

# Little Giants: Exploring the Potential of Small LLMs as Evaluation Metrics in Summarization in the Eval4NLP 2023 Shared Task

Neema Kotonya    Saran Krishnasamy  
Joel Tetreault    Alejandro Jaimes

Dataminr Inc.

{neema.kotonya,skrishnasamy,jtetreault,ajaimes}@dataminr.com

## Abstract

This paper describes and analyzes our participation in the 2023 Eval4NLP shared task, which focuses on assessing the effectiveness of prompt-based techniques to empower Large Language Models to handle the task of quality estimation, particularly in the context of evaluating machine translations and summaries. We conducted systematic experiments with various prompting techniques, including standard prompting, prompts informed by annotator instructions, and innovative chain-of-thought prompting. In addition, we integrated these approaches with zero-shot and one-shot learning methods to maximize the efficacy of our evaluation procedures. Our work reveals that combining these approaches using a “small”, open source model (orca\_mini\_v3\_7B) yields competitive results.

## 1 Introduction

Large Language Models (LLMs) have revolutionized the field of Natural Language Processing (NLP) by demonstrating remarkable proficiency in a multitude of generative tasks (Brown et al., 2020). Beyond their capabilities in text generation, LLMs offer the potential to automate the evaluation of generated text, particularly in domains such as machine translation and summarization.

Previous research efforts have explored LLM-based evaluation metrics, yielding promising results. Notable examples include the development of metrics like the GEMBA metric for translation quality assessment (Kocmi and Federmann, 2023), work on the effectiveness of LLMs as an alternative to human evaluation for NLP tasks by Chiang and Lee (2023), and the INSTRUCTSCORE metric for summarization evaluation (Xu et al., 2023). However, a significant gap exists in the systematic evaluation and exploration of prompting techniques available for metric usage with LLMs. In fact, there is scant work in this area to date. Excep-

tions include the work of Mendonça et al. (2023) for dialogue evaluation, Yang et al. (2023) and GEMBA for MT evaluation, G-EVAL, a chain-of-thought based framework for the evaluation of generated texts that leverages GPT-4 (Liu et al., 2023), and GPTSCORE for text generation evaluation (Fu et al., 2023).

This paper presents our contribution to addressing this gap through our participation in the *Prompting Large Language Models as Explainable Metrics* shared task (Leiter et al., 2023), which was conducted as part of The 4th Workshop on Evaluation & Comparison of NLP Systems, hosted at ACL 2023. We delve into various prompting approaches and techniques, offer a comprehensive overview of the results we have obtained, and provide a thorough analysis of our findings (our team name is COMPETITIONENTRANTS).

We exclusively participated in the Small Models Track, focusing on models with parameters less than or equal to 25 billion, for the Summarization task. Consequently, all our experiments and reported results are derived from the orca\_mini\_v3\_7B model. Among our various approaches, the best-performing one, employing a standard prompt in a zero-shot setting, achieved a score of 0.44 during the test phase of the shared task. While this performance is notable, it places us marginally behind the leaderboard’s highest score of 0.50 by a margin of 0.06. We also find that Chain-of-Thought (CoT) prompting (Wei et al., 2022) also aids in explicating the evaluation task to the model. This technique holds the potential to enhance the interpretability and explainability of quality estimation models.

## 2 Task Description

The primary objective of the shared task is to investigate prompt-based methodologies for LLMs in the development of automated quality metrics in

a reference-free setup tailored to natural language generation tasks, specifically summarization and machine translation. These quality scores are constructed using fine-grained scoring or error labels. The underlying rationale behind reference-free evaluation metrics is to provide assessment scores that are comparable to those of reference-based metrics while reducing reliance on often noisy and costly reference labels. A typical reference-free metric operates by taking a source (SRC) and a hypothesis (HYP) as inputs and subsequently generates a metric score, thereby providing an alternative to traditional reference-based evaluation methodologies.

## 2.1 Task Setup

The shared task is structured into two distinct sub-tasks, each contributing to the overarching goal:

1. **Prompting Strategies for Evaluation of Summarization and Translation:** This entails designing prompts and instructions that enable the assessment of the quality of generated content. The outcome of this sub-task serves as a critical component in the reference-free evaluation process.
2. **Score Aggregation:** The second sub-task focuses on the creation of a score aggregator mechanism. The primary objective here is to devise a method that computes an overall quality estimation score based on the outputs generated in Task 1. This aggregator consolidates individual quality assessments into a comprehensive quality estimation, ultimately providing a holistic evaluation of translation and summarization performance.

## 2.2 Datasets

Data is provided for the tasks of summarization and machine translation:

- **Summarization:** The training and development data for this track is derived from the datasets detailed in SummEval (Fabbri et al., 2020) with the scores being the average of human annotations across the four aspects - coherence, consistency, fluency, and relevance.
- **Machine Translation:** The training and development datasets are derived from the Multi-dimensional Quality Metrics (MQM) annotations of the WMT22 metrics shared task (Freitag et al., 2022) for machine translation.

For the test data, we are provided with a new reference-free dataset with sentence/summary-level quality scores for summarization and MT. As part of the test phase, 2 new language pair datasets, English-Chinese and English-Spanish are introduced for the machine translation track. Table 1 shows the counts of the train, development, and test datasets.

task	train	dev	test	
Summarization	320	1,280	825	
Translation	en-de	11,046	7,364	1,425
	zh-en	15,750	10,500	-
	en-es	-	-	1,834
	en-zh	-	-	1,297

Table 1: Train, Dev, and Test dataset sizes for summarization and machine translation tasks. Entries with - indicate that the dataset wasn’t provided as part of this task.

