

The Eval4NLP 2023 Shared Task on Prompting Large Language Models as Explainable Metrics



Christoph Leiter*, Juri Opitz†, Daniel Deutsch‡, Yang Gao◇
Rotem Dror††, Steffen Eger*

* Bielefeld University, Germany † Heidelberg University, Germany
‡ Google, US ◇ Google Research, UK †† University of Haifa, Israel
christoph.leiter@uni-bielefeld.de opitz.sci@gmail.com
danddeutsch@google.com gaostayyang@google.com
rdror@is.haifa.ac.il steffen.eger@uni-bielefeld.de

Abstract

Generative large language models (LLMs) have seen many breakthroughs over the last year. With an increasing number of parameters and pre-training data, they have shown remarkable capabilities to solve tasks with minimal or no task-related examples. Notably, LLMs have been successfully employed as evaluation metrics in text generation tasks. Approaches often differ in the input prompts, the samples that are selected for demonstration and the construction process of scores from the output. Within this context, we introduce the Eval4NLP 2023 shared task that asks participants to explore such approaches for machine translation evaluation and summarization evaluation. Specifically, we select a list of allowed LLMs and disallow fine-tuning to ensure a focus on prompting. We evaluate the approaches of the participants on a new reference-free test-set spanning 3 language pairs for machine translation as well as a summarization dataset. Further, we present an overview of the approaches taken by the participants, present their results on the test set and analyze paths for future work. Finally, as a separate track, we perform a small-scale human evaluation of the plausibility of explanations given by the LLMs. We make parts of our code and datasets available.¹

1 Introduction

The ChatGPT revolution in late 2022 has ignited a wide public and scientific debate about the possibilities (and limitations) of generative AI in various fields and application scenarios (Leiter et al.,

2023b; Eger et al., 2023), including education (Halaweh, 2023), logic (Liu et al., 2023a), medicine (Dave et al., 2023), math (Frieder et al., 2023), programming (Rozière et al., 2023) and science (Belouadi et al., 2023).

The immense research interest has also triggered the exploration of numerous approaches that leverage generative large language models (LLMs) as *evaluation metrics* (Kocmi and Federmann, 2023; Liu et al., 2023b; Fu et al., 2023; Xu et al., 2023b; Fernandes et al., 2023) for natural language generation (NLG) tasks like machine translation (MT) and summarization. Recent LLM based approaches differ, for example, in their prompting strategies, e.g., in the way that natural language instructions are used to trigger the LLM to compute metric scores. For example, GEMBA (Kocmi and Federmann, 2023) uses zero-shot prompting to directly predict scores or quality labels in the output. In contrast, AutoMQM (Fernandes et al., 2023) instructs LLMs to predict fine-grained error labels and uses these to compute the final scores. These works have contributed to the exploration of prompting for NLG evaluation, but an exhaustive exploration of approaches remains unaddressed. Further, many approaches leverage closed source LLMs while much fewer use open source LLMs. Those approaches relying on open source LLMs put a large focus on acquiring training data (e.g. Xu et al., 2023b) and fine-tune models to specific tasks. Given this typical focus on fine-tuning and motivated by promising work on prompting techniques² (e.g. Wei et al.,

¹<https://github.com/eval4nlp/SharedTask2023/tree/main>

²Various websites track the development of prompting techniques, e.g. <https://www.promptingguide.ai>.

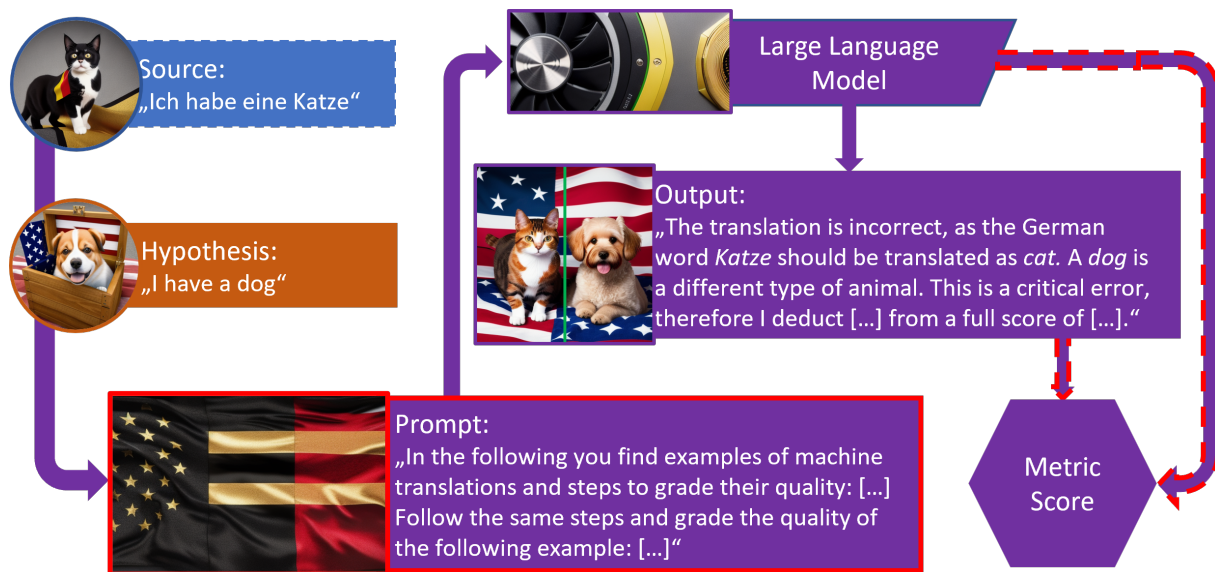


Figure 1: Using a generative LLM as MT evaluation metric. In this example, the metric is reference-free. I.e. it grades the translated sentence based on its source sentence. The input sentences are wrapped into a prompt that is given to an LLM. The LLM generates an output and a final score could for example be constructed from this textual output or from other values involved in the process. The red borders indicate the focus of our shared task. Participants should evaluate the best prompts and the best approaches to construct scores from model output.

2022; Yao et al., 2023; Wang et al., 2023; Zhou et al., 2023), we notice a research gap in the thorough examination of *prompting and score composition in the domain of NLG metrics*, especially for **open-source** generative LLMs.

The Eval4NLP 2023 shared tasks aims to fill this gap by disallowing participants to fine-tune models and by restricting model usage to a fixed list of LLMs (see Figure 1). Hence, participants may only vary how models are prompted, how scores are extracted, and how models are used in combination. To make the task more inclusive, we consider large and small(er) LLM’s in two separate tracks. This is different from shared tasks without model restriction, where the largest models often perform best, for example, the WMT metrics shared task (e.g. Freitag et al., 2022).

The goal of the shared task is to design evaluation metrics for MT and summarization, which we select as sub-tasks of NLG, while adhering to the model restrictions. Our contributions are the following:

- We design a novel, restricted evaluation setting that allows to focus on *prompting and*

ai/, <https://github.com/prompts-lab/Awesome-Prompt-Engineering>, <https://github.com/DukeLuo/awesome-awesome-prompts>, <https://github.com/snwfdhmp/awesome-gpt-prompt-engineering>, <https://github.com/dqxiiu/ICL-PaperList>, <https://github.com/EgoAlpha/prompt-in-context-learning>

score extraction in building evaluation metrics. This might aid inexpensive development of new metrics without fine-tuning or could benefit the selection of metric architectures with fine-tuning.

- We organized a CodaLab (Pavao et al., 2023) / Codabench (Xu et al., 2022) competition where participants could submit their system scores in a dev- and test-phase. The dev-phase has received 44 participant registrations, of which 9 teams have submitted contributions to the test-phase leaderboard and system papers. This paper summarizes their approaches and findings and presents their final ranking.
- We collect a novel dataset from Wikipedia articles created past the 15.07.2023 with the goal of minimizing the use of data that has been used to pre-train LLaMA2 (Touvron et al., 2023) released on 17.07.2023. This is because some of the allowed models are fine-tuned versions of LLaMA2.
- In line with the Eval4NLP 2021 shared task (Fomicheva et al., 2021), we consider the *explainability* of the designed metrics. The generative nature of LLMs allows to return natural language or formatted explanations of its output. While these explanations are not necessarily faithful, they also offer value if

they are plausible (Leiter et al., 2023a) or might support the generation process itself (Wei et al., 2022).

Our paper is structured into 8 sections. §2 gives an overview of how our shared task is related to other competitions. §3 describes the competition setup and §4 / §5 describe the datasets and annotation process for the test phase respectively. In §6, we highlight the approaches tested by the participants, especially those for the test set submissions. §7 presents the final scores of the participants on the test set and further analyses. Finally, §8 discusses future work and provides a conclusion.

2 Related Work

In this paragraph, we describe other work that is related to our shared task. In specific, we give a brief overview of evaluation metrics, highlight the recent development on metrics that are based on generative LLMs and describe related shared tasks.

