot learning for pre-trained instruction-following LLMs, suitable when there is no or limited train data. We create the ClaimMatch dataset with short texts and perform extensive experiments on it. We use manually created prompt templates for zero-shot and few-shot (10 train examples) settings, aimed specifically for CM or for two similar tasks: PD and NLI. We investigate if the latter two tasks can help tackle CM task. Different LLMs from our benchmark comparison require different prompt engineering approaches, but in general PD and NLI work better than CM templates.

Our results help better understand that zero-shot and few-shot methods with templates crafted for PD and NLI tasks are appropriate for CM. With instruction-following LLMs, it can be resolved on the existing data with using NLI and PD tasks approaches. Few-shot training examples are still needed to experiment with the better settings. Hence, advanced prompt engineering methods should be further carefully developed for CM. We suggest the pipeline for zero-shot and few-shot CM, and test it with best prompt templates on longer texts. Our work makes the first contribution of prompt-based zero-shot and few-shot learning with instruction-based LLMs for CM, and provides possible further research questions and directions.

The creation of more datasets for claim matching will greatly benefit further development of the task. One possible reason why PD templates are well suited to the CM task is that existing datasets are not challenging enough. For example, two matching claims “London is the capital city of the UK” and “UK’s capital is London” can be easily identified as a paraphrase of each other. However, matching claims are not necessarily paraphrases of each other, e.g. “The president of the USA is Joe Biden” and “The president of the United States is Barack Obama”. The former is correct at the time of writing, whereas the latter is not correct, while they are matching claims because they can both be fact-checked with the same piece of evidence.

As an avenue for future work, to enhance model explainability, chain-of-thought methods can be implemented to let a model decompose the input and better understand nested sub-claims. Combinations of prompts could be also further examined.

## Limitations

The results depend on the choice of prompt templates and their formulations on surface level. Use of hand-crafted templates depends on the researcher’s intuition, leading to limitations in the experimentation and automating the prompt generation that is left for future work. Likewise, choice of few-shot train examples is also important. Alternative prompt-learning and -engineering methods should be implemented, automatically searching for best prompts and templates. Different LLMs require different prompt styles/selection of examples/prompt formatting. More fine-grained techniques, depending on a particular LLM, could improve their results and help understand why different prompt templates work better for different models. Fine-tuning approaches to LLMs, with more training data, could also reduce in the future the impact of the train examples selection. More detailed analysis of model error types is also the focus of further research. Using “ensemble” instructions, as a method, should be further investigated. On short texts, models in their explanations demonstrate that they mostly consider both system and user instruction combined.

Reproducibility of classification results, with the same LLM parameters, remains a challenging issue. But, for example, for five runs of GPT-3.5 and Gemini with PD-6 template, the results differed only slightly: mean F1 scores 84.7% and 90.5%, with standard error 0.25 and 0.10, respectively. The borderline examples are ambiguous, despite the gold label about their match.

Our experiments were conducted with English data only, which could be further extended to other languages using multilingual LLMs. CM is a new task, and only limited datasets are available. We used all two existing datasets, as a basis for our research. Dataset from (Nakov et al., 2022) contains English and Arabic parts. Dataset from (Kazemi et al., 2022) includes English, Hindi, Spanish, and Portuguese. We provide the first study on our CM approach and use only English parts, for prompt engineering and evaluation purposes, as we do not have native speakers in the aforementioned languages. There is indeed need to create more datasets, and we are planning to create more CM datasets, including the multilingual ones.