## 2.3 Large language models

The following six Huggingface LLMs were permitted for use in the shared task, two larger models (65B and 70B parameter models), which we denote with the following symbol ●, and four smaller models, denoted by ★, each of which has 13B parameters or fewer:

1. Guanaco-65B-GPTQ●<sup>1</sup>
2. WizardLM-13B-V1.1-GPTQ★<sup>2</sup>
3. Nous-Hermes-13b★<sup>3</sup>
4. Platypus2-70B-Instruct-GPTQ●<sup>4</sup>
5. OpenOrca-Platypus2-13B★<sup>5</sup>
6. Orca\_mini\_v3\_7b★<sup>6</sup>

Despite having access to these LLMs, our work faced computational constraints that influenced our choice of models for experimentation. As a result, we focused primarily on experimenting with two small LLMs: orca\_mini\_v3\_7B model and Nous-Hermes-13b. During the submission phase to the shared task’s leaderboard, the final test results we presented were exclusively derived from the orca\_mini\_v3\_7B model. The shared task guidelines explicitly forbade model fine-tuning.

<sup>1</sup><https://huggingface.co/TheBloke/guanaco-65B-GPTQ>

<sup>2</sup><https://huggingface.co/TheBloke/WizardLM-13B-V1.1-GPTQ>

<sup>3</sup><https://huggingface.co/NousResearch/Nous-Hermes-13b>

<sup>4</sup><https://huggingface.co/TheBloke/Platypus2-70B-Instruct-GPTQ>

<sup>5</sup><https://huggingface.co/Open-Orca/OpenOrca-Platypus2-13B>

<sup>6</sup>[https://huggingface.co/pankajmathur/orca\\_mini\\_v3\\_7b](https://huggingface.co/pankajmathur/orca_mini_v3_7b)

## 2.4 Evaluation

For the evaluation process, we used Codalab as the platform for submitting our system entries. Notably, the organizers of the evaluation, as detailed by [Kocmi and Federmann \(2023\)](#), provide direct assessment baselines for these LLMs. These baselines serve as reference points for evaluating the performance of our system and other participants in the shared task.

To quantify the performance of our system and the competing teams, the shared task organizers utilized the Kendall rank coefficient as the evaluation metric. The Kendall rank coefficient stands as an alternative to more traditional correlation metrics like Pearson’s  $r$  and Spearman’s  $\rho$  correlations. It finds particular utility in situations where the data fails to meet specific assumptions or when dealing with relatively small sample sizes.

## 3 Approaches

Three main classes of strategies are employed to enhance prompt effectiveness and interpretability for evaluating generated summaries. The first strategy, Core Prompts, encompasses three one-step methods for generating prompts. The first two borrow from existing literature, and the final uses an LLM to simply generate a prompt from scratch. Next, in Section 3.2, we introduce three methods (one manual, two automatic) to take prompts in Section 3.1 and further refine them. This is akin to paraphrasing in bulk. Finally, in Section 3.3, we outline two simple approaches for further refining prompts generated earlier. In total, the three classes span 8 different approaches, and approaches can be combined across the classes.

### 3.1 Core Prompts

**(1) Standard Prompting:** Our initial approach was to formulate prompts (a total of 9 prompts) that assess summary quality across the four dimensions outlined in [Kryscinski et al. \(2019\)](#): fluency, coherence, consistency, and relevance. These prompts task the model with generating quality scores for summaries, for different score ranges. In Table 2, we provide an example of standard prompting. In this example, the prompt specifies that the summarization should be rated from 1 and 5, with increments of 0.5 permitted.

Given the following summary for a news article, evaluate this summary for its fluency, coherence, consistency and relevance. Provide an overall score for the quality of this summary in the range 1 (worst) to 5 (best). Possible scores are 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5 and 5.

Table 2: An example of a standard prompt for summarization quality evaluation that stipulates scores should be in the range 1 to 5, and intermediate scores should be in 0.5 intervals.

**(2) Annotator Instructions as Seed Prompts:** To facilitate summarization evaluation, we then employed the instructions provided to expert annotators in [Fabbri et al. \(2020\)](#) (See Table 3). These served as foundational “seed prompts” for subsequent prompt refinement. We conducted an assessment of this seed prompt by utilizing a subset of examples from the Eval4NLP training dataset, noting that these instructions exhibit relatively favorable performance on the training data.

[. . .] In this task you will evaluate the quality of summaries written for a news article  
To correctly solve this task, follow these steps:  
1. Carefully read the news articles, be aware of the information it contains.  
2. Read the proposed summary.  
3. Rate each summary on a scale from 1 (Worst) to 5 (Best) by its relevance, consistency, fluency, and coherence.  
Relevance: “The rating measures how well the summary captures the key points of the article. Consider whether all and only the important aspects are contained in the summary.” [. . .]

Table 3: Instructions provided to expert annotators in [Fabbri et al. \(2020\)](#).