NLG evaluation metrics The evaluation of NLG systems is necessary to compare them to other Systems and generally evaluate their applicability in intended scenarios. Manual/human evaluation is expensive, time consuming and often infeasible for larger datasets. Hence, automatic metrics are constructed. Many early metrics like BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) measure the lexical overlap between the generation and a human written reference. Metrics that use manually annotated references are called *reference-based*, while metrics that evaluate the generation quality based on the source text are called *reference-free* (in MT also Quality Estimation, QE). The early metrics that are based on lexical overlap have limitations in their ability to capture semantics of generated text (e.g. Reiter, 2018). For example, a generation might not be graded as good if it uses paraphrases of the reference texts. Newer metrics are usually based on language models that are able to embed the meanings of tokens (e.g. Zhang et al., 2020; Zhao et al., 2019; Sellam et al., 2020; Rei et al., 2020). These metrics achieve strong correlations to human judgments of generation quality (e.g. Freitag et al., 2022). Embedding based metrics have also enabled reference-free evaluation. This has the added benefit of no longer needing human reference generations and therefore enables further use cases, such as checking generation quality on the fly (e.g. Zerva et al., 2022), training with metrics as supervision signal (e.g. Wu et al., 2018) and

using metrics during decoding (Fernandes et al., 2022). However, the usage of black-box systems in the evaluation process also poses new challenges. For example, it can be difficult to understand why metrics exhibit certain behavior, they might lack robustness and fail in unexpected scenarios and they might show social biases (e.g. Leiter et al., 2023a). Surveys on NLG metrics are presented by (e.g. Celikyilmaz et al., 2021; Sai et al., 2022).

Generation-based evaluation metrics Related work includes other generation-based metrics. Beginning with PRISM (Thompson and Post, 2020) and BARTScore (Yuan et al., 2021), generation-based metrics have shown strong performance. These two metrics use the **generation probability** of paraphrases or translations as metric scores. Newer work that follows the same principle with more high-performing LLMs has shown improved scores (e.g. Fu et al., 2023). Another branch of generation-based metrics has originated with recent GPT models and shows that models can directly perform the task of grading machine generated text from in-context task descriptions (e.g. Kocmi and Federmann, 2023; Chiang and Lee, 2023; Fu et al., 2023; Xu et al., 2023b; Yang et al., 2023; Lu et al., 2023). We will refer to these metrics as **output-based**. Here, the rating is usually returned directly in the generated output text or constructed from it. Another branch of these models employs generative LLMs for ranking between better and worse generations (Zheng et al., 2023; Shen et al., 2023; Ji et al., 2023).

This recent surge of approaches has motivated our shared task. During the runtime of the shared task, other state-of-the-art approaches have been published (e.g. Fernandes et al., 2023). The systems submitted to our competition are different from most generation-based metrics in thoroughly exploring the usage of fixed recent open-source LLMs since ChatGPT without the usage of fine-tuning.

Evaluation Shared Tasks Our shared task is also related to other shared tasks that consider the evaluation of evaluation metrics for NLG, especially for MT and summarization. For MT, the established WMT workshop comprises multiple shared tasks on MT evaluation. Especially, the *WMT metrics shared task* (e.g. Mathur et al., 2020; Freitag et al., 2021b, 2022) and the *WMT shared task on quality estimation* (e.g. Specia et al., 2020, 2021; Zerva et al., 2022) are related to ours. The main track of



Figure 2: Schematic overview of possible approaches to compute scores from a generative LLM. Zero-shot approaches do not present examples in the prompt, while few-shot approaches present them. Chain-of-thought (Wei et al., 2022) approaches trigger the LLM to generate an explanation of its process before returning the final score. Fine-grained approaches, e.g. Fernandes et al. (2023), first construct a detailed error analysis and then construct a final score from them. Translation probability approaches, e.g. Fu et al. (2023), use the probability of generating a paraphrase as a translation. In a majority vote approach the results from multiple prompts could be combined. Self-refinement approaches could trigger a model multiple times to refine its output.

the *WMT metrics shared task* considers the system- and segment-level evaluation quality of MT metrics — that is, how well can metrics reflect the quality of whole MT systems or single segment translations. Recent years also put a focus on evaluating the robustness of metrics towards certain linguistic phenomena. The main track of the *WMT metrics shared task* consists of a reference-based evaluation, i.e., metrics compare the machine translation to human-written reference translations. Recent editions also contain a track for reference-free evaluation, where submitted metrics should directly compare the machine translation to its source text. Since 2021, the *WMT metrics shared task* has acquired its test data using the fine-grained MQM evaluation scheme (Lommel et al., 2014; Freitag et al., 2021a) that has been shown to be more accurate than crowd-sourced direct assessment annotations. The *WMT shared task on quality estimation* sets its main focus on the reference-free evaluation of machine translations. In recent years, their test sets are also annotated with MQM. Additionally, the quality estimation workshop has, for example,

conducted tasks on word-level error prediction and span-level error severity prediction.

Like the WMT QE shared task, our task is the reference-free evaluation of machine translations. The biggest difference of our shared task is that we fix the allowed models. That means, participants may only use models from a list we provide to them. Hence, participants have to focus on a thorough exploration of prompting and score extraction rather than fine-tuning and dataset creation. A second difference is that we include summarization as a subtask. As a third difference, our shared task has a subtrack to evaluate explanations that are created as a byproduct of scoring with generative LLM’s for plausibility. This last point offers parallels to the Eval4NLP 2021 shared task (Fomicheva et al., 2021) and its successor subtask at the WMT 2022 shared task (Zerva et al., 2022) on quality estimation. These tasks treated human word-level error annotations as explanations of translation quality and evaluated their correlations to manual annotations. In our subtask, we allow for any kind of explanation. Background information on explain-

ability for machine translation metrics can be found in [Leiter et al. \(2023a\)](#).

3 Shared Task Setup

As described in §1, the goal of our shared task is to leverage generative LLMs as (explainable) metrics for MT and summarization.³ Thereby, participants are not allowed to fine-tune their models and only certain models are allowed. Figure 1 shows the general setup of using generative LLMs as metrics, illustrated with an example from MT. The figure shows that final scores could be constructed from the generated model output or from other variables involved in the inference process. Specifically, recent work on prompting and metrics offer a wide range of possibilities to influence score construction even without fine-tuning. Some of them are shown in Figure 2.

LLM sizes We organize two tracks based on the model sizes. Models smaller than 25B parameters are considered as **small**, and models bigger than 25B parameters as **large**. Table 1 gives an overview of the allowed models. We mainly choose these models based on their good average performance on the Huggingface Open LLM Leaderboard.⁴ For Platypus2, Guanaco and WizardLM, we use 4-bit quantized versions with GPTQ ([Frantar et al., 2023](#)) to lower the system requirements to run them. Of these models, only the Guanaco model was explicitly fine-tuned with multilingual data. The models Wizard, Nous and Guanaco were allowed for use from the start of the competition, while the other 3 models were added to the list 20 days later. In another track, we explore the explanatory value of explanations created as a byproduct of the scoring process (see §7).

Phases Our shared task was conducted in two phases. First, we hosted a dev-phase on CodaLab⁵ ([Pavao et al., 2023](#)) from 07.08.23 to 30.09.23. In this phase, participants were developing their approaches and could already evaluate their scores on a leaderboard. While the standing in the dev-phase does not influence the ranking of the shared task, the phase aided the creation of a competitive atmosphere, acted as an advertisement for the competition and allowed us to gauge the number of

³We treat MT and summarization as separate tracks.

⁴https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard

⁵<https://codalab.lisn.upsaclay.fr/competitions/15072>

interested participants. The main part of the competition was the test-phase conducted from 26.09.23 to 01.10.23. Due to performance problems and unforeseen issues with extending the competition setup on CodaLab, the test phase was migrated to its successor Codabench⁶ ([Xu et al., 2022](#)). Submissions to the dev-phase and test-phase both had to contain at least a file with newline separated scores that grade each sample of our datasets. The test-phase additionally required to enter a team name, to indicate the track for each submission and to provide additional files with (1) a short system description, (2) newline separated prompts for each input, and (3) optionally newline separated explanations.

We describe the shared task datasets in §4.

4 Datasets

During the dev-phase of our shared task, we provided participants with a train- and a dev-set. For the test-phase, we further added a test-set.

Train- & Dev-set Our train- and dev-sets are constructed from two datasets. For MT, we select the en-de and zh-en MQM partitions of the WMT 2022 metrics shared task ([Freitag et al., 2022](#)). For summarization, we select SummEval ([Fabbri et al., 2021](#)). We conduct our task in a reference-free setting, that is, we do not provide human written reference translations or summaries. Hence, we remove the references provided with WMT and SummEval. SummEval has separate scores for relevance, factuality, coherence and consistency for each sample. We construct a single score per example by averaging these separate scores. Further changes to the original datasets include the split into train- and dev-partitions as well as shuffling. In the dev phase participants could experiment with generalizable (prompting) approaches.

