

# Towards Better Evaluation of Instruction-Following: A Case-Study in Summarization

Ondrej Skopek\*  
Google Research

Rahul Aralikkatte  
Google Research

Sian Gooding  
Google Research

Victor Cărbune  
Google Research

## Abstract

Despite recent advances, evaluating how well large language models (LLMs) follow user instructions remains an open problem. While evaluation methods of language models have seen a rise in prompt-based approaches, limited work on the correctness of these methods has been conducted. In this work, we perform a meta-evaluation of a variety of metrics to quantify how accurately they measure the instruction-following abilities of LLMs. Our investigation is performed on grounded query-based summarization by collecting a new short-form, real-world dataset riSum, containing 300 document-instruction pairs with 3 answers each. All 900 answers are rated by 3 human annotators. Using riSum, we analyze the agreement between evaluation methods and human judgment. Finally, we propose new LLM-based reference-free evaluation methods that improve upon established baselines and perform on par with costly reference-based metrics that require high-quality summaries.

## 1 Introduction

Large Language Models (LLMs) have shown human-level performance in many NLP tasks. Recent advances in instruction tuning (Ouyang et al., 2022; Brown et al., 2020) and alignment (Stiennon et al., 2020; Zhou et al., 2023) have dramatically increased the ability of these models to follow instructions. In addition to being used to tackle unseen tasks in zero-shot setups (Chung et al., 2022), these models are now also used as surrogates to human annotators, especially for NLG tasks (Chiang and Lee, 2023; Wu et al., 2023), where human evaluations are time-consuming and expensive.

Consider the instruction “Briefly describe the purpose of the assignment and assumption agreement mentioned in the paragraph” from Figure 1.

\*Correspondence to: [oskopek@google.com](mailto:oskopek@google.com)

There are several dimensions to evaluate a generated output on: (i) *Coherence*: whether it is understandable and free of grammatical mistakes, (ii) *Faithfulness*: whether facts in the output are supported by the document, (iii) *Style*: whether specific formatting requirements (lists, brevity, ...) are met, and (iv) *Alignment*: whether it semantically fulfills the instruction.

Analyzing these different facets for each model output increases the cognitive load of annotators, thereby increasing the likelihood of errors or low-quality evaluations (Goyal et al., 2022). It also increases the turnaround time and hence annotations become expensive. An increasingly popular alternative is to ask LLMs to evaluate the generated outputs. Recent work like Liu et al. (2023a) and Fu et al. (2023) show that LLMs can produce human-like evaluations of text by using clever prompting techniques (Wei et al., 2022b; Yao et al., 2023). But preliminary studies have shown that LLMs can be inconsistent in their evaluations and can easily be influenced (Wang et al., 2023a; Shen et al., 2023). Gehrmann et al. (2023) have also looked at evaluation flaws and have recommended that metric developers should focus on metrics with smaller, but better defined scopes (like instruction-following).

Hence, there is an urgent need for a standard framework to analyze the specifics of instruction-following abilities of LLMs. SummEval (Fabri et al., 2021) proposes something similar for vanilla summarization. Doing this for instruction-following can be tricky because we would like to not only evaluate the LLMs as task solvers “*Summarize this document in 20 words or less*”, but also as task evaluators “*Does the summary satisfy the conditions of the instruction?*”. The meta-evaluation framework should be robust and ideally reference-free (Liu et al., 2023a). Reference-free evaluation for text generation has been widely studied (Liu et al., 2022; Hessel et al., 2021; Ke et al., 2022), but to the best of our knowledge, there has

<b>Document:</b>	
Lee, You should be receiving a package shortly containing the following: (...)	
3. Assignment and assumption agreement to move the equipment from TurboPark to the CAED I. There will be one for CAED II as well. This document is being reviewed by the bank, so I'm not convinced it is in final form. You will note that there is an acknowledgement section for GE. I cut and pasted from the consent to assignment from the TurboPark documents, but shortened the whole thing considerably. Here's that document:	
4. Signature pages (signed by Enron) from the ESA deal, both the facility agreement and the override letter. Obviously, we need your signature. I will forward the final CA facility agreements to you once again, along with the blacklines against what you initialled.	
<b>Instructions:</b>	
<ul style="list-style-type: none"> <li>• Briefly describe the purpose of the assignment and assumption agreement mentioned in the paragraph.</li> <li>• Explain the changes made to the GE acknowledgement section in the context of the TurboPark documents.</li> <li>• Summarize the final steps regarding the CA facility agreements and signature pages.</li> </ul>	
<b>Answers (for instruction #1):</b>	
F-PaLM 2-S	The assignment and assumption agreement is to move the equipment from TurboPark to the CAED I.
F-PaLM 2-Sc	The purpose of the assignment and assumption agreement is to move the equipment from TurboPark to the CAED I.
GPT-3.5	(...) The purpose of the assignment and assumption agreement is not specified.

Figure 1: Randomly sampled example from *riSum* (data source: AESLC). Highlighted how GPT-4 transforms parts of the input document into grounded instructions.

been no prior work on reference-free evaluations for instruction-following.

In this work, we take the first steps towards building such a framework. To make this problem tractable, we choose to limit our scope to the task of query-based summarization. We consider this to be an appropriate initial task since (i) numerous domains to source documents from exist, (ii) the space of appropriate instructions is broad, while still (iii) maintaining groundedness of both instructions and answers into facts present in the documents. We leave the expansion of the dataset in size and domain/instruction scope to future work.

**Contributions** For this purpose, we release a rated, instructed summarization dataset *riSum*<sup>1</sup>, consisting of 900 instruction-summary pairs with 3 human ratings each (Figure 1).

We introduce several reference-free evaluation methods which perform on-par with expensive reference-based methods and outperform existing reference-free baselines in terms of correlation with human judgement.

Lastly, we leverage *riSum* to perform an extensive meta-evaluation, quantifying how well different evaluation methods are able to replace human judgments by statistically ranking model outputs.