## References

- Carlos Alvarez, Maxwell Bennett, and Lucy Wang. 2024. [Zero-shot scientific claim verification using LLMs and citation text](#). In *Proceedings of the Fourth Workshop on Scholarly Document Processing (SDP 2024)*, pages 269–276, Bangkok, Thailand. Association for Computational Linguistics.
- Stephen Bach, Victor Sanh, Zheng Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak, Abheesh Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, Zaid Alyafeai, Manan Dey, Andrea Santilli, Zhiqing Sun, Srulik Ben-david, Canwen Xu, Gunjan Chhablani, Han Wang, Jason Fries, Maged Al-shaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangru Tang, Dragomir Radev, Mike Tian-jian Jiang, and Alexander Rush. 2022. [Prompt-Source: An integrated development environment and repository for natural language prompts](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 93–104, Dublin, Ireland. Association for Computational Linguistics.
- Alberto Barrón-Cedeño, Tamer Elsayed, Preslav Nakov, Giovanni Da San Martino, Maram Hasanain, Reem Suwaileh, Fatima Haouari, Nikolay Babulkov, Bayan Hamdan, Alex Nikolov, Shaden Shaar, and Zien Sheikh Ali. 2020. [Overview of checkthat!2020: Automatic identification and verification of claims in social media](#). *Preprint*, arXiv:2007.07997v1.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. [Sparks of artificial general intelligence: Early experiments with gpt-4](#). *Preprint*, arXiv:2303.12712.
- Jifan Chen, Grace Kim, Aniruddh Sriram, Greg Durrett, and Eunsol Choi. 2024. [Complex claim verification with evidence retrieved in the wild](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3569–3587, Mexico City, Mexico. Association for Computational Linguistics.

- Eun Cheol Choi and Emilio Ferrara. 2024a. [Automated claim matching with large language models: Empowering fact-checkers in the fight against misinformation](#). In *Companion Proceedings of the ACM Web Conference 2024, WWW '24*, page 1441–1449, New York, NY, USA. Association for Computing Machinery.
- Eun Cheol Choi and Emilio Ferrara. 2024b. [Fact-gpt: Fact-checking augmentation via claim matching with llms](#). In *Companion Proceedings of the ACM Web Conference 2024, WWW '24*, page 883–886, New York, NY, USA. Association for Computing Machinery.
- Hai Dang, Lukas Mecke, Florian Lehmann, Sven Goller, and Daniel Buschek. 2022. [How to prompt? opportunities and challenges of zero- and few-shot learning for human-ai interaction in creative applications of generative models](#). *Preprint*, arXiv:2209.01390.
- Yong Deng, Chenxiao Dou, Liangyu Chen, Deqiang Miao, Xianghui Sun, Baochang Ma, and Xiangang Li. 2022. [BEIKE NLP at SemEval-2022 task 4: Prompt-based paragraph classification for patronizing and condescending language detection](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 319–323, Seattle, United States. Association for Computational Linguistics.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021a. [Making pre-trained language models better few-shot learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021b. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2022. [PPT: Pre-trained prompt tuning for few-shot learning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8410–8423, Dublin, Ireland. Association for Computational Linguistics.
- Yi-Li Hsu, Jui-Ning Chen, Yang Fan Chiang, Shang-Chien Liu, Aiping Xiong, and Lun-Wei Ku. 2024. [Enhancing perception: Refining explanations of news claims with LLM conversations](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2129–2147, Mexico City, Mexico. Association for Computational Linguistics.
- Shengding Hu, Ning Ding, Huadong Wang, Zhiyuan Liu, Jingang Wang, Juanzi Li, Wei Wu, and Maosong Sun. 2022. [Knowledgeable prompt-tuning: Incorporating knowledge into prompt verbalizer for text classification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2225–2240, Dublin, Ireland. Association for Computational Linguistics.
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. [How can we know what language models know?](#) *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Ashkan Kazemi, Kiran Garimella, Devin Gaffney, and Scott Hale. 2021. [Claim matching beyond English to scale global fact-checking](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4504–4517, Online. Association for Computational Linguistics.
- Ashkan Kazemi, Zehua Li, Veronica Perez-Rosas, Scott A. Hale, and Rada Mihalcea. 2022. [Matching tweets with applicable fact-checks across languages](#). *Preprint*, arXiv:2202.07094.
- Teven Le Scao and Alexander Rush. 2021. [How many data points is a prompt worth?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2627–2636, Online. Association for Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Miaoran Li, Baolin Peng, Michel Galley, Jianfeng Gao, and Zhu Zhang. 2024. [Self-checker: Plug-and-play modules for fact-checking with large language models](#). *Preprint*, arXiv:2305.14623.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- Robert Logan IV, Ivana Balazevic, Eric Wallace, Fabio Petroni, Sameer Singh, and Sebastian Riedel. 2022. [Cutting down on prompts and parameters: Simple few-shot learning with language models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2824–2835, Dublin, Ireland. Association for Computational Linguistics.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2022. [Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity](#). In *Proceedings of the*