Test-set We collect a novel test set for the test-phase of our shared task. It consists of 3 language pairs for MT: en-de, en-es, en-zh and a summarization part. We only choose high-resource languages, as the LLaMA(2)-based models have seen limited multilingual data during their pre-training and fine-tuning. Hence, high-resource languages can indicate an upper bound of what these models can achieve without further fine-tuning. To reduce the possibility that our chosen LLMs were trained

⁶<https://www.codabench.org/competitions/1359/>

Mode	Release Date	Track
Platypus2-70B-Instruct-GPTQ ⁷ (Lee et al., 2023a)	11.08.23	Large
Guanaco-65B-GPTQ ⁸ (Detmers et al., 2023)	25.05.23	Large
WizardLM-13B-V1.1-GPTQ ⁹ (Xu et al., 2023a)	07.07.23	Small
Nous-Hermes-13b ¹⁰	03.06.23	Small
OpenOrca-Platypus2-13B ¹¹ (Lee et al., 2023b; Mukherjee et al., 2023)	11.08.23	Small
orca_mini_v3_7b ¹² (Mathur, 2023; Mukherjee et al., 2023)	07.08.23	Small

Table 1: Generative LLMs whose usage was allowed in the Eval4NLP 2023 shared task.

on parts of the test set, we gather Wikipedia articles created after 15.07.23 as source texts.¹³

Figure 3 shows the score distributions of our datasets. We can see that all language pairs exhibit a pattern of centering around values divisible by 5. This makes sense, as MQM weighs major errors with 5 points. Also, in *en-es*, samples have generally received a higher score; i.e., fewer major errors were annotated. Finally, our summarization dataset, which uses a combined annotation scheme (see §5) does not show this pattern.

5 Annotation

In this section, we describe the annotation process of our dataset. For MT annotation, we hire one annotator per language pair: one Master student who speaks Spanish as mother tongue with English certifications, one NLP Bachelor student, who is a native English speaker that lives in Germany since many years, and one data and discourse studies Master student, who is a native Chinese speaker who uses English on a daily basis. For summarization annotation, we hire one NLP Bachelor student as well as a data and discourse studies Master student with a prior master in linguistics. Both annotators annotated the same data. All annotators demonstrated their suitability for the role in initial test rounds with further applicants. The distribution of our final MT dataset is shown in Table 3. The total annotation costs were ca. 5000€.

We use Google’s Anthea¹⁴ as annotation tool, because of its support for MQM annotations (Lommel et al., 2014; Freitag et al., 2021a). As we mostly annotate single sentences for MT, we modify Anthea to provide context via a Wikipedia URL that can be consulted if annotators are unsure about a translation. For summarization, annotations were

conducted in a modified version of Anthea with a new template (we show a screenshot of the UI in Appendix C).

For both data sets, we perform fine-grained annotations. In MT this has been shown to yield more reliable human annotations than other annotation schemes (Freitag et al., 2021a). Also, the fine-grained annotations could be used later-on to verify automatically generated explanations. As we only received 2 submissions for the explainability track, we do not consider apply this in this report.

MT We construct the MT dataset from random source sentences with a minimum length of 110 characters, as tokenized by the NLTK sentence tokenizer¹⁵. In a few cases, multiple sentences are concatenated due to missing spaces between dots. We obtain machine translations with 4 different translation models (see Table 2). Further, we use MQM as annotation scheme and conducted the annotation process in multiple batches to allow for corrections in subsequent batches. The batch sizes varied between 200 and 600 samples. For the first batch, we changed parts of the process during the annotation.

Specifically, we had accidentally chosen an incorrect tokenization for the first few samples of the first batch.¹⁶ This may have led to coarser annotation and to ignoring some punctuation issues. We still use these samples, as punctuation errors only have a very small weight in MQM and a coarser annotation does not change the severity assigned to errors. Hence, we assume that the impact on the MQM scores is minimal. Another change between annotation versions is that the first batch contains

¹⁵<https://www.nltk.org/api/nltk.tokenize.html>

¹⁶For the evaluation phase, we keep the annotations of the first batch, as small issues in source sentences should not invalidate the possibility of creating good translations; instead, we remove every sentence from the final dataset that has at least one major source error. We do this as major source errors might cause ambiguity in the annotation process. For example, if the source is unreadable, it is unclear which quality should be expected from the translation.

¹³Limitations of this approach are discussed in §8

¹⁴<https://github.com/google-research/google-research/tree/master/anthea>




Figure 3: Score distributions of our datasets. The annotation process is described in §5.

unordered sentences, while in the second version, all translations of a single source follow each other (in a random order). This has majorly improved the annotation speed as annotators do not need to reread the source sentences anymore. Further, the annotators commented on difficult source texts in the first batch. Therefore, in the following batches, we pre-filter the Wikipedia source articles by their quality classes¹⁷ and keep only c-class and better articles. Furthermore, we employ languagetool¹⁸ to filter for the grammatical correctness of the source sentences.

To verify the quality of the dataset, members of our team who are native speakers of the respective

¹⁷https://en.wikipedia.org/wiki/Wikipedia:Content_assessment

¹⁸<https://languagetool.org/de>

target languages have annotated small subsets of 30-50 samples of the datasets. Table 4 shows the agreement on these subsets. For en-es, either the MT models were more performant, the annotator might have been missing some errors or annotating them less strictly, as suggested by Figure 3.

Summarization We select random sections from Wikipedia that have a length of 150 to 800 tokens as measured by the tokenizer of *bart-large-cnn*. The summarization models we use are listed in Table 2. To create a dataset that offers as much explanatory value on the summary quality as possible, we perform a fine-grained evaluation inspired by MQM. However, we cannot simply reuse all criteria of the MQM commonly used in MT, as instead of fulfilling the criteria of adequacy, sum-

MT Models ¹⁹	Summarization Models
mbart50_en2m (Fan et al., 2021)	sshleifer/distilbart-cnn-12-6 ²⁰ (Shleifer and Rush, 2020)
mbart50_m2m (Fan et al., 2021)	facebook/bart-large-cnn ²¹ (Lewis et al., 2020)
m2m_100_418M (Tang et al., 2021)	google/bigbird-pegasus-large-bigpatent ²² (Zaheer et al., 2020)
m2m_100_1.2B (Tang et al., 2021)	facebook/bart-large-xsum ²³ (Lewis et al., 2020)
	mT5_multilingual_XLSum ²⁴ (Hasan et al., 2021)

Table 2: An overview of the translation and summarization models we have used to create our datasets.

Type	Train	Dev	Test
en-de	11046	7364	1425
en-es	-	-	1834
en-zh	-	-	1161 (1297)
zh-en	15750	10500	-
summarization	320	1280	671 (825)

Table 3: Number of samples in our datasets. In the case of the brackets, we filtered out potentially malformed examples after the test phase was conducted.

Type	Agreement
en-de	0.458
en-es	0.239
en-zh	0.480
summarization	0.625

Table 4: Kendall agreement between annotators. For MT, the agreement was calculated on 30-50 samples. For summarization, it was calculated on 373 examples.

maries need to capture the most relevant facts (relevance) and only represent correct facts (factuality). Specifically, we orient ourselves on the quality criteria for summaries by Dang (2005); Fabbri et al. (2021): relevance, factuality, and readability, where readability includes the property of coherence and fluency. We note that readability is already covered to a large degree by the MT MQM annotation guidelines. We change them by removing adequacy and adding coherence. Coherence has the following sub-categories: referential clarity, redundancy, structure, and meaning. The meaning category refers to cases where the summary changes the meaning of the source text without hallucinating, e.g., by concatenating facts in the wrong order.

One common approach to determine the relevance and factuality of summaries is the *pyramid approach* (Nenkova and Passonneau, 2004). Here, small atomic facts of many human written references are collected and ordered in a pyramid, based

on their occurrence count. Instead we introduce a more resource efficient approach, where we use a reference-free method for annotating the summaries’ relevance and factuality. Inspired by Liu et al. (2023c), who manually split the source text into atomic facts, we leverage the NLTK sentence tokenizer to split the source text into enumerated sentences. In some cases, sentences were not split correctly. In sentences of the final test set, we have corrected them manually. We treat each sentence as a single fact.²⁵ Next, we annotate the relevance of each of these facts, i.e., how likely would the annotator use the fact in a given sentence if they should write a summary themselves. Then, we annotate which source sentence is reflected in which part of the summary. By doing so, we can weigh the relevance of each fact that appears in the summary. Finally, we annotate each fact not represented in the original source text as a hallucination. Based on these components, we build a heuristic that is negative for bad summaries and positive for good summaries. The equation is shown in Figure 4. α , β and γ can be chosen to determine the influence of each sub-score for relevance, hallucinations and readability, respectively. There are many design choices regarding the weighting of each component and different normalization approaches. We find that these generally only have a small impact on the final ranking of our shared task (see Appendix A). Longer summaries can contain more facts and would hence receive higher scores in this heuristic. We address this issue by generating summaries of similar lengths using max token settings. The example in Figure 5 shows this annotation process.

Like with MT, we annotated in several batches. After the first batch, as for MT, we took measures to improve the source quality and ordered the sources to allow for faster annotations. After a check on the annotation quality, some misunderstandings of the

²⁵ Splitting each sentence into more granular facts, might further improve the fine-grained score composition but would require more effort in determining distinct facts.

$$\sum_{i \in \text{Facts in Summary}} \alpha * \text{relevance}(i) + \beta * \frac{|\text{Hallucinated Characters}|}{|\text{Characters in the summary}|} + \gamma * \text{MQM} \quad (1)$$

Figure 4: A heuristic for fine-grained reference-free evaluation of summaries. We set $\alpha = 3$, $\beta = 5$ and $\gamma = 1$.