**Model naming** In this work, we rely on different LLMs for a variety of tasks. Specifically, we use GPT-3.5 (Ouyang et al., 2022) and GPT-4<sup>2</sup> (Ope-

<sup>1</sup> The dataset will be made available at [goo.gle/risum](http://goo.gle/risum).

<sup>2</sup> OpenAI model id: gpt-4-0314

Data source	Min	Med	Max
AESLC emails (Zhang and Tetreault, 2019)	118	172.0	469
arXiv abstracts (Clement et al., 2019)	122	145.5	224
BBC news (Narayan et al., 2018)	173	272.5	473
CNN/DM news (Hermann et al., 2015)	244	465.5	532
Common Crawl (Raffel et al., 2019)	127	282.5	506
ForumSum threads (Khalman et al., 2021)	158	320.0	519
Reddit posts (Völske et al., 2017)	156	299.0	552
SAMSum dialogues (Gliwa et al., 2019)	127	189.5	384
Task-Oriented dialogues (Lee et al., 2022)	161	329.5	605
Yelp reviews (Zhang et al., 2015)	119	140.5	357

Table 1: Data sources from which *riSum* is sampled and the minimum (Min), median (Med), and maximum (Max) sampled document length (in words). 10 documents were sampled without replacement from each of the 10 data sources.

nAI, 2023) models from the GPT LLM family, and PaLM 2-S and PaLM 2-L models from the PaLM family (Anil et al., 2023). The models are also fine-tuned on the Flan corpus as described in Anil et al. (2023, Appendix A.2), denoted as F-PaLM 2-S and F-PaLM 2-L. Finally, these models are further fine-tuned using standard methods and data known to improve instruction-following (Taori et al., 2023), denoted as F-PaLM 2-Sc and F-PaLM 2-Lc.

## 2 Data Collection

### 2.1 Dataset collection

**Data sourcing** To create *riSum*, a total of 100 documents are chosen from 10 existing datasets of different domains to ensure the data is as diverse as possible. The documents are uniformly sampled

$\alpha$ method	Follows Instruction?		How Well?	
	Mean $\pm$ SE	$\geq 50\%$	Mean $\pm$ SE	$\geq 50\%$
Global $_{n=900}$	54.3		11.4	
Local $_{n=295}$	62.1 $\pm$ 3.6	67.5	31.9 $\pm$ 4.5	56.9

Table 2: Krippendorff  $\alpha$  values (in %) of *riSum* human ratings.  $\geq 50\%$  denotes the % of pairs where  $\alpha \geq 0.5$ .

from each dataset, restricting to documents with a word count between 100 and 500 words (Table 1).

**Instruction generation** To procure instructions for each document, we first evaluate the quality of generations from four models: F-PaLM 2-Sc, F-PaLM 2-Lc, GPT-3.5, and GPT-4. We randomly sample 10 documents from the dataset and let each model generate 3 instructions per document. Each of the 40 ( $10 \times 4$  models) document-instructions pairs was rated “good”, “neutral”, or “bad” by three evaluators in a side-by-side setting. In this evaluation, GPT-4 outperformed the other models on 6/10 documents, therefore we used it to sample instructions for all documents in the dataset. This results in a total of 300 document-instruction pairs.

**Answer generation** Subsequently, three different models<sup>3</sup> are used to generate answers for each of the document-instruction pairs, yielding the final dataset with 900 data points.

**Human evaluation** Finally, each document-instruction-output triplet individually is evaluated by at least three human annotators. They are asked two questions:

1. Does the output follow the instruction? (Y/N).
2. Rate the output on a scale of 1 to 5. 1 indicates the output does not follow the instruction at all, 5 indicates the instruction is followed strictly.

See Appendix C for a description of the annotator UI, Appendix D for annotator guidelines, and Appendix E for the instruction-generation prompt.

## 2.2 Analysis of Human Ratings

For analyzing annotator agreement (Table 2), we leverage locally and globally computed Krippendorff  $\alpha$  (Krippendorff, 2019). For the first boolean question, we use the nominal distance function (indicator function) and for the second ordinal question, we use the interval distance method (squared difference). For local application, we compute a

<sup>3</sup> F-PaLM 2-S, F-PaLM 2-Sc, and GPT-3.5. We do not use GPT-4 as it was used to generate the instructions.




Figure 2: Histogram of local Krippendorff  $\alpha$  for document-instruction pairs.

localized  $\alpha$  for each document-instruction pair and then aggregate the results over all pairs. We omit 5 document-instruction pairs from the analysis for which the Krippendorff  $\alpha$  is not defined because there is no annotator overlap among the 3 ratings for each of the 3 model outputs.

We note that around 67% of the dataset has high levels of agreement on the first question and 57% on the second question. The tail of disagreement is long (Figure 2), but we hypothesize that given the difficulty of rating outputs in these diverse and highly specific texts, disagreements would be non-negligible even with higher replication rates. At the expense of gathering only relative information, ranking two responses against each other instead of rating single responses may help. Given the diversity of domains and instructions, hiring domain experts for future ratings could help increase quality and agreement, whilst also increasing costs.

Additionally, factoring out independent rating dimensions (e.g. language level, factuality) may help quantify common mistakes types in LLM instruction following and identity misalignment areas with respect to human expectations at the expense of a slower and more expensive rating process.

In Table 3a, we present aggregate numbers for annotator preferences among the three model outputs. We explore the mean of ratings, majority consensus votes (ties broken randomly), and a global mean over individual ratings (no aggregation). In Table 3b, the three model outputs are ranked for each document-instruction pair and the ranking indices are then averaged across the dataset, with ties broken randomly. Both tables are averaged over 100,000 runs to eliminate noise from tie-breaking.

Label	Agg.	F-PaLM 2-S	F-PaLM 2-Sc	GPT-3.5
FI	Mean	80.8 ± 1.8	85.0 ± 1.6	94.0 ± 1.0
	Maj.	81.3 ± 2.2	86.8 ± 2.0	96.7 ± 1.0
	None	80.9 ± 1.3	85.0 ± 1.2	94.0 ± 0.8
HW	Mean	61.1 ± 1.8	65.0 ± 1.6	73.5 ± 1.2
	Maj.	59.8 ± 2.2	64.4 ± 2.1	74.8 ± 1.5
	None	61.1 ± 1.2	65.1 ± 1.2	73.6 ± 0.9

(a) Average annotator response per answer (%; higher is better).