- 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 8086–8098, Dublin, Ireland. Association for Computational Linguistics.
- Watheq Mansour, Tamer Elsayed, and Abdulaziz Al-Ali. 2022. Did i see it before? detecting previously-checked claims over twitter. In *Advances in Information Retrieval*, pages 367–381, Cham. Springer International Publishing.
- Simona Mihaylova, Iva Borisova, Dzhovani Chemishev, Preslav Hadzhitsanev, Momchil Hardalov, and Preslav Nakov. 2021. Dips at checkthat! 2021: Verified claim retrieval. In *CEUR Workshop Proceedings (CEUR-WS.org) CLEF 2021 – Conference and Labs of the Evaluation Forum, September 21–24, 2021, Bucharest, Romania*.
- Preslav Nakov, Giovanni Da San Martino, Firoj Alam, Shaden Shaar, Hamdy Mubarak, and Nikolay Babulov. 2022. Overview of the CLEF-2022 Check-That! lab task 2 on detecting previously fact-checked claims. In *Proceedings of the Working Notes of CLEF 2022 - Conference and Labs of the Evaluation Forum*.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. 2021. True few-shot learning with language models. In *Advances in Neural Information Processing Systems 34 - 35th Conference on Neural Information Processing Systems, NeurIPS 2021*, Advances in Neural Information Processing Systems, pages 11054–11070.

- Neural information processing systems foundation. Publisher Copyright: © 2021 Neural information processing systems foundation. All rights reserved.; 35th Conference on Neural Information Processing Systems, NeurIPS 2021 ; Conference date: 06-12-2021 Through 14-12-2021.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah Smith, and Mike Lewis. 2023. [Measuring and narrowing the compositionality gap in language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5687–5711, Singapore. Association for Computational Linguistics.
- Albert Pritzkau. 2021. Nlytics at checkthat! 2021: Detecting previously fact-checked claims by measuring semantic similarity. In *CEUR Workshop Proceedings (CEUR-WS.org) CLEF 2021 – Conference and Labs of the Evaluation Forum, September 21–24, 2021, Bucharest, Romania*.
- Guanghui Qin and Jason Eisner. 2021. [Learning how to ask: Querying LMs with mixtures of soft prompts](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5203–5212, Online. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021a. [Exploiting cloze-questions for few-shot text classification and natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021b. [It’s not just size that matters: Small language models are also few-shot learners](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2339–2352, Online. Association for Computational Linguistics.
- Shaden Shaar, Firoj Alam, Giovanni Da San Martino, and Preslav Nakov. 2022. [The role of context in detecting previously fact-checked claims](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1619–1631, Seattle, United States. Association for Computational Linguistics.
- Shaden Shaar, Firoj Alam, Giovanni Da San Martino, and Preslav Nakov. 2021a. [Assisting the human fact-checkers: Detecting all previously fact-checked claims in a document](#). *Preprint*, arXiv:2109.07410v1.
- Shaden Shaar, Nikolay Babulkov, Giovanni Da San Martino, and Preslav Nakov. 2020. [That is a known lie: Detecting previously fact-checked claims](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3607–3618, Online. Association for Computational Linguistics.
- Shaden Shaar, Fatima Haouari, Watheq Mansour, Maram Hasanain, Nikolay Babulkov, Firoj Alam, Giovanni Da San Martino, Tamer Elsayed, and Preslav Nakov. 2021b. [Overview of the clef-2021 checkthat! lab task 2 on detecting previously fact-checked claims in tweets and political debates](#). In *CEUR Workshop Proceedings (CEUR-WS.org) CLEF 2021 – Conference and Labs of the Evaluation Forum, September 21–24, 2021, Bucharest, Romania*, page 264–291.
- Sinong Wang, Han Fang, Madian Khabza, Hanzi Mao, and Hao Ma. 2021. [Entailment as few-shot learner](#). *Preprint*, arXiv:2104.14690.
- Ye Wang, Yanmeng Wang, Baishun Ling, Zexiang Liao, Shaojun Wang, and Jing Xiao. 2022. [PINGAN omniscient at SemEval-2022 task 4: Multi-prompt training for patronizing and condescending language detection](#). In *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, pages 313–318, Seattle, United States. Association for Computational Linguistics.
- Albert Webson and Ellie Pavlick. 2022. [Do prompt-based models really understand the meaning of their prompts?](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2300–2344, Seattle, United States. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Yeo Wei Jie, Ranjan Satapathy, Rick Goh, and Erik Cambria. 2024. [How interpretable are reasoning explanations from prompting large language models?](#) In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2148–2164, Mexico City, Mexico. Association for Computational Linguistics.
- Xia Zeng, Amani S Abumansour, and Arkaitz Zubiaga. 2021. Automated fact-checking: A survey. *Language and Linguistics Compass*, 15(10):e12438.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and Ed Chi. 2023. [Least-to-most prompting enables complex reasoning in large language models](#). *Preprint*, arXiv:2205.10625.