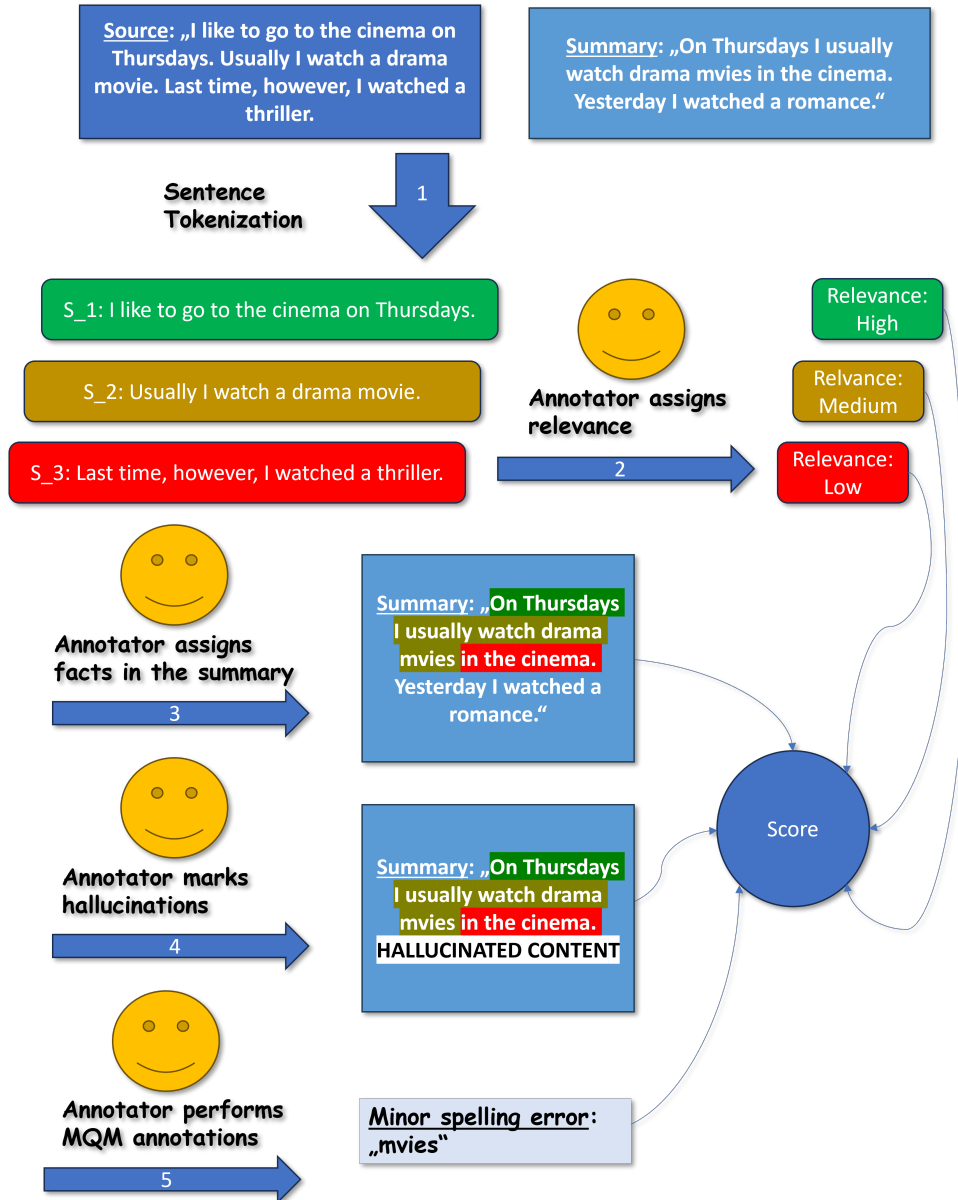


Figure 5: An example of the summarization annotation process.

annotation classes were uncovered and discussed. In the final evaluation, we drop all examples labeled before this discussion, such that we keep a total of 671 samples. Further, one annotator showed a larger annotation speed and a more consistent understanding of the task. In the test set, we use the annotations of this annotator.

Table 4 shows the agreement between the annotators. It is high for relevance and factuality

annotations and lower for the MQM part.

Evaluation

Following earlier WMT tasks on segment-level evaluation, we compute Kendall’s tau correlation (KENDALL, 1945) to compare the system generated scores to human scores. We further report the Spearman and Pearson correlations.²⁶ Future work

²⁶For these evaluations of correlations, we use the implementations of the python scipy library: <https://scipy.org/>

could explore if the usage of other and possibly more suited variants of Kendall, as suggested by [Deutsch et al. \(2023\)](#), might affect the rankings of our competition.

6 Shared Task Approaches

The test phase of our shared task received submissions from 12 different teams, 9 of which submitted system papers. Here, we summarize the approaches of these 9 systems and announce their final standings. Table 5 gives an overview of the participating teams and of the tracks they are participating in.²⁷ This table can be used as a mapping for the scores reported in §7.

We divide the approaches taken by the participants into *probability-based*, *output-based* and *agent-based*.²⁸ Besides their final approaches, the participants have explored a large number of possible variations. Afterwards, we introduce the baseline approaches, we compare the participants with.

Probability-based Probability-based approaches calculate how likely a paraphrase or translation of an input is generated with an LLM. Probability based approaches are explored by [Zhang et al. \(2023\)](#) and [Pradhan and Todi \(2023\)](#). [Zhang et al. \(2023\)](#) define 10 different prompts to translate a source sentence with an LLM. They combine this approach with demonstrating samples in the input prompt selected by (among others) SBERT ([Reimers and Gurevych, 2019](#)). Further, they use ensembles to recombine the scores of multiple prompts and models. [Pradhan and Todi \(2023\)](#) use the probability-based approach with own prompts and prompts designed by the authors of GPTScore ([Fu et al., 2023](#)).

Output-based All submitted papers explore the direct usage of an LLM’s natural language output as score. [Zhang et al. \(2023\)](#) test the same sample selection and ensembling strategies described above with 4 different prompts in an output-based setting. [Larionov et al. \(2023\)](#) follow a similar approach to [Zhang et al. \(2023\)](#) and retrieve demonstration examples by finding similar examples with LABSE ([Feng et al., 2022](#)) embeddings in an output-based setting. [Pradhan and Todi \(2023\)](#)

try one approach in which they present a prompt that triggers the prediction of a single score and one approach that triggers the model to first rate summary qualities for consistency, coherence, fluency and relevancy. Then they aggregate these scores in 3 different ways. [Baswani et al. \(2023\)](#) quantify Orcamini themselves to run an even smaller model (which is close to violating the allowed settings of the shared task). They provide a detailed explanation to their model that triggers it to produce fine-grained scores and a combined score in the same output. [Kim et al. \(2023\)](#) choose rating guidelines from related work — concretely, the human guidelines (HG) for SummEval, the machine guidelines for G-Eval ([Liu et al., 2023b](#)) and evaluation steps generated by GPT4 ([OpenAI, 2023](#)). They test various adaptations to this prompt, explore the usage of examples in the prompt and the usage of coarse-grained vs. fine-grained and aggregated scores. On the test set, they add a shortcut for very bad summarizations and employ bucketing for their scores. [Akkasi et al. \(2023\)](#) explore evaluating 6 different criteria over all model combinations. [Kotonya et al. \(2023\)](#) explore 8 prompt types: 3 base prompts and their extensions with chain-of-thought ([Wei et al., 2022](#)), zero-shot and few-shot settings. [Mahmoudi \(2023\)](#) explores various zero-shot and few-shot settings with Orcamini. Finally, [Mahmoudi \(2023\)](#); [Baswani et al. \(2023\)](#) generate explanations as an additional request to their model.

Agent-based While they also use an output-based setup, we place [Lu and Yu-Ting \(2023\)](#) in a separate group. They define 4 characters that should be played by a model and a list of 10 properties. For example they define “Internet Troll” as a critical character or “Teacher” as more knowledgeable character, with the intention that different viewpoints can help to judge generation quality better. Then, they evaluate the combined 40 settings and use XGBoost ([Chen and Guestrin, 2016](#)) to combine their scores. While they did not add their top submissions to the final leaderboard they present their reasonably good final scores in their paper.

Baselines As baselines, we use the widely used metrics BERTScore (with XLMR-large embeddings) ([Zhang et al., 2020](#)), SBERT ([Reimers and Gurevych, 2019](#)) cosine-similarity (with XLMR-large embeddings), SUPERT ([Gao et al., 2020](#)), GEMBA ([Kocmi and Federmann, 2023](#)) and

²⁷While the first and last authors of [Larionov et al. \(2023\)](#) are members of the NLLG group, we did not share any internal details that would have given them an advantage. They developed their approach independently.

²⁸View §2 for the distinction of probability-based and output-based.

Team	Authors	Tracks
Pradhan/Todi	(Pradhan and Todi, 2023)	S, SU
Kotonya et. al.	(Kotonya et al., 2023)	S, SU
DSBA	(Kim et al., 2023)	S, L, SU
HIT-MI&T Lab	(Zhang et al., 2023)	S, MT
IUST_NLP_Lab	(Mahmoudi, 2023)	S, SU, E
LTRC	(Baswani et al., 2023)	S, MT, SU, E
NLLG	(Larionov et al., 2023)	L, MT, SU
TaiwanSenior	(Lu and Yu-Ting, 2023)	S, MT
iML	(Akkasi et al., 2023)	S, L, SU

Table 5: Overview of shared task submissions. The letters are abbreviations for the following tracks: S(small model track), L(arge model track), M(achine)T(ranslation track), SU(mmarization track), E(xplainability track).

Comet-Kiwi-XXL (Rei et al., 2023). Further, we include one baseline for every allowed model that uses the DA score prompt of GEMBA (Kocmi and Federmann, 2023) (with a slight modification for summarization). The models are further specified in Appendix D.