**(3) Prompt Generation via LLMs:** Additionally, we employed a separate LLM to generate a prompt entirely from scratch. The intuition behind this approach is that an LLM-derived prompt may yield improved results over a manually-crafted prompt. Similar intuition is followed in previous works that use LLMs to produce high-quality labels for LLM-generated texts ([Zellers et al., 2019](#); [Fu et al., 2023](#)). For prompt generation we author simple prompts, instructing the LLM to generate a sequence of instructions based on the requirements (score range and aspects to consider for summarization) specified in the prompt. (see Table 4).

Write a set of instructions to evaluate the quality of the summary of a news article according to its coherence, consistency, fluency, and relevance for each sentence in the summary with respect to the news article. Each aspect (coherence, consistency, fluency, and relevance) should be scores from 1 to 5. 1 is the worst possible score, 5 is the best possible score. Instructions:

Table 4: Instructions for prompting LLMs to generate prompts for summarization quality evaluation.

### 3.2 Prompt Refinement

To further enhance the prompts’ quality and effectiveness from Section 3.1, we employ three key strategies, one manual and two automatic:

**(4) Manual Prompt Rewriting:** This method involves meticulous manual rewriting of the instructions (done by the authors). We created prompt variations to elicit fine-grained answers, seek explanations for the provided answers, and employ templates specifying the desired answer format. These steps ensure that the instructions are refined to enhance the clarity and comprehensibility of the prompts. We also experiment with prompts that instruct the LLM to output both scores and explanations, similar to other works that explore both prediction (which in our case is quality estimation) and explanation jointly (Camburu et al., 2018; Kotonya and Toni, 2020; Wei et al., 2022). However, we find that prompting for explanations in addition to quality estimation yields poor results (see Table 12 in Appendix C).

**(5) Instruction Enhancement via LLMs:** In this strategy, we provide the seed prompt as context and prompt a separate LLM to enhance the existing instructions. We utilize various phrases such as “Improve the following instructions”, “Rewrite the following instructions to yield better responses”, “Write a more precise set of instructions”, and “Rewrite the instructions below in order to yield the best results” (see Table 5).

### System: You are an AI assistant that follows instruction extremely well. Help as much as you can.

### User :

**Improve the following instructions:**

“In this task you will evaluate the quality of summaries written for a news article

To correctly solve this task, follow these steps:

1. Carefully read the news article, be aware of the information it contains.

2. Read the proposed summary.

3. Rate each summary on a scale from 1 (Worst) to 5 (Best) by its relevance, consistency, fluency, and coherence... .”

**New instructions:**

Table 5: Example of a prompt supplied for LLM-based prompt refinement, where the instruction used is “Improve the following instructions.”

**(6) Chain-of-Thought (CoT) Prompting:** We also harness the CoT prompting technique, which guides the model’s evaluation process through a sequence of intermediate reasoning steps leading to the determination of the quality score for the provided summary. The main advantages of CoT prompts are that their specificity should reduce the number of inconsistencies in the generated response, yield responses that correlate strongly with human judgments and also allow for more complex reasoning. The CoT additions are marked in blue in Tables 6 and 7.

### 3.3 Exploration of Inference Settings

Given the constraint of not permitting model fine-tuning, we explore various inference settings to optimize model performance:

**(7) Zero-shot Approach:** In this setting, the model is evaluated and prompted to generate responses without any prior training specific to the evaluation task.

**(8) Few-shot and One-shot Approaches:** These approaches involve leveraging a limited amount of training data to serve as exemplars to direct the model’s responses. While we experimented with a one-shot setting, it is important to note that increasing the number of examples in the prompt had the unintended consequence of slowing down inference.

<p>1. Coherence: Assess how well the summary conveys a clear and logical message.</p> <p>2. Consistency: Check if the summary accurately represents the main points of the news article.</p> <p>3. Fluency: Evaluate the smoothness and readability of the summary.</p> <p>4. Relevance: Determine if the summary is relevant to the news article’s topic.</p> <p>For each sentence in the summary, assign a score from 1 to 5 for each aspect (coherence, consistency, fluency, and relevance).</p> <p>Example:</p> <p>Sentence 1: “The company announced a new product line.”</p> <p>Coherence: 4</p> <p>Consistency: 3</p> <p>Fluency: 3</p> <p>Relevance: 4</p> <p><b>Total Score: (Coherence + Consistency + Fluency + Relevance) / 4</b></p> <p>Total Score: <math>(4 + 3 + 3 + 4) / 4 = 14 / 4 = 3.5</math></p> <p>So, the summary has an overall score of 3.5 out of 5.</p>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Table 6: Example of a prompt generated for summarization quality estimation. These instructions demonstrate step-by-step, with the aid of an example, how the final score should be calculated.

<p><b>Example:</b></p> <p>1. Read the news article: “A new study found that regular exercise can significantly improve mental health.”</p> <p>2. Read the summary: “A study discovered that exercise has a significant impact on mental health.”</p> <p>3. Evaluate the summary based on the aspects:</p> <p>a. Coherence: 5 (The summary maintains a clear and logical flow of ideas.)</p> <p>b. Consistency: 5 (The main points of the news article are accurately represented.)</p> <p>c. Fluency: 5 (The summary is written in a smooth and easy-to-understand manner.)</p> <p>d. Relevance: 5 (The summary conveys the essential information from the news article.)</p> <p>4. Assign scores for each aspect: [ . . . ]</p>
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Table 7: Example of a chain-of-thought prompt generated for summarization quality estimation. In this chain of thought prompt, descriptions are generated for each of the four aspects (coherence, consistency, fluency, and relevance).