## **A Appendix**

Table 4 shows 5 matching and 5 non-matching claim pairs from ClaimMatch, whereas Table 5 shows 10 randomly selected claim pairs used for few-shot training.

Claim 1	Claim 2	Class label
1. Here is a homeless, elderly couple from California sleeping on concrete with no water from a toilet or even a toilet at all. If only they were illegals, the Democrats might actually care about their well-being — derek schwartz (derek_mafs) July 13, 2019	Is This an ‘Elderly Homeless Couple’ in California? Viral social media posts claimed Democratic lawmakers might care about the couple’s well-being if they were "illegals.". A photograph depicts an elderly homeless couple sleeping on the concrete in California.	Positive
2. netflix what is this madness?? the literal source of my happiness, a.k.a. Friends, is leaving?? heartbroken — Noelle Michaud (noellemichaud14) December 2, 2018	Did Netflix Announce They Will Be Removing ‘Friends’ from Their Streaming Catalog? The online streaming service caused consternation in December 2018 when they briefly signalled the hit 90’s sitcom would no longer be available after 1 January 2019. Netflix announced that ‘Friends’ would no longer be available for customers to stream after the end of 2018.	Positive
3. The Sheriff in Union County, Arkansas is putting Nike t-shirts on people they arrest and making them wear them during mugshots. Source says it is to mock Nike and Colin Kaepernick. Disgusting. — Shaun King (shaunking) October 11, 2018	Did a Sheriff’s Dept. Dress Arrestees in Nike Shirts for Their Mugshots? Was the use of Nike shirts for mugshots a subtle jab at the company’s political choices, or a practical use of available material? The Union County Sheriff’s Department in Arkansas dressed numerous arrestees in Nike shirts for their mugshots as a jab at the company.	Positive
4. Should we even be surprised anymore? Peter Strzok’s sister-in-law works with Christine Blasey Ford’s brother. (Corrected version with Jill Strzok) — PaleoHorse (PaleoHorse) September 26, 2018	Did Christine Blasey Ford’s Brother Work with a Close Relative of Peter Strzok’s? Conspiracy theorists claim they’ve found ‘three degrees of separation’ between fired FBI agent Peter Strzok and Brett Kavanaugh’s sexual assault accuser Christine Blasey Ford. Christine Blasey Ford’s brother, Tom Blasey, worked in the same company as a close relative (wife, sister, or sister-in-law) of former FBI agent Peter Strzok.	Positive
5. Shocking video from 2017 shows Nancy Pelosi describing the Democrat attack plan called the “Wrap-up Smear” (the technique used against Justice Brett Kavanaugh ) movingUSforward — THE SCOOP (TheScoop_US) October 8, 2018	Did Nancy Pelosi Admit Democrats Use a Tactic Called the ‘Wrap-Up Smear’?. Unreliable sources claimed Pelosi admitted on video that Democrats use such a tactic, but in reality she ascribed it to Republicans. In a C-SPAN video, U.S. House Minority Leader Nancy Pelosi revealed that Democrats use a political smear tactic she called the "wrap-up smear."	Positive