Label	Agg.	F-PaLM 2-S	F-PaLM 2-Sc	GPT-3.5
FI	Mean	2.12 ± 0.05	2.05 ± 0.05	1.83 ± 0.05
	Maj.	2.10 ± 0.05	2.02 ± 0.05	1.87 ± 0.05
	None	2.12 ± 0.05	2.05 ± 0.05	1.83 ± 0.05
HW	Mean	2.16 ± 0.05	2.09 ± 0.05	1.75 ± 0.05
	Maj.	2.16 ± 0.05	2.04 ± 0.05	1.80 ± 0.05
	None	2.16 ± 0.05	2.09 ± 0.05	1.75 ± 0.05

(b) Model ranking per (doc., instr.) pair (1–3, lower is better).

Table 3: Aggregate model quality according to human ratings. “Mean” aggregation takes the mean of human ratings for each model output ( $n = 300$ ), “Maj.” takes the majority vote with ties broken randomly ( $n = 300$ ), and “None” performs no aggregation ( $n = 900$ ). Averaged across 100,000 runs. FI is the binary rating “Follows Instruction?”, HW is the qualitative rating of “How Well?”. Ratings are normalized to 0–1 and reported as %.

### 3 Evaluation Methods

We propose and evaluate several methods that model annotator preferences, focusing our analysis on reference-based vs. reference-free methods and their effectiveness in different data regimes.

#### 3.1 Reference-based methods

Reference-based methods require access to at least one reference answer which can be considered the “gold standard” for each document-instruction pair. Given numerous prior work noting that summaries written by crowd workers exhibit limitations associated with lack of annotator expertise in the domain (Gillick and Liu, 2010), especially at narrower tasks like query-based summarization (Jiang et al., 2018), we use LLM-generated references for benchmarking reference-based methods instead.

The requirement of having access to high-quality references fundamentally limits the utility of the methods. In all our reference-based experiments, we use GPT-4 and F-PaLM 2-Lc generated summaries as references. Since we use GPT-3.5 and F-PaLM 2-S, and F-PaLM 2-Sc to generate candidate answers for evaluations, we use larger variants of these models to generate the “gold” references, which ensures that they are generally of higher quality (see e.g. Table 19 of Anil et al., 2023).

**BLEURT (model-based)** Sellam et al. (2020) take a (candidate, reference) answer pair as input and aim to model semantic similarity between the two texts. In all results below, we use the BLEURT<sub>20</sub> model (Pu et al., 2021). In scenarios with multiple reference answers, we take the maximum BLEURT<sub>20</sub> score across all reference answers.

**ROUGE (n-gram-based)** Lin (2004) also take (candidate, reference) pairs as input and measure

n-gram overlap to provide a numerical estimate of how well the candidate resembles the reference. We report the geometric mean of ROUGE<sub>1</sub>, ROUGE<sub>2</sub>, and ROUGE<sub>Lsum</sub> and refer to this method as ROUGE<sub>avg</sub>. Similar to BLEURT, in a scenario with multiple reference answers, we report the maximum ROUGE<sub>avg</sub> score for a given candidate.

#### 3.2 Reference-free baseline methods

We investigate popular heuristics (e.g. length of the generated response) and several LM-based approaches, varying the amount of data used. Fine-tuning a model on a subset of the collected data would also yield a viable evaluation method, but we leave that for future exploration.

**Length-based heuristics** The simplest reference-free method we use is based on length heuristics. The length of the model output is a common source of bias in human ratings when evaluating the quality of summaries, where longer answers are often preferred over shorter ones, since the former usually contains more information. Therefore, it is a natural baseline for assessing the degree to which the collected ratings suffer from this type of bias. We simply count the words and sentences using NLTK (Bird et al., 2009) and meta-evaluate how they would behave if they were used as a proxy for generated answer quality.

**Model-based methods** We benchmark the following state-of-the-art model-based methods on the *riSum* dataset: (i) BARTSCORE and BARTSCORE<sub>CNN</sub> (Yuan et al., 2021), and (ii) T5<sub>ANLI</sub> (Honovich et al., 2022). Both are encoder-decoder Transformer models and have around 400M and 11B parameters respectively.



### 3.3 LLM-based reference-free methods

The following methods depend on an underlying LLM for evaluation. Though we use PaLM 2 models in our experiments, these methods are model agnostic, and any LLM can be used in their place. For the following methods, we leverage either the base PaLM 2-S/L models, or the instruction-tuned F-PaLM 2-Sc/Lc.

**Constrained Softmax** We feed the underlying model two prompts: one for the “Follows instruction? (Y/N)” question, and another for the “How well? (1-5)” question. The prompts used correspond to the task descriptions provided to annotators (Prompts presented in Figure 9 of Appendix E).

Instead of sampling tokens to obtain the ratings, we use the model to compute the negative log-likelihood of all the possible rating values (“Yes”/“No” for the first question,  $\{1, 2, 3, 4, 5\}$  for the second question) and pick the most likely token as the rating. This approach has multiple advantages over generating tokens directly:

1. *Correctness*: The model can never output a rating that is not from the list of options.
2. *Efficiency*: All our rating values are a single token in the model’s vocabulary, which makes the scoring extremely efficient. Additionally, repeated sampling is not necessary to obtain a more precise estimate of the model’s rating.
3. *Uncertainty*: By re-normalizing the likelihoods across all rating values, we obtain a rating distribution, which lets us precisely quantify the confidence the model assigns to ratings. For an unbiased estimate with respect to the logits, we fix the softmax temperature to 1.

Finally, we return the expected value for each of the question’s distributions:

$$\mathbb{E}[R] = \sum_{j=1}^{|r|} r_j \cdot \text{softmax}(r|d, i, a)_j,$$

where  $R$  is the random variable representing the rating,  $r$  represents the rating values:  $\{0, 1\}$  for Question 1,  $\{1, 2, 3, 4, 5\}$  for Question 2.  $(d, i, a)$  represent the document, instruction, and answer.