7 Results and Analysis

In this section, we first report statistics of the shared task. Then we will present and discuss the final system ranking. Note that we include submissions of participants on the test-set-leaderboard that did not submit a system paper. However, we do not describe their approaches in §5. Lastly, we will discuss the implications of these results on the development of generation-based metrics.

Statistics The dev-phase on CodaLab has received 44 registrations, 13 of which have submitted their scores. In total, there have been 1048 submissions on the dev-set suggesting that some participants might have optimized their method on the dev-set. Especially, one participant submitted 417 submissions on the dev set. The test-phase on Codabench has received 21 registrations and 248 submissions from 11 participants. We have restricted the number of allowed submissions per day to 10. Allowing a higher number would enable participants to optimize their approaches on the test-set too much, such that the results would not reflect the generalization capability anymore. On the other hand, we wanted to give participants the option to try out multiple approaches they designed. Further, Codabench would sometimes fail to compute scores and still deduct one submission. Hence, 10 submissions per day allows us to continue in these cases. Two participants have used

up a contingent of ≈ 50 submissions. Of the 11 test-phase participants, 9 have submitted a system paper. The first authors are from China, India (2), Korea, Taiwan, Canada, Iran, Germany and the United Kingdoms. That means, many authors are from developing countries. Also, many authors are students. Hence, their resource availability was limited, leading many of them to opting for smaller models.

Correlation with humans Here, we present the results that the participants achieve on the test sets. A mapping between team names and authors can be found in Table 5. Table 6 shows the final ranking of the *small* MT subtask. Compared to the other participants, Zhang et al. (2023) leads by a large margin on all correlation measures. Even significantly outperforming the recent COMET-kiwi-XXL and only being matched by GEMBA with GPT-4. This is surprising, as the scores they report on the dev-set are not this strong. However, also on the dev-set they beat the large model baselines that use the 6 models we allow in the shared task. The test-set approach that Zhang et al. (2023) report in their paper builds on ensembling probability-based scores from prompts to OpenOrca-Platypus. These prompts contain 3 up to the maximum number of possible example demonstrations. Future work should explore whether their approach can uphold its strong performance across other datasets and settings. The ranking is then followed by various baseline models and team LTRC.

Table 7 shows the final ranking of the *large* MT subtask. For this subtask, the baselines have not been beaten. Table 8 shows the final ranking of the *small* summarization subtask. Kim et al. (2023) and Akkasi et al. (2023) lead this track. Both use

Team	Kendall			Pearson			Spearman		
	de	zh	es	de	zh	es	de	zh	es
<i>baselineGEMBA</i>	0.492	0.384	0.409	0.506	0.356	0.251	0.625	0.496	0.512
HIT-MI&T Lab	0.491	0.375	0.417	0.655	0.528	0.453	0.656	0.511	0.553
<i>baselineCometKiwiXXL</i>	0.421	0.345	0.288	0.562	0.443	0.331	0.583	0.484	0.403
<i>baselineBertscore</i>	0.239	0.174	0.221	0.344	0.236	0.179	0.344	0.252	0.312
<i>baselineSBERT</i>	0.209	0.167	0.226	0.246	0.210	0.081	0.304	0.242	0.320
LTRC	0.194	0.144	0.112	0.232	0.133	0.031	0.233	0.173	0.132
<i>baselineNous</i>	0.189	0.011	0.112	0.183	0.044	0.045	0.230	0.013	0.136
<i>baselineOrcaPlaty</i>	0.189	0.011	0.112	0.183	0.044	0.045	0.230	0.013	0.136
seanstilwell	0.120	NaN	NaN	0.164	NaN	NaN	0.152	NaN	NaN
<i>baselineWizard</i>	0.101	0.065	0.079	0.047	0.057	0.026	0.121	0.077	0.093
<i>baselineOrcaMini</i>	0.073	0.188	0.065	0.030	0.102	0.009	0.088	0.225	0.077
TaiwanSenior	0.041	NaN	NaN	-0.037	NaN	NaN	0.051	NaN	NaN

Table 6: Results of the *small* model track for MT. For our main metric Kendall, we write results that are significantly better than the following, with $p \leq 0.05$, as measured by a permute-both significance test (Deutsch et al., 2021). GEMBA was not included in the significance test. Teams with paper submissions are bolded.

Team	Kendall			Pearson			Spearman		
	de	zh	es	de	zh	es	de	zh	es
<i>baselinePlaty_large</i>	0.362	0.293	0.264	0.312	0.270	0.129	0.445	0.364	0.320
<i>baselineGuanaco_large</i>	0.350	0.219	0.241	0.344	0.176	0.125	0.445	0.273	0.300
NLLG	0.245	0.139	0.179	0.257	0.196	0.155	0.335	0.190	0.238
kaiwalya_large	0.174	0.113	0.125	0.161	0.141	0.052	0.209	0.138	0.147

Table 7: Results of the *large* model track for MT. For our main metric Kendall, we write results that are significantly better than the following, with $p \leq 0.05$, as measured by a permute-both significance test (Deutsch et al., 2021). Teams with paper submissions are bolded.

Team	kd	ps	sp
DSBA	0.633	0.783	0.782
iML	0.615	0.763	0.772
<i>baselineBertscore</i>	0.578	0.771	0.765
IUST_NLP_Lab	0.573	0.722	0.722
<i>baselineOrcaMini</i>	0.560	0.681	0.706
<i>baselineSupertMpnet2</i>	0.554	0.736	0.747
<i>baselineOrcaPlaty</i>	0.552	0.666	0.674
<i>baselineNous</i>	0.552	0.666	0.674
Kotonya et. al.	0.546	0.680	0.682
LTRC	0.531	0.691	0.679
<i>baselineSupertFull</i>	0.516	0.686	0.706
<i>baselineSupert5</i>	0.492	0.654	0.678
<i>baselineSBERT</i>	0.465	0.625	0.645
Pradhan/Todi	0.436	0.032	0.610
<i>baselineWizard</i>	0.411	0.534	0.536
Haaland	0.221	0.514	0.280

Table 8: Results of the *small* model track for summarization. *kd* stands for Kendall, *ps* stands for Pearson and *sp* stands for Spearman. For our main metric Kendall, we write results that are significantly better than the following, with ≤ 0.05 , as measured by a permute-both significance test (Deutsch et al., 2021). Teams with paper submissions are bolded.

Team	kd	ps	sp
iML	0.612	0.738	0.768
DSBA	0.603	0.756	0.766
<i>baselinePlaty_large</i>	0.600	0.740	0.753
NLLG	0.471	0.643	0.638
<i>baselineGuanaco_large</i>	0.402	0.492	0.504

Table 9: Results of the *large* model track for summarization. *kd* stands for Kendall, *ps* stands for Pearson and *sp* stands for Spearman. For our main metric Kendall, we write results that are significantly better than the following, with ≤ 0.05 , as measured by a permute-both significance test (Deutsch et al., 2021). Teams with paper submissions are bolded.

carefully crafted prompts to achieve their results.

Table 9 shows the final ranking of the *large* summarization subtask. Here, Akkasi et al. (2023) is the winning team. Interestingly, for MT and summarization, the small models have beaten the large models. One potential reason might be that the large models take much longer to run and therefore they could not be examined with the same care. Further, it is interesting that the OrcaMini baseline and Mahmoudi (2023) beats many other models despite its parameter count being the lowest of the allowed models’. Generally, many teams opted for the usage of small models. Some teams only use the OrcaMini model, due to resource constraints. This highlights the importance of the inclusiveness of research in the metrics domain. We show a further analysis of the impact of the summarization subcategories in Appendix B.

Performance The best performing approaches of the participants achieve a similar Kendall correlation as our team members when we were testing the inter-annotator agreement on a small subset of samples (see §3). This suggests that these approaches are already close to the performance of native speakers with little training with the annotation process (as compared to our main annotators with a strong language background and more annotation experience on the task). This is an intriguing finding and highlights the potential of current open source models with and without fine-tuning. Especially, as many prompting approaches, like tree-of-thoughts or self-refinement still remain to be explored. Further, it shows that for closed source models like ChatGPT or GPT4 similar opportunities may exist and lead to new state-of-the-art metrics. The results also show that comparably small hardware can already be enough to create strong new metrics.

Explainability Only 2 participants (Baswani et al., 2023; Mahmoudi, 2023) have submitted entries with complementary explanations to the Codabench leaderboard. Both directly prompted the model to give reasoning for the model’s decision. Thus, we perform the human experiment on explainability only on a small scale of 50 annotations for randomly selected samples of our summarization dataset. Two annotators of our team were presented with source, summary, MQM annotations (to help to identify problems), the scores of the participants and the explanations of the participants. They annotated which of two explanations they pre-

fer. One annotator preferred explanations of one system, let's call it A, in 27 cases and explanations of the other in 23 cases. The other annotator preferred system A in 24 cases and the other system in 26 cases. In these annotations the annotators agree in 56% of cases. These findings show that the annotators did not have a clear preference between the systems. Also, we notice that many explanations tend to be vague and return texts such as "The summary has a good coherence and fluency". In some cases, the explanations correctly describe problems. We show one example explanation of [Baswani et al. \(2023\)](#) in Table 10. Here, the explanation correctly captures the word repetition.

8 Conclusion

We discuss future work and then summarize the shared task in a conclusion.