**4 Results**

Table 8 shows results for a battery of approaches. One should note that we tried many combinations of the approaches with different seed prompts, and the number of experiments is quite large. For the sake of simplicity, we report on key combinations that we uncovered. Prompt ID refers to the specific prompt that was used and the exact text can be

found in the Appendix.

Among the approaches, Prompt P1, which employs a standard manual prompt in a zero-shot setting with a grading scale ranging from 1 to 5, emerges as the top performer, achieving a notable score of 0.3211 on the development dataset. This was surprising as this is essentially one of the most straightforward approaches to try. However, it is perhaps unsurprising as the Direct Assessment baseline provided by the Shared Task is also a simple manual prompt in a zero-shot setting (Kocmi and Federmann, 2023).

Following closely behind, we find approaches that leverage annotator instructions as seed prompts (P12) and prompts generated using LLMs (P14 and P10), all in zero-shot settings, also stand out. These prompts, in contrast to standard ones, contain a wealth of details about the evaluation metrics, offering intricate guidance to the model.

However, our exploration of a one-shot setting indicates that this approach does not yield as promising results (as much as 0.1 behind our best approach combination). Further experimentation with the choice of examples provided to the model may be warranted to enhance its performance.

Notably, the incorporation of chain-of-thought prompting appears to be a beneficial strategy, as evidenced by the strong performance observed (our third, fourth, and fifth-best experiments).

On the lower end of the result scores, we find standard prompts (P5, P4, P7, P3, P8) that utilize grading scales with exceptionally high precision

Prompt ID	Approach	Score
P1	(1) Standard Prompt (7) Zero-shot	<b>0.3211</b>
P2	(1) Standard Prompt (7) Zero-shot	0.3075
P12	(2) Annotator Instructions (6) CoT (7) Zero-shot	0.2837
P14	(3) LLM-generated (6)CoT (7) Zero-shot	0.2827
P10	(3) LLM-generated (6)CoT (8) One-shot	0.2687
P6	(1) Standard Prompt (7) Zero-shot	0.2597
P9	(6) CoT (7) Zero-shot	0.2477
P11	(3) LLM-generated (6)CoT (8) One-shot	0.2245
P13	(6) CoT (8) One-shot	0.2244
P5	(1) Standard Prompt (7) Zero-shot	-0.0163
P4	(1) Standard Prompt (7) Zero-shot	-0.0172
P7	(1) Standard Prompt (7) Zero-shot	-0.0209
P3	(1) Standard Prompt (7) Zero-shot	-0.0255
P8	(1) Standard Prompt (7) Zero-shot	-0.0329
Baseline	Direct Assessment	0.3065
Baseline	Random	-0.0340

Table 8: Results of prompts for evaluating summarization. Score is the dev score obtained from the Codalab submission. The Prompt ID map to the full Prompts in the Appendix. The details for each of the Approaches can be found in Section 3.

or qualitative labels as quality measures. These prompts, while designed with attention to detail, exhibit comparatively lower scores in the evaluation, suggesting the importance of striking a balance between precision and other factors when designing prompts for quality estimation tasks.

Prompt P6 achieves the highest Kendall correlation coefficient on the test set with a score of 0.4423. Furthermore, P1 and P2 also achieve competitive scores of 0.4419 and 0.4422 respectively.

## 5 Discussion

The evaluation results we have presented furnish compelling evidence regarding the proficiency of LLMs in the domain of quality estimation for summarization. Our findings underscore the capacity of these models to provide valuable insights into the quality of generated outputs, thereby contributing to the advancement of evaluation methodologies within the field of NLP.

### 5.1 Insights

**Scoring Rubric Matters:** Our experimentation with scoring rubrics revealed an intriguing trend. While assessing the precision and range of scores requested from the model, we observed that quantitative scores with lower precision exhibited favorable performance. Surprisingly, the use of quali-

tative labels such as “Very Poor”, “Poor”, “Average”, “Good”, and “Very Good” to describe quality yielded comparatively less favorable results as can be seen from the dev scores in Table 9. This suggests that when instructing LLMs for quality estimation, a preference for quantitative, less granular scoring may be more effective.

**Effect on Performance through Examples:** To enhance performance, we incorporated explicit examples into zero-shot prompts for each score on the evaluation scale. Contrary to our expectations, the inclusion of examples did not yield a noticeable improvement in model performance. This observation highlights the nuanced nature of prompt design and underscores the need for tailored approaches that align with the unique characteristics of the task.