Additionally, we discuss a variant called *Constrained Softmax n-shot*, where we contextualize the model with  $n$  examples (document-instruction-answer-rating tuples) in each of the prompts.

**Self-Agreement** In this method, we test if the model is consistent with itself across rating gen-

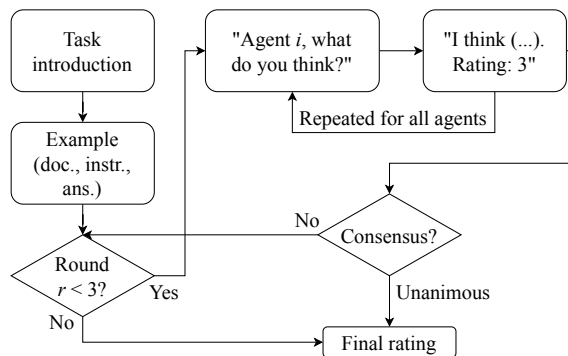


Figure 3: Multi-LLM agreement communication flow.

erations by repeatedly sampling the rating from the LLM  $n = 7$  times. To diversify the samples, we experiment with various softmax temperatures, only to find that lower temperatures yield better results<sup>4</sup>. The final rating is the arithmetic mean of the individual samples. We contextualize the model with  $k = 3$  examples in the prompt (see Figure 7 in Appendix E). We also investigate the following variants:

- *no intro* Omitting the description of the task in the prompt and using only the  $k$  examples.
- *rationale* Asking the model to generate Chain of Thought-like “rationales” for the given rating to each few-shot example (Wei et al., 2022b).
- *random* Using the same hand-crafted examples (not occurring in the dataset) vs. picking  $k$  random examples from the remaining documents in the dataset.

**Multi-LLM Agreement** Recent works (Bakker et al., 2022; Park et al., 2023) have used LLMs in conversational settings where all participant LLMs communicate with each other and try to achieve a common goal. We propose a consensus-based metric where  $k$  LLM instances<sup>5</sup> debate amongst each other and try to arrive at a common assessment. Though there are no restrictions on the LLMs to use, we evaluate the simplest case where each instance is the same LLM. The rules of communication are set as follows (Figure 3):

1. The models communicate amongst each other in a controlled manner for up to 3 rounds and try to arrive at a consensus. After at most 3 rounds, one of three outcomes occurs: (i) *unanimous agreement*: all 3 models agree. If this happens

<sup>4</sup> Temperature is set to 0.1 for all reported Self-Agreement and Multi-LLM Agreement experiments.

<sup>5</sup>  $k = 3$  in all our experiments.

in the earlier rounds, the process ends immediately, (ii) *majority agreement*: one model disagrees with the other two, or (iii) *disagreement*: all 3 models disagree with each other.

2. In each round, all models provide a rating and a brief rationale. The models do not have access to the other model outputs till the end of a round<sup>6</sup>. Before the start of rounds two and three, they receive the ratings and rationales of all models from the previous rounds.

The prompt for the models is presented in Figure 8 of Appendix E. This method is referred to as *Multi-LLM Agreement* henceforth. We repeat the process  $n = 3$  times for added stability.

## 4 Evaluating Agreement with Annotators

As discussed in Section 2, we asked annotators to provide a binary Yes/No rating answering whether a model output *follows the instruction* and a qualitative rating from 1 to 5, representing *how well* it follows the instruction. Using meta-evaluation methods described below, we then study agreement between annotators and our evaluation methods.

### 4.1 “Follows Instruction?”

For the binary rating, we compute a macro-averaged *Area Under ROC Curve* (AUC ROC) statistic for each evaluation method. Using AUC ROC, we analyze the effectiveness of each method if they were used as binary classifiers for “Does the output follow the instruction?”, thereby assessing the degree to which they can replace human ratings. Since our classes are imbalanced towards “Yes” (Table 3a) we opt for the macro-averaged version of ROC AUC so that we can better detect which methods can accurately predict the “No” class.

### 4.2 “How well?”

**Rank-based evaluation** To analyze the ability of evaluation metrics to rank model outputs in relation to each other, we compute *Kendall’s  $\mathcal{T}_b$  rank distance*  $d_{\mathcal{T}_b}$  among the model outputs for each document-instruction pair. When the ranking produced by a metric is independent from human ranking, the value of  $d_{\mathcal{T}_b}$  will be equal to 0.5 in expectation. Values below 0.5 represent rankings that are similar to the human ranking, values above 0.5 represent orderings that are similar to the inverse of the

human ranking. As opposed to the  $\mathcal{T}_b$  rank correlation coefficient,  $d_{\mathcal{T}_b}$  has values in the range of  $[0, 1]$  and can be interpreted as a distance function (lower is better):  $d_{\mathcal{T}_b} = (1 - \mathcal{T}_b) / 2$ . Compared to other forms of  $\mathcal{T}$ ,  $\mathcal{T}_b$  adjusts for ties: situations, where a metric or annotators give the same rating to two or more model outputs for one document-instruction pair.

For our human ratings,  $\mathcal{T}_b$  is not defined for 9 out of 300 document-instruction pairs: the mean of the 3 annotators’ ratings is constant for all 3 models, making it impossible to rank the models. We report the mean and standard error of the rank distance  $d_{\mathcal{T}_b}$  across all non-constant pairs.

**Linear value correlation** Additionally, we would like evaluation method outputs to align with annotators’ notions of “good” or “bad”. To study this, we compute Pearson’s distance across all document-instruction-answer tuples:  $d_{|r|} = 1 - |r|$ , where  $r$  is Pearson’s correlation coefficient between an evaluation method’s values and the mean annotator rating. Values of  $d_{|r|}$  range from 0 to 1; the lower the value, the higher the linear correlation with human ratings.