Future Work We have considered high resource languages for the MT task. Future work could evaluate low-resource languages, especially once more generative LLMs are released that are trained across a wide range of languages. Also, if this shared task topic is repeated in the future, we might encourage and set rewards for pipeline-based solutions. In other words, currently most approaches of the shared task are based on single prompts or probability outputs; instead many interesting approaches like tree of thoughts ([Yao et al., 2023](#)) explore pipelines in which the output is generated iteratively or in parallel. Future work might also create larger or more diverse datasets for our evaluation scheme. Another point is that our current work only contains a small analysis of explainability that remained indecisive on the explanation quality between two participants. This could be extended in future work.

Conclusion This work describes the Eval4NLP 2023 shared task on *prompting LLMs as explainable metrics*. We have constructed a fine-grained dataset for MT and summarization evaluation, with a novel annotation scheme for the latter. Further, we have organized a competition following the novel restriction to specify allowed models and disallow fine-tuning in a MT and summarization evaluation setting. By running a small and a large model track, we have enabled participation for participants with fewer resources, leading to an inclusive shared task setting.

The top scores of the participants highlight a

number of interesting findings that we summarize here:

- **Small Models:** The results on the test set show that the best solutions built on small models outperform those that are built on larger models. This is contradicting usual patterns and an interesting finding for metric efficiency.
- **Probability-based vs. Output based:** The MT ranking is lead by a probability-based method, while the summarization ranking is lead by two prompt-based methods. For MT, this could be caused by the models' understanding of other languages being smaller than its capability of translation, therefore favoring paraphrasation based methods.
- **Simplicity helps:** Many baseline systems achieved high ranks, despite using a simple prompting approach. Participants often report that demonstrating examples reduced their performance. Hence, lean metrics are easier to design and can still be very powerful. The best ranked systems, however, explore more intricate prompts.

The contributions of our participants highlight once more how current LLMs can achieve state-of-the-art performance, even without any task-specific fine-tuning.

Acknowledgements

We thank our participants for the active contribution and discussion. Further, we thank our annotators for their effort in creating our test sets. Christoph Leiter is financed by the BMBF project "Metrics4NLG". Steffen Eger is financed by DFG Heisenberg grant EG 375/5-1.

Limitations

One potential limitation of our work lies in the usage of data from Wikipedia after 15.07. While the selected articles were indeed selected after 15.07, texts could still be copied from other places, some texts were automatically translated from other languages were an entry existed and some texts might even be generated. Another issue of our work lies in the comparably small dataset with low agreements for the small test conducted on the Spanish annotations. Due to time restrictions, we could not do further evaluations. Still we believe that

Source	Summary	Explanation
<p>In 1918, Wimble built a small ship at Hastings with help from a friend and sailed to the West Indies to seek his fortune after his family faced financial hardship. In 1922, he acquired land in the Bahamas which enabled him to begin trading with the English colonies in mainland North America. He also acquired land in North Carolina, which was formally granted to him by George Burrington's council on August 4, 1723. Wimble later moved to Boston, Massachusetts where he married Rebecca Waters, the daughter of a prominent local, on March 26, 1724. Their first son, James, was born on December 20, 1724. He owned land in the South End which he presumably operated as a distillery. While in Boston, he continued his business of moving trade goods between North Carolina and various British trade posts in the West Indies. This business enabled him to increase his land holdings in North Carolina and purchase a brigantine, which he named "Rebecca" after his wife. In 1932, Wimble lost his ship and all of its cargo to a hurricane after being forced by Governor Woodes Rogers of the Bahamas to use his ship to protect vessels and salt ponds in Rum Cay. Wimble was forced to sell a portion of his belongings, land, and slaves to cover the loss and began the process of trying to collect damages from Woodes Rogers' commandeering of his ship.²⁹</p>	<p>Wimble is the subject of this patent application. Wimble is the son of Wimble. Wimble is the son of Wimble. Wimble is the son of Wimble. Wimble is the son of Wimble. Wimble is the son of Wimble. Wimble is the son of Wimble. Wimble is the son of Wimble. Wimble is the son of Wimble. Wimble is the son of Wimble.</p>	<p>Explanation: The summary is not relevant to the source text, as it contains no information about Wimble's life, family, or business ventures. Instead, it repeats the name "Wimble" multiple times, which is not a summary of the source text.</p>

Table 10: Explanation generated with the approach by Baswani et al. (2023)

. It correctly identifies the issue of the word *Wimble* repeating often.

our annotators were capable in their languages and thorough with their analysis of the samples. As another limitation, pre-filtering with language tool and later on sorting out severe source errors, might miss out on more subtle errors causing problems in the test set.

References

- Abbas Akkasi, Kathleen C. Fraser, and Majid Komeili. 2023. Reference-free summarization evaluation with large language models. In *Proceedings of the 4th Workshop on Evaluation and Comparison for NLP systems*.
- Pavan Baswani, Ananya Mukherjee, and Manish Shrivastava. 2023. Ltrc_iith’s 2023 submission for prompting large language models as explainable metrics task. In *Proceedings of the 4th Workshop on Evaluation and Comparison for NLP systems*.
- Jonas Belouadi, Anne Lauscher, and Steffen Eger. 2023. [Automatikz: Text-guided synthesis of scientific vector graphics with tikz](#).
- Asli Celikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2021. [Evaluation of text generation: A survey](#).
- Tianqi Chen and Carlos Guestrin. 2016. [XGBoost](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM.
- 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.
- Hoa Trang Dang. 2005. [Overview of duc 2005](#).
- Tirth Dave, Sai Anirudh Athaluri, and Satyam Singh. 2023. [Chatgpt in medicine: an overview of its applications, advantages, limitations, future prospects, and ethical considerations](#). *Frontiers in Artificial Intelligence*, 6.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#).
- Daniel Deutsch, Rotem Dror, and Dan Roth. 2021. [A statistical analysis of summarization evaluation metrics using resampling methods](#). *Transactions of the Association for Computational Linguistics*, 9:1132–1146.
- Daniel Deutsch, George Foster, and Markus Freitag. 2023. [Ties matter: Modifying kendall’s tau for modern metric meta-evaluation](#).
- Steffen Eger, Christoph Leiter, Jonas Belouadi, Ran Zhang, Aida Kostikova, Daniil Larionov, Yanran Chen, and Vivian Fresen. 2023. [Nllg quarterly arxiv report 06/23: What are the most influential current ai papers?](#) *ArXiv*, abs/2308.04889.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. [SummEval: Re-evaluating summarization evaluation](#). *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *Journal of Machine Learning Research*, 22(107):1–48.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Ariavzhagan, and Wei Wang. 2022. [Language-agnostic BERT sentence embedding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.
- Patrick Fernandes, Daniel Deutsch, Mara Finkelstein, Parker Riley, André F. T. Martins, Graham Neubig, Ankush Garg, Jonathan H. Clark, Markus Freitag, and Orhan Firat. 2023. [The devil is in the errors: Leveraging large language models for fine-grained machine translation evaluation](#).
- Patrick Fernandes, António Farinhas, Ricardo Rei, José G. C. de Souza, Perez Ogayo, Graham Neubig, and Andre Martins. 2022. [Quality-aware decoding for neural machine translation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1396–1412, Seattle, United States. Association for Computational Linguistics.
- Marina Fomicheva, Piyawat Lertvittayakumjorn, Wei Zhao, Steffen Eger, and Yang Gao. 2021. [The Eval4NLP shared task on explainable quality estimation: Overview and results](#). In *Proceedings of the 2nd Workshop on Evaluation and Comparison of NLP Systems*, pages 165–178, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Elias Frantar, Saleh Ashkboos, Torsten Hoeffler, and Dan Alistarh. 2023. [Gptq: Accurate post-training quantization for generative pre-trained transformers](#).
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021a. [Experts, errors, and context: A large-scale study of human evaluation for machine translation](#). *Transactions of the Association for Computational Linguistics*, 9:1460–1474.