6. Polls are starting to look really bad for Obama. Looks like he’ll have to start a war or major conflict to win. Don’t put it past him! — Donald J. Trump (realDonaldTrump) October 17, 2012	Is Amazon Bankrupting the United States Postal Service? President Donald Trump has claimed the giant online retailing company is jeopardizing USPS’s financial health. Amazon’s use of the USPS for package delivery services has resulted in a severe financial loss to the Postal Service.	Negative
7. Suspended by facebook. For blowing the whistle. On something they have known privately for 2 years. — Christopher Wylie (chrisinsilico) March 18, 2018	Did George Soros Pay ‘March for Our Lives’ Protesters \$300 Each? As hundreds of thousands of students and supporters took to the streets to protest gun violence, conspiracy theorists blamed a familiar bogeyman. Organizers working for billionaire George Soros ran Craigslist ads offering \$300 each to individuals participating in the 24 March 2018 "March for Our Lives" protests.	Negative
8. Earlier, Planned Parenthood suggested ‘we need a Disney princess who’s had an abortion’ and then they deleted it. So, in case you missed it: — Chet Cannon (Chet_Cannon) March 27, 2018	Did an Instagram Post by Rihanna Cause Snapchat’s Stock Market Value to Fall? The megastar’s exhortation to her 61 million followers to "throw away" the app likely did make a significant contribution to a \$750 million one-day loss. Rihanna’s Instagram message to followers to throw away the Snapchat app caused the company’s share value to fall by hundreds of millions of dollars in one day.	Negative
9. Yeti cuts ties with the NRA Foundation — Mark R. Levin (marklevinshow) April 23, 2018	Did Donald Trump Tweet Warnings About Obama Ordering Syrian Airstrikes? In 2013, Donald Trump issued multiple warnings to his predecessor about attacking Syria without congressional approval. Donald Trump leveled the same criticisms against President Barack Obama over Syrian strikes that were later used against him.	Negative
10. Today we are flipping our iconic McDonalds logo in honour of women everywhere. IWD2018 — Steve Easterbrook (SteveEasterbrk) March 8, 2018	Does Delta Airlines Give Planned Parenthood Members Discounted Rates? A Georgia senator has claimed — without any supporting evidence — that Delta offers discounts on air travel to Planned Parenthood members. Delta Air Lines gives members of Planned Parenthood discounted rates on air travel.	Negative