## 5 Results and Analysis

We compare the effectiveness of evaluation methods on the three rating dimensions, based on the reported numbers for the binary rating “Follows Instruction?” and for the qualitative rating “How well?” in Table 4. For both rating tasks, the two length-based heuristics perform the worst out of all methods, which suggests that the instructions are of good quality, as annotators are not strongly influenced by the length of model outputs.

### 5.1 Predicting “Follows Instruction?”

First, we focus on how good of a binary classifier the methods are. We report the AUC ROC and its standard error (Section 4.1) with respect to the human majority vote labels.

**Reference-based methods** Having access to several reference answers that follow the instruction continues to be a good indicator when combined with ROUGE or BLEURT. However, the results show that, when we have access to a capable LLM like F-PaLM 2-Lc, it is better to use it directly as a reference-free evaluator, than sampling reference summaries from it and using reference-based metrics like ROUGE<sub>avg</sub> and BLEURT<sub>20</sub>.

<sup>6</sup> Empirically, models tend to agree more easily with each other when shown other models’ ratings before the round ends.

Evaluation Method	Follows Instruction?	How Well?	
	AUC ROC % $\uparrow$	$\overline{d_{\mathcal{T}_b}}$ % $\downarrow$	$\overline{d_{ r }}$ % $\downarrow$
<b>Reference-based baseline methods (Section 3.1)</b>			
BLEURT <sub>20</sub> [references: GPT-4]	<b>78.5</b> $\pm$ 1.9	41.1 $\pm$ 1.9	50.9 $\pm$ 2.9
BLEURT <sub>20</sub> [references: F-PaLM 2-Lc]	71.8 $\pm$ 2.3	48.8 $\pm$ 1.9	54.4 $\pm$ 2.8
ROUGE <sub>avg</sub> [references: GPT-4]	<b>79.5</b> $\pm$ 1.9	<b>35.4</b> $\pm$ 1.9	52.6 $\pm$ 2.8
ROUGE <sub>avg</sub> [references: F-PaLM 2-Lc]	71.1 $\pm$ 2.3	46.7 $\pm$ 1.9	60.8 $\pm$ 2.7
<b>Reference-free baseline methods (Section 3.2)</b>			
Sentence Count	39.5 $\pm$ 3.1	54.8 $\pm$ 1.8	72.7 $\pm$ 2.3
Word Count	42.2 $\pm$ 3.1	51.4 $\pm$ 2.0	71.0 $\pm$ 2.4
BARTSCORE (Yuan et al., 2021)	68.4 $\pm$ 2.5	45.0 $\pm$ 1.9	74.7 $\pm$ 2.2
BARTSCORE <sub>CNN</sub> (Yuan et al., 2021)	69.7 $\pm$ 2.4	43.7 $\pm$ 1.9	70.3 $\pm$ 2.4
T5 <sub>ANLI</sub> (Honovich et al., 2022)	71.9 $\pm$ 2.3	<b>38.8</b> $\pm$ 1.9	64.7 $\pm$ 2.5
<b>LLM-based reference-free methods (Section 3.3)</b>			
PaLM 2-S Constrained Softmax	74.0 $\pm$ 2.2	43.6 $\pm$ 1.9	80.0 $\pm$ 2.0
PaLM 2-L Constrained Softmax	77.8 $\pm$ 2.0	39.9 $\pm$ 1.9	46.4 $\pm$ 3.0
F-PaLM 2-Sc Self-Agreement	67.2 $\pm$ 2.5	42.8 $\pm$ 1.7	56.7 $\pm$ 2.7
F-PaLM 2-Lc Self-Agreement	<b>81.7</b> $\pm$ 1.7	<b>37.1</b> $\pm$ 1.7	<b>39.5</b> $\pm$ 3.1
F-PaLM 2-Lc Self-Agreement (+ no intro)	<b>79.7</b> $\pm$ 1.9	<b>38.4</b> $\pm$ 1.8	45.7 $\pm$ 3.0
F-PaLM 2-Lc Self-Agreement (+ rationale)	75.0 $\pm$ 2.1	43.2 $\pm$ 1.3	50.5 $\pm$ 2.9
F-PaLM 2-Sc Self-Agreement (random)	69.0 $\pm$ 2.4	42.7 $\pm$ 1.7	58.3 $\pm$ 2.7
F-PaLM 2-Lc Self-Agreement (random)	<b>80.4</b> $\pm$ 1.8	<b>37.0</b> $\pm$ 1.8	<b>42.2</b> $\pm$ 3.0
F-PaLM 2-Lc Self-Agreement (random + no intro)	<b>78.2</b> $\pm$ 1.9	39.5 $\pm$ 1.8	50.2 $\pm$ 2.9
F-PaLM 2-Sc Multi-LLM Agreement	66.4 $\pm$ 2.5	45.7 $\pm$ 1.2	61.8 $\pm$ 2.6
F-PaLM 2-Lc Multi-LLM Agreement	67.1 $\pm$ 2.5	46.0 $\pm$ 1.2	58.7 $\pm$ 2.7

Table 4: AUC ROC Curve measures how well methods predict Yes/No annotator responses on “Follows Instruction?” ( $n = 900$ ). For “How Well?” (1–5 rating), we report Kendall’s rank distance  $d_{\mathcal{T}_b}$  comparing evaluation methods’ ranking of answers to that of annotators’ ( $n = 291$ ) and Pearson’s distance from mean annotator responses  $d_{|r|}$  ( $n = 900$ ). All values are in %,  $\pm$  signifies standard error,  $\uparrow$  signifies higher is better ( $\downarrow$  lower is better). Methods highlighted in bold have overlapping confidence intervals with the best method per column. Non-deterministic methods (Self-Agreement, Multi-LLM Agreement) have been re-run  $5\times$  and the mean is reported.

**Reference-free methods** As expected, performance of each evaluation method improves with model size. We observe that standard error is usually higher ( $> 2.0$ ) when using PaLM 2-S compared to PaLM 2-L ( $< 2.0$ ), across different methods. Combined with generally lower performance, methods using PaLM 2-S as the underlying model are more noisy and produce less meaningful evaluations compared to methods using PaLM 2-L.