- 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.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, George Foster, Alon Lavie, and Ondřej Bojar. 2021b. [Results of the WMT21 metrics shared task: Evaluating metrics with expert-based human evaluations on TED and news domain](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 733–774, Online. Association for Computational Linguistics.
- Simon Frieder, Luca Pinchetti, Alexis Chevalier, Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiewicz, Philipp Christian Petersen, and Julius Berner. 2023. [Mathematical capabilities of chatgpt](#).
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. [Gptscore: Evaluate as you desire](#).
- Yang Gao, Wei Zhao, and Steffen Eger. 2020. [SUPERT: Towards new frontiers in unsupervised evaluation metrics for multi-document summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1347–1354, Online. Association for Computational Linguistics.
- Mohanad Halaweh. 2023. [Chatgpt in education: Strategies for responsible implementation](#). *Contemporary Educational Technology*.
- Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, and Rifat Shahriyar. 2021. [XLsum: Large-scale multilingual abstractive summarization for 44 languages](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4693–4703, Online. Association for Computational Linguistics.
- Yunjie Ji, Yan Gong, Yiping Peng, Chao Ni, Peiyan Sun, Dongyu Pan, Baochang Ma, and Xiangang Li. 2023. [Exploring chatgpt’s ability to rank content: A preliminary study on consistency with human preferences](#).
- M. G. KENDALL. 1945. [THE TREATMENT OF TIES IN RANKING PROBLEMS](#). *Biometrika*, 33(3):239–251.
- JoongHoon Kim, Sangmin Lee, Seung Hun, Saeran Park, Jiyeon Lee, Kiyoon Jeong, and Pilsung Kang. 2023. [Which is better? exploring prompting strategy for llm-based metrics](#). In *Proceedings of the 4th Workshop on Evaluation and Comparison for NLP systems*.
- 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, Saran Krishnasamy, Joel R. Tetreault, and Alejandro Jaimes. 2023. [Little giants: Exploring the potential of small llms as evaluation metrics in summarization in the eval4nlp 2023 shared task](#). In *Proceedings of the 4th Workshop on Evaluation and Comparison for NLP systems*.
- Daniil Larionov, Vasiliy Viskov, George Kokush, Alexander Panchenko, and Steffen Eger. 2023. [Team nllg submission for eval4nlp 2023 shared task: Retrieval-augmented in-context learning for nlg evaluation](#). In *Proceedings of the 4th Workshop on Evaluation and Comparison for NLP systems*.
- Ariel N. Lee, Cole J. Hunter, and Nataniel Ruiz. 2023a. [Platypus: Quick, cheap, and powerful refinement of llms](#).
- Ariel N. Lee, Cole J. Hunter, Nataniel Ruiz, Bley Goodson, Wing Lian, Guan Wang, Eugene Pentland, Austin Cook, Chanvichet Vong, and "Teknium". 2023b. [Openorcaplatypus: Llama2-13b model instruct-tuned on filtered openorca v1 gpt-4 dataset and merged with divergent stem and logic dataset model](#). <https://huggingface.co/OpenOrca/OpenOrca-Platypus2-13B>.
- Christoph Leiter, Piyawat Lertvittayakumjorn, M. Fomicheva, Wei Zhao, Yang Gao, and Steffen Eger. 2023a. [Towards explainable evaluation metrics for machine translation](#). *ArXiv*, abs/2306.13041.
- Christoph Leiter, Ran Zhang, Yanran Chen, Jonas Belouadi, Daniil Larionov, Vivian Fresen, and Steffen Eger. 2023b. [Chatgpt: A meta-analysis after 2.5 months](#).
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Hanmeng Liu, Ruoxi Ning, Zhiyang Teng, Jian Liu, Qiji Zhou, and Yue Zhang. 2023a. [Evaluating the logical reasoning ability of chatgpt and gpt-4](#).
- Yang Liu, Dan Iter, Yichong Xu, Shuhang Wang, Ruochen Xu, and Chenguang Zhu. 2023b. [G-eval: Nlg evaluation using gpt-4 with better human alignment](#).

- Yixin Liu, Alex Fabbri, Pengfei Liu, Yilun Zhao, Linyong Nan, Ruilin Han, Simeng Han, Shafiq Joty, Chien-Sheng Wu, Caiming Xiong, and Dragomir Radev. 2023c. [Revisiting the gold standard: Grounding summarization evaluation with robust human evaluation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4140–4170, Toronto, Canada. Association for Computational Linguistics.
- Arle Lommel, Aljoscha Burchardt, and Hans Uszkorait. 2014. [Multidimensional quality metrics \(mqm\): A framework for declaring and describing translation quality metrics](#). *Tradumàtica: tecnologies de la traducció*, 0:455–463.
- Qingyu Lu, Baopu Qiu, Liang Ding, Kanjian Zhang, Tom Kocmi, and Dacheng Tao. 2023. [Error analysis prompting enables human-like translation evaluation in large language models: A case study on chatgpt](#).
- Yuan Lu and Lin Yu-Ting. 2023. [Characterised llms affect its evaluation of summary and translation](#). In *Proceedings of the 4th Workshop on Evaluation and Comparison for NLP systems*.
- Ghazaleh Mahmoudi. 2023. [Exploring prompting large language models as explainable metrics](#). In *Proceedings of the 4th Workshop on Evaluation and Comparison for NLP systems*.
- Nitika Mathur, Johnny Wei, Markus Freitag, Qingsong Ma, and Ondřej Bojar. 2020. [Results of the WMT20 metrics shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 688–725, Online. Association for Computational Linguistics.
- Pankaj Mathur. 2023. [orca_mini_v3_7b: An explain tuned llama2-7b model](#). https://https://huggingface.co/psmathur/orca_mini_v3_7b.
- Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. [Orca: Progressive learning from complex explanation traces of gpt-4](#).
- Ani Nenkova and Rebecca Passonneau. 2004. [Evaluating content selection in summarization: The pyramid method](#). In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 145–152, Boston, Massachusetts, USA. Association for Computational Linguistics.
- OpenAI. 2023. [Gpt-4 technical report](#). *ArXiv*, abs/2303.08774.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Adrien Pavao, Isabelle Guyon, Anne-Catherine Letournel, Dinh-Tuan Tran, Xavier Baro, Hugo Jair Escalante, Sergio Escalera, Tyler Thomas, and Zhen Xu. 2023. [Codalab competitions: An open source platform to organize scientific challenges](#). *Journal of Machine Learning Research*, 24(198):1–6.
- Abhishek Pradhan and Ketan Kumar Todi. 2023. [Understanding large language model based metrics for text summarization](#). In *Proceedings of the 4th Workshop on Evaluation and Comparison for NLP systems*.
- Ricardo Rei, Nuno M. Guerreiro, José Pombal, Daan van Stigt, Marcos Treviso, Luisa Coheur, José G. C. de Souza, and André F. T. Martins. 2023. [Scaling up cometkiwi: Unbabel-ist 2023 submission for the quality estimation shared task](#).
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, 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.
- Ehud Reiter. 2018. [A structured review of the validity of BLEU](#). *Computational Linguistics*, 44(3):393–401.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2023. [Code llama: Open foundation models for code](#).
- Ananya B. Sai, Akash Kumar Mohankumar, and Mitesh M. Khapra. 2022. [A survey of evaluation metrics used for nlg systems](#). *ACM Comput. Surv.*, 55(2).
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Chenhui Shen, Liying Cheng, Yang You, and Lidong Bing. 2023. [Are large language models good evaluators for abstractive summarization?](#)
- Sam Shleifer and Alexander M. Rush. 2020. [Pre-trained summarization distillation](#).

- Lucia Specia, Frédéric Blain, Marina Fomicheva, Erick Fonseca, Vishrav Chaudhary, Francisco Guzmán, and André F. T. Martins. 2020. [Findings of the WMT 2020 shared task on quality estimation](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 743–764, Online. Association for Computational Linguistics.
- Lucia Specia, Frédéric Blain, Marina Fomicheva, Chrysoula Zerva, Zhenhao Li, Vishrav Chaudhary, and André F. T. Martins. 2021. [Findings of the WMT 2021 shared task on quality estimation](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 684–725, Online. Association for Computational Linguistics.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2021. [Multilingual translation from denoising pre-training](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3450–3466, Online. Association for Computational Linguistics.
- Brian Thompson and Matt Post. 2020. [Automatic machine translation evaluation in many languages via zero-shot paraphrasing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 90–121, Online. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. 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*.
- Lijun Wu, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2018. [A study of reinforcement learning for neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3612–3621, Brussels, Belgium. Association for Computational Linguistics.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023a. [WizardLM: Empowering large language models to follow complex instructions](#).
- Wenda Xu, Danqing Wang, Liangming Pan, Zhenqiao Song, Markus Freitag, William Yang Wang, and Lei Li. 2023b. [Instructscore: Explainable text generation evaluation with finegrained feedback](#).
- Zhen Xu, Sergio Escalera, Adrien Pavão, Magali Richard, Wei-Wei Tu, Quanming Yao, Huan Zhao, and Isabelle Guyon. 2022. [Codabench: Flexible, easy-to-use, and reproducible meta-benchmark platform](#). *Patterns*, 3(7):100543.
- 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](#).
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#).
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [Bartscore: Evaluating generated text as text generation](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 27263–27277. Curran Associates, Inc.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. [Big bird: Transformers for longer sequences](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 17283–17297. Curran Associates, Inc.
- Chrysoula Zerva, Frédéric Blain, Ricardo Rei, Piyawat Lertvittayakumjorn, José G. C. de Souza, Steffen Eger, Diptesh Kanojia, Duarte Alves, Constantin Orăsan, Marina Fomicheva, André F. T. Martins, and Lucia Specia. 2022. [Findings of the WMT 2022 shared task on quality estimation](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 69–99, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Rui Zhang, Fuhai Song, Hui Huang, Jinghao Yuan, Muyun Yang, and Tiejun Zhao. 2023. [Hit-mi&lab’s submission to eval4nlp 2023 shared task](#). In *Proceedings of the 4th Workshop on Evaluation and Comparison for NLP systems*.

team	s_kd	s_ps	s_sp
DSBA	0.623	0.675	0.772
iML	0.602	0.642	0.757
IUST_NLP_Lab	0.566	0.678	0.712
bertscore	0.546	0.711	0.729
baselineOrcaMini	0.545	0.640	0.684
Kotonya et.al.	0.543	0.745	0.675
baselineOrcaPlaty	0.527	0.589	0.650
baselineNous	0.527	0.589	0.650
LTRC	0.522	0.655	0.666
baselineSBERT	0.438	0.524	0.611
Pradhan/Todi	0.424	0.030	0.594
baselineWizard	0.408	0.489	0.531
Haaland	0.265	0.732	0.332
cometXXL	-0.009	0.091	-0.015
baselineSUPERT	-0.028	-0.040	-0.040

Table 11: Results of the *small* model track for summarization with Equation 6.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. 2019. [MoverScore: Text generation evaluating with contextualized embeddings and earth mover distance](#). 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 563–578, Hong Kong, China. Association for Computational Linguistics.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric. P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#).

Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2023. [Large language models are human-level prompt engineers](#).

A Impact of the summarization heuristic

Here, we consider the impact of using alternative heuristics for summarization, by studying their effect on the ranking of summarization systems. The results for Equation 6 are shown in Table 11. The results for Equation 7 are shown in Table 12. We can see that the top rankings remain the same.

team	s_kd	s_ps	s_sp
DSBA	0.551	0.490	0.695
iML	0.533	0.454	0.687
ISUT_NLP_Lab	0.512	0.546	0.649
bertscore	0.497	0.569	0.663
baselineOrcaMini	0.485	0.517	0.612
Kotonya et.al.	0.480	0.690	0.604
LTRC	0.476	0.534	0.609
baselineOrcaPlaty	0.462	0.446	0.581
baselineNous	0.462	0.446	0.581
Pradhan/Todi	0.422	0.023	0.591
baselineSBERT	0.384	0.371	0.539
baselineWizard	0.361	0.381	0.478
Haaland	0.295	0.800	0.368
cometXXL	0.015	0.159	0.021
baselineSUPERT	0.003	-0.018	0.004

Table 12: Results of the *small* model track for summarization with Equation 7.

B Impact of subcategories

We also study the impact of subcategories on the final ranking of summarization. That means, we calculate the ranking with each of α , β , γ set to 1, while the others are 0. The results are shown in Tables 13, 14 and 15. Intriguingly, when only the MQM score is evaluated, the model by *Haaland* has the highest correlation. However, they did not submit a system description or a system paper. Further, all baselines in this setting perform relatively weak. The best baseline is comet, potentially as it has been trained on MQM scores. The results for relevance and hallucinations are rather unsurprising with one time *DSBA* being the winning team and the other time *iML*.

C Screenshot of the annotation interface

Figure 8 shows a screenshot of the Anthea annotation interface.

D Model Details

For SBert, we use embeddings of XLM-R to include multilinguality³⁰. For SUPERT we report the standard metric using bert-large-nli-stsb-mean-tokens³¹ with 5 and all source sentences as pseudo-references. Further, we upgrade SUPERT to use all-

³⁰<https://huggingface.co/sentence-transformers/stsb-xlm-r-multilingual>

³¹<https://huggingface.co/sentence-transformers/bert-large-nli-stsb-mean-tokens>

$$\sum_{i \in \text{Facts in Summary}} \alpha * \text{relevance}(i) + \beta * \frac{|\text{Hallucinated Characters}|}{|\text{Characters in the summary}|} + \gamma * \text{MQM} \quad (2)$$

Figure 6: A heuristic for fine-grained reference-free evaluation of summaries. Alternatively, we set $\alpha = 1$, $\beta = 1$ and $\gamma = 1$.

$$\frac{\sum_{i \in \text{Facts in Summary}} \alpha * \text{relevance}(i)}{|\text{Facts in Source}|} + \beta * \frac{|\text{Hallucinated Characters}|}{|\text{Characters in the summary}|} + \gamma * \text{MQM} \quad (3)$$

Figure 7: An alternative heuristic for fine-grained reference-free evaluation of summaries. We set $\alpha = 1$, $\beta = 1$ and $\gamma = 1$. Further, we divide the relevance part by the number of facts in the source as normalization.

Source ([en-EN])	Text ([en-EN])	Evaluations								
Context: https://en.wikipedia.org/wiki/True_Faith_(comics)										
<p>S_0: The story was one of several Crisis stories to be rapidly repackaged as a graphic novel, an unusual step for British comics of the time.</p> <p>S_1: Rian Hughes, art director for Crisis, handled the design, which included a new cover commissioned from Pleece and a foreword by Grant Morrison.</p> <p>S_2: At this stage the story began attracting negative attention from Christian groups, and The Sunday Times ran an interview with Dave Roberts (then-editor of Christian magazine Alpha), describing it as "an incitement to religious hatred" and comparing it to Salman Rushdie's The Satanic Verses as an example of blasphemy.</p> <p>S_3: Fleetway's managing director John Davidge, who had controversially blocked Crisis from publishing Peter Milligan and Brendan McCarthy's "Skin", ordered the collected edition withdrawn after a short period on sale; a press release stated this was due to the story being "inappropriate" for the book market.</p> <p>S_4: Purportedly Davidge made the decision on the direct order of Fleetway owner Robert Maxwell.</p> <p>S_5: The print run - variously reported as being 5,000 or 20,000 copies - was destroyed.</p> <p>S_6: The Economist ran an article on the controversy on 19 January 1991, speculating that "no one would think twice about it" if the story was a novel, and felt that the furore was largely caused by the ongoing belief in the press that the comic medium was still exclusively aimed at children.</p> <p>S_7: New Statesman also discussed the controversy, quoting a BBC Radio 4 interview where Eric Cartmel had his comments to generate criticism.</p>	<p>FACTS: "The Devil in the White City" was one of the most controversial stories in British comic book history, and was withdrawn from print in 1991 by publisher Fleetway.</p> <p>ERRORS: "The Devil in the White City" was one of the most controversial stories in British comic book history, and was withdrawn from print in 1991 by publisher Fleetway.</p>	<p>Click on the start of an error span not yet marked</p> <table border="1"> <thead> <tr> <th>Major severity (M)</th> <th>Minor severity (m)</th> <th>Fact (F)</th> <th>Relevance (R)</th> </tr> </thead> <tbody> <tr> <td></td> <td></td> <td></td> <td></td> </tr> </tbody> </table> <p>Relevance: Medium Relevance: 0.81 The story was one of several Crisis stories to be rapidly repackaged as a graphic novel, an unusual step for British comics of the time.</p> <p>Feedback/notes (optional):</p>	Major severity (M)	Minor severity (m)	Fact (F)	Relevance (R)				
Major severity (M)	Minor severity (m)	Fact (F)	Relevance (R)							

Figure 8: The modified anthea annotation interface for summarization.

team	s_kd	s_ps	s_sp
Haaland	0.334	0.796	0.379
DSBA	0.172	0.401	0.210
Kotonya et. al.	0.166	0.642	0.200
IUST_NLP_LAB	0.164	0.472	0.200
cometXXL	0.163	0.184	0.215
Pradhan/Todi	0.158	0.022	0.205
LTRC	0.154	0.462	0.191
iML	0.146	0.362	0.174
baselineWizard	0.133	0.327	0.163
baselineOrcaMini	0.126	0.447	0.155
baselineOrcaPlaty	0.100	0.370	0.120
baselineNous	0.100	0.370	0.120
bertscore	0.097	0.481	0.130
baselineSBERT	0.071	0.293	0.094
baselineSUPERT	0.023	-0.013	0.030

Table 13: Results of the *small* model track for summarization, when only predicting MQM.

team	s_kd	s_ps	s_sp
DSBA	0.600	0.730	0.727
iML	0.596	0.720	0.722
bertscore	0.562	0.687	0.724
IUST_NLP_LAB	0.553	0.637	0.677
baselineOrcaMini	0.549	0.595	0.669
baselineOrcaPlaty	0.536	0.606	0.638
baselineNous	0.536	0.606	0.638
Kotonya et. al.	0.522	0.525	0.634
LTRC	0.511	0.608	0.635
baselineSBERT	0.464	0.594	0.616
Pradhan/Todi	0.397	0.023	0.543
baselineWizard	0.393	0.479	0.491
Haaland	0.164	0.280	0.197
baselineSUPERT	-0.041	-0.059	-0.056
cometXXL	-0.065	-0.083	-0.092

Table 14: Results of the *small* model track for summarization, when only predicting relevance.

team	s_kd	s_ps	s_sp
iML	0.516	0.599	0.606
bertscore	0.471	0.480	0.595
DSBA	0.454	0.576	0.537
baselineOrcaPlaty	0.432	0.483	0.506
baselineNous	0.432	0.483	0.506
Pradhan/Todi	0.414	0.041	0.532
baselineOrcaMini	0.406	0.417	0.487
baselineSBERT	0.403	0.477	0.525
IUST_NLP_LAB	0.391	0.421	0.469
LTRC	0.353	0.382	0.429
Kotonya et.al.	0.348	0.220	0.417
baselineWizard	0.267	0.331	0.323
baselineSUPERT	-0.031	-0.043	-0.041
Haaland	-0.067	-0.127	-0.077
cometXXL	-0.198	-0.212	-0.265

Table 15: Results of the *small* model track for summarization, when only predicting hallucinations.

mpnet-base-v2³², which improves its performance. For COMET, we use comet-kiwi-xxlm³³, which achieved strong results on reference-free evaluation. For GEMBA we use the GEMBA library³⁴ and make small modifications to support GPT-4 requests. Finally, for BERTScore, we use xlm-roberta-large³⁵.

³²<https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

³³<https://huggingface.co/Unbabel/wmt23-cometkiwi-da-xxl>

³⁴<https://github.com/MicrosoftTranslator/GEMBA>

³⁵<https://huggingface.co/xlm-roberta-large>