**Simpler Prompts Suffice:** A notable finding emerged from our exploration of prompt complexity. While we originally hypothesized that detailed prompts derived from SummEval annotation guidelines would outperform simpler prompts based on the same four quality dimensions (fluency, coherence, consistency, relevance), our results did not substantiate this hypothesis. However, it is worth noting that this approach shows promise, particularly when the aspect being evaluated is ambiguous to the model. The provision of detailed prompts with examples and context holds the potential to im-

Prompt*	Dev score
Provide an overall score for the quality of this summary in the range 1 (worst) to 5 (best). Possible scores are 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5 and 5.	0.32
Provide an overall score for the quality of this summary in the range 0 (worst) to 100 (best). Possible scores are 0, 10, 20, 30, 40, 50, 60, 70, 80, 90 and 100.	0.31
Provide an overall score for the quality of this summary in the range 1 (worst) to 5 (best) that is an average of the scores (also from 1 to 5) for fluency, coherence, consistency and relevance.	-0.03
Provide an overall score for the quality of this summary in the range 0 (worst) to 100 (best). Possible scores are 0, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50, 55, 60, 65, 70, 75, 80, 85, 90, 95 and 100.	-0.02
Provide an overall score for the quality of this summary in the range Very Poor (worst) to Very Good (best). Possible scores are Very Poor, Poor, Average, Good, Very Good.	-0.02
Provide an overall score for the quality of this summary in the range Incomprehensible (worst) to Excellent (best). Possible scores are Incomprehensible, Poor, Average, Good, Excellent.	-0.03

Table 9: Standard prompting with different score ranges\*The prompts in the table are appended to *Given the following summary for a news article, evaluate this summary for its fluency, coherence, consistency and relevance.* along with the *input* to form the full instruction to the model.

prove performance, suggesting a fruitful direction for future research.

## 5.2 Hallucinations during prompt refinement

We conducted an experiment where we leveraged a separate LLM to generate instructions for assessing the quality of summarizations. This approach yielded promising results initially, as the generated instructions were detailed and exhibited potential when tested with a subset of examples from the training dataset. However, during our extensive evaluation process, we encountered instances of hallucinations within the generated instructions, prompting us to incorporate further refinement steps into our methodology. Hallucinations in this context refer to the phenomenon where the model produces content in the instruction that deviates from the original intent or context, thereby introducing inaccuracies or inconsistencies (Maynez et al., 2020).

One recurring form of hallucination pertained to the numbering within the instructions. In particular, we observed instances where the numbering in the latter part of the instruction was hallucinated, leading to inconsistencies. For instance, as shown in Table 10, the numbering sequence (2, 3, 4, 4) in the model’s response was incongruent with the original (1, 2, 3, 4) prompt. Additionally, we identified another hallucination wherein the phrase “Answer:” was erroneously introduced in the model’s output.

This unintended addition in the response further underscores the need for vigilance when using LLMs to generate instructions. For our experiments, we manually removed LLM-generate prompts with hallucinations.

## 6 Conclusion

In this paper, we have presented our system entry for the Eval4NLP 2023 Shared Task on Prompting Large Language Models as Explainable Metrics. Our approach revolves around assessing the capabilities of the small `orca_mini_v3_7B` model in the reference-free evaluation of generated summaries. Through a series of extensive experiments, we have explored the model’s adaptability and responsiveness to various prompting techniques, including zero-shot, one-shot, and chain-of-thought prompts. These experiments were conducted using meticulously designed prompts that emphasize key evaluation criteria. Our findings suggest that the use of a standard prompting (when using a small model) is performant while the chain-of-thought prompt-based approach holds significant promise for evaluating the quality of generated summaries. However, we do note that slight variations in the underlying prompt text could swing performances rather widely, so more research is definitely required to understand this phenomenon.

Response without hallucinations	Response with hallucinations
<p>1. Relevance: 2 Explanation: The summary does not capture the key points of the article. It only mentions the man gently cupping the fish, but it does not mention the relationship between the man and the fish, the video, or the fish’s behavior.</p> <p>2. Consistency: 2 Explanation: The summary does not reproduce all facts accurately. It does not mention the fish’s hybrid origin or the bullying of the other fish.</p> <p>3. Fluency: 2 Explanation: The summary is not well-written and grammatically correct. It contains grammatical errors and is not easy to understand.</p> <p>4. Coherence: 2 Explanation: The summary does not sound naturally and collectively. It is difficult to understand and does not flow well as a whole.</p>	<p>2. Relevance: 3 Explanation: The summary captures the key points of the article, but it is missing some important aspects, such as the half-and-half scarves and the fact that the match was played at Altrincham’s Moss Lane. <b>Answer:</b></p> <p>3. Consistency: 3 Explanation: The summary does reproduce all facts accurately, but it makes up untrue information by stating that the match was played at Old Trafford. <b>Answer:</b></p> <p>4. Fluency: 3 Explanation: The individual sentences are well-written and grammatically correct, but the summary could be more coherent and sound more naturally. <b>Answer:</b></p> <p>4. Coherence: 3 Explanation: The summary is not as coherent as it could be, as it jumps from one topic to another without a clear flow.</p>

Table 10: Hallucinations observed in prompts that are refined using LLMs.

## References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris 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*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. e-snli: Natural language inference with natural language explanations. *Advances in Neural Information Processing Systems*, 31.
- Cheng-Han Chiang and Hung-yi Lee. 2023. [Can large language models be an alternative to human evaluations?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15607–15631, Toronto, Canada. Association for Computational Linguistics.
- Alexander R Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2020. Summeval: Re-evaluating summarization evaluation. *arXiv preprint arXiv:2007.12626*.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, Eleftherios Avramidis, Tom Kocmi, George Foster, Alon Lavie, and André F. T. Martins. 2022. [Results of WMT22 metrics shared task: Stop](#) using BLEU – neural metrics are better and more robust. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 46–68, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. [Gptscore: Evaluate as you desire](#). *CoRR*, abs/2302.04166.
- Tom Kocmi and Christian Federmann. 2023. [Large language models are state-of-the-art evaluators of translation quality](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203, Tampere, Finland. European Association for Machine Translation.
- Neema Kotonya and Francesca Toni. 2020. [Explainable automated fact-checking: A survey](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5430–5443, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Wojciech Kryscinski, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Neural text summarization: A critical evaluation](#). 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 540–551, Hong Kong, China. Association for Computational Linguistics.
- Christoph Leiter, Juri Opitz, Daniel Deutsch, Yang Gao, Rotem Dror, and Steffen Eger. 2023. The eval4nlp 2023 shared task on prompting large language models as explainable metrics. In *Proceedings of the 4th Workshop on Evaluation and Comparison for NLP systems*.



Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruo Chen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using GPT-4 with better human alignment](#). *CoRR*, abs/2303.16634.

Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.

John Mendonça, Patrícia Pereira, Helena Moniz, Joao Paulo Carvalho, Alon Lavie, and Isabel M Trancoso. 2023. [Simple LLM prompting is state-of-the-art for robust and multilingual dialogue evaluation](#). In *Proceedings of The Eleventh Dialog System Technology Challenge*, pages 133–143, Prague, Czech Republic. 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.

Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Yang Wang, and Lei Li. 2023. [Instructscore: Towards explainable text generation evaluation with automatic feedback](#).

Hao Yang, Min Zhang, Shimin Tao, Minghan Wang, Daimeng Wei, and Yanfei Jiang. 2023. [Knowledge-prompted estimator: A novel approach to explainable machine translation assessment](#).

Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. [Defending against neural fake news](#). *Advances in neural information processing systems*, 32.

## A Software Libraries and Hardware

We use PyTorch v2.0.1 compiled with CUDA 11.7. We use version 0.0.64 of the guidance library<sup>7</sup> for prompt generation and prompt evaluation. We conduct all experiments on NVIDIA T4 and NVIDIA A10G GPUs.

---

<sup>7</sup><https://pypi.org/project/guidance/>

## B Prompt Templates

### B.1 Summarization

Below we provide our prompt templates which we use for the experiments described in this paper.

Prompt ID	Prompt
P1	Given the following summary for a news article, evaluate this summary for its fluency, coherence, consistency and relevance. Provide an overall score for the quality of this summary in the range 1 (worst) to 5 (best). Possible scores are 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5 and 5.
P2	Given the following summary for a news article, evaluate this summary for its fluency, coherence, consistency and relevance. Provide an overall score for the quality of this summary in the range 0 (worst) to 100 (best). Possible scores are 0, 10, 20, 30, 40, 50, 60, 70, 80, 90 and 100.
P3	Given the following summary for a news article, evaluate this summary for its fluency, coherence, consistency and relevance. Provide an overall score for the quality of this summary in the range 1 (worst) to 5 (best) that is an average of the scores (also from 1 to 5) for fluency, coherence, consistency and relevance.
P4	Given the following summary for a news article, evaluate this summary for its fluency, coherence, consistency and relevance. Provide an overall score for the quality of this summary in the range 0 (worst) to 100 (best). Possible scores are 0, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50, 55, 60, 65, 70, 75, 80, 85, 90, 95 and 100.
P5	Given the following summary for a news article, evaluate this summary for its fluency, coherence, consistency and relevance. Provide an overall score for the quality of this summary in the range 0 (worst) to 100 (best). Possible scores are 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99, 100.
P6	Given the following summary for a news article, evaluate this summary for its fluency, coherence, consistency and relevance. Provide an overall score for the quality of this summary in the range -100 (worst) to 100 (best). Possible scores are -100, -50, 0, 50, 100.
P7	Given the following summary for a news article, evaluate this summary for its fluency, coherence, consistency and relevance. Provide an overall score for the quality of this summary in the range Very Poor (worst) to Very Good (best). Possible scores are Very Poor, Poor, Average, Good, Very Good.
P8	Given the following summary for a news article, evaluate this summary for its fluency, coherence, consistency and relevance. Provide an overall score for the quality of this summary in the range Incomprehensible (worst) to Excellent (best). Possible scores are Incomprehensible, Poor, Average, Good, Excellent.

P9	<p>### System: You are an AI assistant that follows instruction extremely well. Help as much as you can.</p> <p>### User:</p> <p>In this task you will evaluate the quality of summaries written for a news article To correctly solve this task, follow these steps:</p> <ol style="list-style-type: none"> <li>1. Carefully read the news articles, be aware of the information it contains.</li> <li>2. Read the proposed summary.</li> <li>3. Rate each summary on a scale from 1 (Worst) to 5 (Best) by its relevance, consistency, fluency, and coherence.</li> </ol> <p>Relevance: “The rating measures how well the summary captures the key points of the article. Consider whether all and only the important aspects are contained in the summary.”</p> <p>Consistency: “The rating measures the facts in the summary are consistent with the facts in the original article. Consider whether the summary does reproduce all facts accurately and does not make up untrue information.”</p> <p>Fluency: “This rating measures the quality of individual sentences, are they well-written and grammatically correct. Consider the quality of individual sentences.”</p> <p>Coherence: “The rating measures the quality of all sentences collectively, to the fit together and sound naturally. Consider the quality of the summary as a whole.”</p> <p>Format the response as follows:</p> <p>Answer:</p> <p>Relevance: &lt;Rating for Relevance&gt;</p> <p>Explanation: &lt;Evidence for Relevance rating&gt;</p> <p>Consistency: &lt;Rating for Consistency&gt;</p> <p>Explanation: &lt;Evidence for Consistency rating&gt;</p> <p>Fluency: &lt;Rating for Fluency&gt;</p> <p>Explanation: &lt;Evidence for Fluency rating&gt;</p> <p>Coherence: &lt;Rating for Coherence&gt;</p> <p>Explanation: &lt;Evidence for Coherence rating&gt;</p> <p>News article: {source_text}</p> <p>Summary: {summary}</p> <p>### Assistant:</p> <p>Answer:</p>
----	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