We also note that Multi-LLM Agreement approaches, while interesting, are outperformed by both Self-Agreement and Constrained Softmax approaches, irrespective of the model size.

For scoring-based approaches (Constrained Softmax), non-instruction-tuned LLMs outperform their instruction-tuned counterparts. When generation is involved, instruction-tuned models outperform their base versions. This applies to rating generation, but also for generating answers directly. We only report numbers of instruction-tuned LLMs for generation-based methods and corre-

spondingly, only report numbers of non-instruction-tuned LLMs for scoring-based approaches.

## 5.2 Predicting “How Well?”

In the case of qualitative ratings, obtaining a ranking of answers that matches the annotators’ ranking proves to be difficult. We note sensitivity in the analysis with respect to how ratings are aggregated per answer (majority vote or mean). To minimize ties and maximize the use of annotator information, we use mean aggregation for the following analysis.

Observing  $d_{\mathcal{T}_b}$  ranking performance, ROUGE<sub>avg</sub> using GPT-4 model-generated answers seems to perform on-par with F-PaLM 2-Lc Self-Agreement based methods, as well as the 11B parameter T5<sub>ANLI</sub> model from Honovich et al. (2022).

**Reference-based methods** In our experiments, BLEURT performs worse than ROUGE at relative ranking of model outputs. Since ROUGE is based on surface form, there is reason to believe that sam-

Evaluation Method	Perfect agreement	Disagreement	Prefers own LM family
Constr. Softmax PaLM 2-L	35.7%	40.3%	56.2%
ROUGE F-PaLM 2-Lc	27.0%	48.0%	93.1%
Constr. Softmax F-PaLM 2-Lc	23.0%	54.3%	71.8%
Self-Agr. F-PaLM 2-Lc	25.7%	23.0%	72.5%

Table 5: Agreement analysis with respect to mean qualitative ranking (“How well?”).

ples from different models in a single LM-family are closer in surface form than samples from different LM-families. In Table 5, we analyze the effectiveness of methods at picking the best answers out of the 3 model outputs. *Perfect agreement* happens when the sets of annotator and metric “winner” models is equal. *Disagreement* occurs when the intersection between annotator and metric winners is empty. Within disagreement, *prefers own LM family* means the metric winners contained *at least one* model output from the LM family the metric is based on.

We observe that when the evaluation model is sufficiently different from the rated models, the likelihood of evaluation models preferring their own LM family goes down. However, when using a similar model, reference-based methods are more biased towards preferring their own LM family. If human-written reference answers are unavailable, using a reference-free metric is preferable.

**Reference-free methods** Similarly to the binary rating, we observe that methods with larger underlying models perform better. Likewise, reference-free methods based on F-PaLM 2-Lc outperform their reference-based counterparts when using the same underlying model. The base PaLM 2-L model with Constrained Softmax performs better and at lower cost than using the instruction-tuned F-PaLM 2-Lc to generate reference summaries. With more available compute, one can further improve performance by leveraging multi-sampling Self-Agreement methods.

Interestingly, using *random* examples in Self-Agreement decreases performance as opposed to hand-crafting a small ( $k = 4$ ) set of held-out examples. Contrary to intuition, using Chain-of-Thought approaches (*rationale*) seems to degrade performance, but when removing the task description (*no intro*) we do not observe a big drop.

When linear correlation  $d_{|r|}$  with human ratings is required, methods that model the qualitative rat-

ing directly outperform more generic methods.

## 6 Related Work

**Measuring instruction following with LLMs** Liu et al. (2023a) use GPT-4 as a backbone model and study the correlation with human ratings on non-query-based summarization, finding a bias towards LLM-generated text. We do not study this aspect, as our rating task focuses on model-generated text. Fu et al. (2023) propose a zero-shot approach for multi-faceted evaluation of text generation.

An increase in interest for improving instruction-following capabilities of LLMs has resulted in the creation of multiple datasets. FLAN (Wei et al., 2022a) and Natural Instructions (Mishra et al., 2022) were two of the earlier datasets which turned standard NLP tasks (e.g. sentiment classification, question-answering) into instruction following tasks. Other works like Self-Instruct (Wang et al., 2023b), Super-NaturalInstructions (Wang et al., 2022), and the H4 instruction dataset (Hugging Face, 2023) curate human-written instruction and answer pairs. Guo et al. (2023) and Qingyi Si (2023) collect instruction-answer pairs from LLM generations. All of them use standard NLP metrics or human annotation to evaluate the model outputs.

**Model-based metrics** A large body of prior work focuses on model-based approaches fine-tuned on human ratings. Usually, encoder models such as BERTSCORE (Zhang et al., 2020) or BLEURT (Selam et al., 2020) are used, but encoder-decoder methods exist as well (BARTSCORE, Yuan et al., 2021). We focus on low-resource zero/few-shot methods using larger, decoder-containing models from PaLM and GPT families.

**Human evaluation** Kryściński et al. (2018); Huang et al. (2020); Shen et al. (2022b) and several others have resorted to human evaluation for analyzing the quality of reference summaries and model outputs. They adopt a Likert-type scale for rating individual aspects of generated text. Fan et al. (2018); Fabbri et al. (2019); Shen et al. (2022a) and others perform side-by-side comparisons of two or more model-generated summaries and use Elo, or other rating systems to build rankings of models.

**LLM evaluation** Many recent works use LLMs as evaluators for summarization tasks. Wu et al. (2023) use LLMs with “different persona” to evaluate summaries from various perspectives. Luo et al.



(2023) examine if LLMs can be used to detect factual inconsistencies. Concurrent to our work, Liu et al. (2023b) curate a human-evaluation dataset consisting of 22,000 summary-level annotations and perform a study of various automatic and LLM-based metrics for summarization and call for more rigorous evaluation of LLM performance.