Table 4: Positive and negative class examples from the ClaimMatch dataset. 9735

Claim 1	Claim 2	Class label
1. Statement 1: How are butterflies surviving the Australian Fires? Julie Favell was putting out water for wildlife that survived the fires when she witnessed common brown butterflies ( <i>Heteronympha merope</i> ) fluttering in a moist wombat hole. Footage by Julie Favell — Center for Bio Div (CenterForBioDiv) January 14, 2020	Statement 2: Are Wombats Inviting Animals Into Their Burrows to Escape Australia Fires? Gather 'round to hear the tale of the wombat hero ... or at least the wombat's big burrow. Wombats are herding animals and inviting them into their burrows in order to escape the wildfires in Australia.	Positive
2. Statement 1: Trump needs to immediately divest from his businesses and comply with the emoluments clause. Iran could threaten Trump hotels *worldwide* and he could provoke war over the loss of revenue from skittish guests. His business interests should not be driving military decisions. — Ilhan Omar (IlhanMN) January 6, 2020	Statement 2: No, U.S. Rep. Ilhan Omar Didn't Give 'Treasonous' Military Advice to Iran. An article by the anti-Muslim activist Robert Spencer prompted threats and incitements of violence and murder against the Minnesota congresswoman in January 2020. In January 2020, U.S. Rep. Ilhan Omar advised Iran to attack Trump-branded hotels in the world, thus committing treason.	Positive
3. Statement 1: Once more for the folks who missed the tweet the first time I posted it. Minecraft is NOT stopping, Mojang is NOT closing. Minecraft — Helen Z PAXSouth (HelenAngel) January 2, 2020	Statement 2: Did Trump Ask Advisers About 'Nuking' Hurricanes? A report from Axios suggested the question was raised on more than one occasion. U.S. President Donald Trump has asked his advisers about the feasibility of stopping hurricanes with nuclear bombs.	Negative
4. Statement 1: Earlier today, a news outlet accurately reported that a subcontractor for one of our vendors was using prison workers to make phone calls on behalf of my campaign. After learning this, we immediately ended our relationship with that company. Full statement below: — Mike Bloomberg (MikeBloomberg) December 24, 2019	Statement 2: Is Minecraft Shutting Down? Internet pranks are not confined to April Fools' Day. The popular video game Minecraft is shutting down in 2020.	Negative
5. Statement 1: A number of fraudulent text messages informing individuals they have been selected for a military draft have circulated throughout the country this week.	Statement 2: Is US Army Sending Texts About a Military Draft? The military has been an all-volunteer force since 1973. The U.S. Army is sending text messages informing people they've been selected for the military draft.	Positive
6. Statement 1: This is the third spill along the Keystone pipeline's route in less than three years. Pipeline projects like this will continue to spill and pollute the earth until we shut them down for good and replace with 100% clean energy. Zeroto100 — Earthjustice (Earthjustice) October 31, 2019	Statement 2: Did Michael Bloomberg's Presidential Campaign Use Prison Labor? The former New York mayor landed in hot water over a vendor's reported use of prison workers to place calls to would-be voters. In 2019, the presidential campaign of Michael Bloomberg worked with a vendor that used prison labor on the campaign's behalf.	Negative
7. Statement 1: America has more governors who've worn blackface than black governors. — Samuel Sinyangwe (samswey) August 30, 2019	Statement 2: Was the Keystone Pipeline Shut Down After Leaking Oil? The pipeline did leak in North Dakota, but claims by alternative media that "no one" was talking about it were demonstrably false. The Keystone Pipeline in North Dakota was shut down after it leaked oil.	Negative
8. Statement 1: Hurricane Dorian washed up bricks of cocaine on Florida's coast — NowThis (nowthisnews) September 10, 2019	Statement 2: Does America Have More Governors Who Have Worn Blackface Than Black Governors? The question gets at the heart of whether an African American currently leads the executive branch of any state in the United States. The U.S. has more governors who have worn blackface than actual black governors.	Negative
9. Statement 1: The US drone attack on Soleimani caught on camera. IranUsa — Olaudah Equiano (RealOlaudah) January 6, 2020	Statement 2: Does This Video Show the Drone Strike That Killed Soleimani? Iran Gen. Qassem Soleimani was killed by a U.S. drone strike in Baghdad on Jan. 3, 2020. A video shows the U.S.-ordered drone strike that killed Iran Gen. Qassem Soleimani.	Positive
10. Statement 1: Picture of engagement rings removed from the fingers of World War- II victims. Imagine how many love stories were buried! Wars are never easy!Nd we r making fun of it.. shameful act.. WWII WW3 — ChaudhRy Saab (tera_Lover_) January 5, 2020	Statement 2: Does This Photo Show Wedding Rings Taken from Holocaust Victims? "Every wedding ring here represents a home broken and a human murdered by the Germans." A photograph documents a cache of wedding rings removed from Holocaust victims.	Positive

Table 5: Few-shot train examples.