P10	<p>### User:</p> <p>Evaluate the summary for a news article by assigning a score from 0 to 100 for each of the following aspects: Coherence, Consistency, Fluency, and Relevance.</p> <ol style="list-style-type: none"> <li>1. Coherence: <ul style="list-style-type: none"> <li>- Read the summary and determine if it is well-structured, easy to understand, and logically connected.</li> <li>- Assign a score from 0 to 100 based on how well the summary is coherent.</li> </ul> </li> <li>2. Consistency: <ul style="list-style-type: none"> <li>- Check if the summary accurately represents the main points and ideas from the original news article.</li> <li>- Assign a score from 0 to 100 based on how consistent the summary is with the original content.</li> </ul> </li> <li>3. Fluency: <ul style="list-style-type: none"> <li>- Evaluate the clarity and smoothness of the summary.</li> <li>- Assign a score from 0 to 100 based on how well the summary is written and easy to read.</li> </ul> </li> <li>4. Relevance: <ul style="list-style-type: none"> <li>- Determine if the summary effectively conveys the most important information from the original news article.</li> <li>- Assign a score from 0 to 100 based on how relevant and informative the summary is.</li> </ul> </li> </ol> <p>Once you have evaluated each aspect, add up the scores and assign a final score from 0 to 100 for the overall summary.</p> <p>News article: {source_text}</p> <p>Summary: {summary}</p> <p>### Assistant:</p> <p>Evaluation:</p>
-----	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

P11	<p>### User:</p> <ol style="list-style-type: none"> <li>1. Read the news article carefully.</li> <li>2. Read the summary of the news article.</li> <li>3. Evaluate the summary based on the following aspects: <ol style="list-style-type: none"> <li>a. Coherence: How well does the summary maintain a clear and logical flow of ideas?</li> <li>b. Consistency: Does the summary accurately represent the main points of the news article?</li> <li>c. Fluency: Is the summary written in a smooth and easy-to-understand manner?</li> <li>d. Relevance: Does the summary convey the essential information from the news article?</li> </ol> </li> <li>4. Assign a score from 1 to 5 for each aspect based on the evaluation.</li> <li>5. Average the scores for each aspect to get the overall score for the summary.</li> </ol> <p>Example:</p> <ol style="list-style-type: none"> <li>1. Read the news article: “A new study found that regular exercise can significantly improve mental health.”</li> <li>2. Read the summary: “A study discovered that exercise has a significant impact on mental health.”</li> <li>3. Evaluate the summary based on the aspects: <ol style="list-style-type: none"> <li>a. Coherence: 5 (The summary maintains a clear and logical flow of ideas.)</li> <li>b. Consistency: 5 (The main points of the news article are accurately represented.)</li> <li>c. Fluency: 5 (The summary is written in a smooth and easy-to-understand manner.)</li> <li>d. Relevance: 5 (The summary conveys the essential information from the news article.)</li> </ol> </li> <li>4. Assign scores for each aspect: <ol style="list-style-type: none"> <li>a. Coherence: 5</li> <li>b. Consistency: 5</li> <li>c. Fluency: 5</li> <li>d. Relevance: 5</li> </ol> </li> <li>5. Average the scores for each aspect: <ol style="list-style-type: none"> <li>a. Coherence: 5</li> <li>b. Consistency: 5</li> <li>c. Fluency: 5</li> <li>d. Relevance: 5</li> </ol> </li> <li>6. Average the scores for each aspect: 5</li> <li>7. Overall score for the summary: 5</li> </ol> <p>News article: {source_text}  Summary: {summary}</p> <p>### Assistant:  Assign scores for each aspect:</p>
-----	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

P12	<p>### User:</p> <p>Evaluate the summary for a news article by assigning a score from 0 to 100 for each of the following aspects: Coherence, Consistency, Fluency, and Relevance.</p> <ol style="list-style-type: none"> <li>1. Coherence: - Read the summary and determine if it is well-structured, easy to understand, and logically connected. - Assign a score from 0 to 100 based on how well the summary is coherent.</li> <li>2. Consistency: - Check if the summary accurately represents the main points and ideas from the original news article. - Assign a score from 0 to 100 based on how consistent the summary is with the original content.</li> <li>3. Fluency: - Evaluate the clarity and smoothness of the summary. - Assign a score from 0 to 100 based on how well the summary is written and easy to read.</li> <li>4. Relevance: - Determine if the summary effectively conveys the most important information from the original news article. - Assign a score from 0 to 100 based on how relevant and informative the summary is.</li> </ol> <p>Once you have evaluated each aspect, add up the scores and assign a final score from 0 to 100 for the overall summary.</p> <p>News article: {source_text} Summary: {summary}</p> <p>### Assistant: Evaluation:</p>
P13	<p>### User:</p> <ol style="list-style-type: none"> <li>1. Coherence: Assess how well the summary conveys a clear and logical message.</li> <li>2. Consistency: Check if the summary accurately represents the main points of the news article.</li> <li>3. Fluency: Evaluate the smoothness and readability of the summary.</li> <li>4. Relevance: Determine if the summary is relevant to the news article's topic.</li> </ol> <p>For each sentence in the summary, assign a score from 1 to 5 for each aspect (coherence, consistency, fluency, and relevance).</p> <p>Example: Sentence 1: "The company announced a new product line." Coherence: 4 Consistency: 3 Fluency: 3 Relevance: 4</p> <p>Total Score: (Coherence + Consistency + Fluency + Relevance) / 4 Total Score: (4 + 3 + 3 + 4) / 4 = 14 / 4 = 3.5</p> <p>So, the summary has an overall score of 3.5 out of 5.</p>