## 7 Conclusion

In this work, we investigate the effectiveness of multiple evaluation methods in quantifying the degree to which LLM-generated text follows user-given instructions. We release *riSum*, a new short-form dataset of 300 document-instruction pairs with 3 answers each. All of the 900 answers are rated by at least 3 human annotators. When analyzing agreement between evaluation methods and human judgment, we find that established metrics, such as ROUGE and BLEURT are not effective at quantifying LLMs' instruction-following ability. LLM-based evaluation methods tend to have stronger correlation with annotator judgment, without requiring high-quality reference answers. We hope that the introduced evaluation framework is adopted by the community for evaluating instruction-following abilities of LLMs, possibly expanding into more tasks, domains, and examples.

## Acknowledgements

The authors would like to thank Lucas Werner for building the human annotation tool which dramatically simplified the creation of the *riSum* dataset, as well as all the 14 annotators and their managers for their tireless and accurate work. Thank you also to Hassan Mansoor and Misha Khalman, who have helped with exploratory work on evaluation methods and early data collection. We also thank the anonymous reviewers for their constructive feedback. Lastly, we would like to thank Matt Sharifi, Jindong Chen and Blaise Agüera y Arcas for their feedback on the paper draft.

## References

Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham,

Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Auroko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. 2023. [PaLM 2 Technical Report](#).

Michiel A. Bakker, Martin J Chadwick, Hannah Sheahan, Michael Henry Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matthew Botvinick, and Christopher Summerfield. 2022. [Fine-tuning language models to find agreement among humans with diverse preferences](#). In *Advances in Neural Information Processing Systems*.

Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python: Analyzing text with the Natural Language Toolkit*. O'Reilly Media, Inc.

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.

Cheng-Han Chiang and Hung-yi Lee. 2023. [Can Large Language Models Be an Alternative to Human Evaluations?](#) In *Proceedings of ACL*.

- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling Instruction-Finetuned Language Models](#).
- Colin B. Clement, Matthew Bierbaum, Kevin P. O’Keefe, and Alexander A. Alemi. 2019. [On the use of arxiv as a dataset](#).
- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. [Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model](#). In *Proceedings of ACL*.
- 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, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical Neural Story Generation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. [GPTScore: Evaluate as You Desire](#).
- Sebastian Gehrmann, Elizabeth Clark, and Thibault Sèlam. 2023. [Repairing the Cracked Foundation: A Survey of Obstacles in Evaluation Practices for Generated Text](#). *Journal of Artificial Intelligence Research*, 77:103–166.
- Dan Gillick and Yang Liu. 2010. [Non-Expert Evaluation of Summarization Systems is Risky](#). In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*, pages 148–151, Los Angeles. Association for Computational Linguistics.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. [SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization](#). In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 70–79, Hong Kong, China. Association for Computational Linguistics.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. [FALTE: A Toolkit for Fine-grained Annotation for Long Text Evaluation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 351–358, Abu Dhabi, UAE. Association for Computational Linguistics.
- Biyang Guo, Xin Zhang, Ziyuan Wang, Minqi Jiang, Jinran Nie, Yuxuan Ding, Jianwei Yue, and Yupeng Wu. 2023. [How Close is ChatGPT to Human Experts? Comparison Corpus, Evaluation, and Detection](#). *arXiv preprint arxiv:2301.07597*.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. [Teaching machines to read and comprehend](#). In *NIPS*, pages 1693–1701.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. [CLIPScore: A reference-free evaluation metric for image captioning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7514–7528, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022. [TRUE: Re-evaluating factual consistency evaluation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3905–3920, Seattle, United States. Association for Computational Linguistics.
- Dandan Huang, Leyang Cui, Sen Yang, Guangsheng Bao, Kun Wang, Jun Xie, and Yue Zhang. 2020. [What have we achieved on text summarization?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 446–469, Online. Association for Computational Linguistics.
- Hugging Face. 2023. [HuggingFaceH4/instruction-dataset](#).
- Youxuan Jiang, Catherine Finegan-Dollak, Jonathan K. Kummerfeld, and Walter Lasecki. 2018. [Effective Crowdsourcing for a New Type of Summarization Task](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 628–633, New Orleans, Louisiana. Association for Computational Linguistics.
- Pei Ke, Hao Zhou, Yankai Lin, Peng Li, Jie Zhou, Xiaoyan Zhu, and Minlie Huang. 2022. [CTRLEval: An unsupervised reference-free metric for evaluating controlled text generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2306–2319, Dublin, Ireland. Association for Computational Linguistics.
- Misha Khalman, Yao Zhao, and Mohammad Saleh. 2021. [ForumSum: A multi-speaker conversation summarization dataset](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4592–4599, Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Klaus Krippendorff. 2019. *Content Analysis: An Introduction to Its Methodology*. SAGE Publications, Inc.
- Wojciech Kryściński, Romain Paulus, Caiming Xiong, and Richard Socher. 2018. [Improving abstraction in text summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1808–1817, Brussels, Belgium. Association for Computational Linguistics.
- Harrison Lee, Raghav Gupta, Abhinav Rastogi, Yuan Cao, Bin Zhang, and Yonghui Wu. 2022. [SGD-X: A Benchmark for Robust Generalization in Schema-Guided Dialogue Systems](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):10938–10946.
- 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.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruo Chen Xu, and Chenguang Zhu. 2023a. [G-Eval: NLG Evaluation using GPT-4 with Better Human Alignment](#).
- Yixin Liu, Alexander R. Fabbri, Pengfei Liu, Yilun Zhao, Linyong Nan, Ruilin Han, Simeng Han, Shafiq Joty, Chien-Sheng Wu, Caiming Xiong, and Dragomir Radev. 2023b. [Revisiting the gold standard: Grounding summarization evaluation with robust human evaluation](#).
- Yizhu Liu, Qi Jia, and Kenny Zhu. 2022. [Reference-free summarization evaluation via semantic correlation and compression ratio](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2109–2115, Seattle, United States. Association for Computational Linguistics.
- Zheheng Luo, Qianqian Xie, and Sophia Ananiadou. 2023. [ChatGPT as a Factual Inconsistency Evaluator for Text Summarization](#).
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2022. [Cross-Task Generalization via Natural Language Crowdsourcing Instructions](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3470–3487, Dublin, Ireland. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don’t Give Me the Details, Just the Summary! Topic-Aware Convolutional Neural Networks for Extreme Summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- OpenAI. 2023. [GPT-4 Technical Report](#).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*.
- Joon Sung Park, Joseph C. O’Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. [Generative agents: Interactive simulacra of human behavior](#).
- Amy Pu, Hyung Won Chung, Ankur Parikh, Sebastian Gehrmann, and Thibault Sellam. 2021. [Learning Compact Metrics for MT](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 751–762, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zheng Lin Qingyi Si. 2023. [Alpaca-CoT: An Instruction Fine-Tuning Platform with Instruction Data Collection and Unified Large Language Models Interface](#). <https://github.com/PhoebusSi/alpaca-CoT>.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *arXiv e-prints*.
- 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, Lidong Bing, Yang You, and Luo Si. 2022a. [SentBS: Sentence-level Beam Search for Controllable Summarization](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10256–10265, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Chenhui Shen, Liying Cheng, Yang You, and Lidong Bing. 2023. [Are Large Language Models Good Evaluators for Abstractive Summarization?](#)
- Chenhui Shen, Liying Cheng, Ran Zhou, Lidong Bing, Yang You, and Luo Si. 2022b. [MReD: A Meta-Review Dataset for Structure-Controllable Text Generation](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2521–2535, Dublin, Ireland. Association for Computational Linguistics.
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020. [Learning](#)