P14	<p>### User:</p> <p>To evaluate the coherence, consistency, fluency and relevance of a summary of a news article, follow these steps:</p> <ol style="list-style-type: none"><li>1. Read the news article carefully to understand its main points.</li><li>2. Read the summary and place it next to the news article so that you can compare them.</li><li>3. First, starts by evaluating the relevance of each sentence in the summary. Ask yourself if the sentence supports the main points of the news article and if it adds new and relevant information. You can assign a score to each sentence that evaluates this criterion and check whether it is coherent with the main points of the news article.</li><li>4. Next, evaluate the consistency of each sentence in the summary with respect to the main points of the news article. Ask yourself if the sentence is consistent with the overall message of the news article. You can assign a score to each criterion and check whether it is consistent with the overall purpose of the news article.</li><li>5. Start evaluating the fluency of each sentence. Check if each sentence flows well and if there are any awkward phrasing or sentence structures. You can assign a score to each sentence based on the fluency of the writing.</li><li>6. Assign a score of between 0 and 100 to each aspect (coherence, consistency, fluency and relevance) to evaluate how well the summary accurately and clearly conveys the article's content.</li></ol> <p>For example, if the summary accurately conveys the article's content and clearly and coherently communicates each aspect, you can assign a score of 100 for each aspect, but if it slightly or not at all conveys, you can assign a score of 0.</p>
-----	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## C Explanations

In Tables 12 and 13 present explanations generated for our experiments on the dev set. The prompt employed to generate the results in Tables 12 and 13 is shown in Table 14.

```
### Answer :  
2 (Worst)  
Relevance: 2  
Explanation: The summary contains some relevant information, but it is not comprehensive enough. It only mentions the guard slipping and falling, but it does not mention the incident being caught on camera or the reason for the slip.  
Consistency: 2  
Explanation: The summary is consistent with the original article, but it does not reproduce all facts accurately. It does not mention the metal protectors nailed to the soles of the guard's boots, which is an important detail.  
Fluency: 2  
Explanation: The summary is grammatically correct, but it is not well-written. It is difficult to understand and lacks proper sentence structure.  
Coherence: 2  
Explanation: The summary is coherent, but it is not well-structured. It is difficult to follow and does not flow naturally.
```

Table 12: Explanation and scores produced for a low-scoring summary. The explanations in the generated response are shown in red. Note that an overall assessment, 2 (*Worst*), is also generated for this example.

```
### Answer :  
Relevance: 4  
Explanation: The summary captures the key points of the article, but it could be more detailed and specific.  
Consistency: 4  
Explanation: The summary accurately reproduces the main facts from the article, but it could be more detailed.  
Fluency: 3  
Explanation: The summary is grammatically correct, but it could be more concise and fluent.  
Coherence: 3  
Explanation: The summary is coherent, but it could be more natural and flowing.
```

Table 13: Explanation and scores produced for a high-scoring summary from the dev set. The explanations in the generated response are shown in red.



```

### User:
In this task you will evaluate the quality of summaries written for a news article
To correctly solve this task, follow these steps:
1. Carefully read the news article, be aware of the information it contains.
2. Read the proposed summary.
3. Rate each summary on a scale from 1 (Worst) to 5 (Best) inclusive by its relevance, consistency, fluency, and coherence.
Relevance: "The rating measures how well the summary captures the key points of the article. Consider whether all and only the important aspects are contained in the summary."
Consistency: "The rating measures the facts in the summary are consistent with the facts in the original article. Consider whether the summary does reproduce all facts accurately and does not make up untrue information."
Fluency: "This rating measures the quality of individual sentences, are they well-written and grammatically correct. Consider the quality of individual sentences."
Coherence: "The rating measures the quality of all sentences collectively, to the fit together and sound naturally. Consider the quality of the summary as a whole."
Format the response as follows:
Answer:
Relevance: <Rating for Relevance>
Explanation: <Evidence for Relevance rating>
Consistency: <Rating for Consistency>
Explanation: <Evidence for Consistency rating>
Fluency: <Rating for Fluency>
Explanation: <Evidence for Fluency rating>
Coherence: <Rating for Coherence>
Explanation: <Evidence for Coherence rating>
News article: {source_text}
Summary: {summary}
### Assistant:

```

Table 14: Instructions that prompt the LLM to generate explanations in addition to quality scores for a summary.