- to Summarize from Human Feedback. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, Red Hook, NY, USA. Curran Associates Inc.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford Alpaca: An Instruction-following LLaMA model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- Michael Völske, Martin Potthast, Shahbaz Syed, and Benno Stein. 2017. **TL;DR: Mining Reddit to learn automatic summarization**. In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 59–63, Copenhagen, Denmark. Association for Computational Linguistics.
- Peiyi Wang, Lei Li, Liang Chen, Dawei Zhu, Binghui Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023a. **Large Language Models are not Fair Evaluators**.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023b. **Self-Instruct: Aligning Language Models with Self-Generated Instructions**.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krma Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sampat, Siddhartha Mishra, Sujan Reddy A, Sumanta Patro, Tanay Dixit, and Xudong Shen. 2022. **Super-NaturalInstructions: Generalization via declarative instructions on 1600+ NLP tasks**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5085–5109, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jason Wei, Maarten Paul Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew Mingbo Dai, and Quoc V. Le. 2022a. **Finetuned Language Models are Zero-Shot Learners**. In *International Conference on Learning Representations*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022b. **Chain of Thought Prompting Elicits Reasoning in Large Language Models**. In *Advances in Neural Information Processing Systems*.
- Ning Wu, Ming Gong, Linjun Shou, Shining Liang, and Daxin Jiang. 2023. **Large Language Models are Diverse Role-Players for Summarization Evaluation**.
- 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.
- Rui Zhang and Joel Tetreault. 2019. **This Email Could Save Your Life: Introducing the Task of Email Subject Line Generation**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 446–456, Florence, Italy. Association for Computational Linguistics.
- 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*.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. **Character-Level Convolutional Networks for Text Classification**. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1, NIPS'15*, page 649–657, Cambridge, MA, USA. MIT Press.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023. **LIMA: Less is more for alignment**.



## A Limitations

While the presented data offers a variety (e.g. diverse origin texts), a drawback to our work is that we only consider the task of instruction-based summarization (e.g. long-form question answering, query-driven summarization, stylistic summarization) as such. The extent to which metrics generalize to other tasks is not yet explored. Furthermore, for language diversity, the proposed benchmarks are restricted to English only. However, we hope that this initial benchmark allows further work to consider a larger range of tasks as well as exploration for how these benchmarks generalize to other languages.

Our correlation with human judgment analysis on the qualitative rating (“How Well?”) has a limitation where the annotators do not provide sufficient signal to distinguish between the 3 answers. This happens in only 9 out of the 300 document-instruction pairs and we chose to skip those pairs in the analysis for this rating task. The motivation for doing this is that our focus is on the cases where there is sufficient signal from the human annotators when an answer is better than another.

We acknowledge that relying on human ratings as a ground truth has drawbacks, especially as summarization is notoriously difficult to evaluate due to the subjective nature. To mitigate this, we provide extensive training and feedback to annotators and are in active communication throughout the annotation process to provide clarifications. The annotators used in our experiment have over a year of experience with rating NLU tasks. However, a limitation is that our annotator pool represents individuals from similar backgrounds, which may mean other populations would have differing quality perspectives. The background statistics of annotators can be found in Appendix C.2.

## B Ethics Statement

The alignment of model behavior with user expectations is a crucial area of research, and we recognize the importance of contributing to the development of benchmarking methods for instruction following. Our work represents a step towards benchmarking how LLMs can self-evaluate their performance in the task of summarization. However, there are still many other aspects of summary quality, such as factuality, that warrant further exploration due to their significant downstream implications.




Figure 4: Number of questions annotated by each human annotator. Annotator IDs pseudonymized to capital letters.

A model’s ability to follow instructions for a specific task, such as summarization, may not reflect the overall proficiency in instruction following. As such, these metrics serve as proxies to estimate the extent to which task instructions are adhered to within the context of summarization. Given the ongoing discussions regarding the risks associated with LLMs, this distinction is relevant.

During dataset construction, it is important to acknowledge the ethical concerns arising from the use of publicly sourced data without explicit permission from the original parties. While the data we employ is derived from previously released datasets, the examples are generated using LLMs trained on large, uncurated, static datasets obtained from the internet.

## C Annotator methodology

### C.1 Annotation UI

In Figure 5 we illustrate the user interface used for collecting the dataset. Annotators follow a multi-step process, by first answering “Does the output follow the instruction?” followed by “Rate the output on a scale of 1 to 5” to qualitatively assess the answer.

The UI also allows annotators to navigate through the provided content and highlight words that appear either in the answer or in the original text. Annotators can use this as a way to verify that content is present in both the output and input.

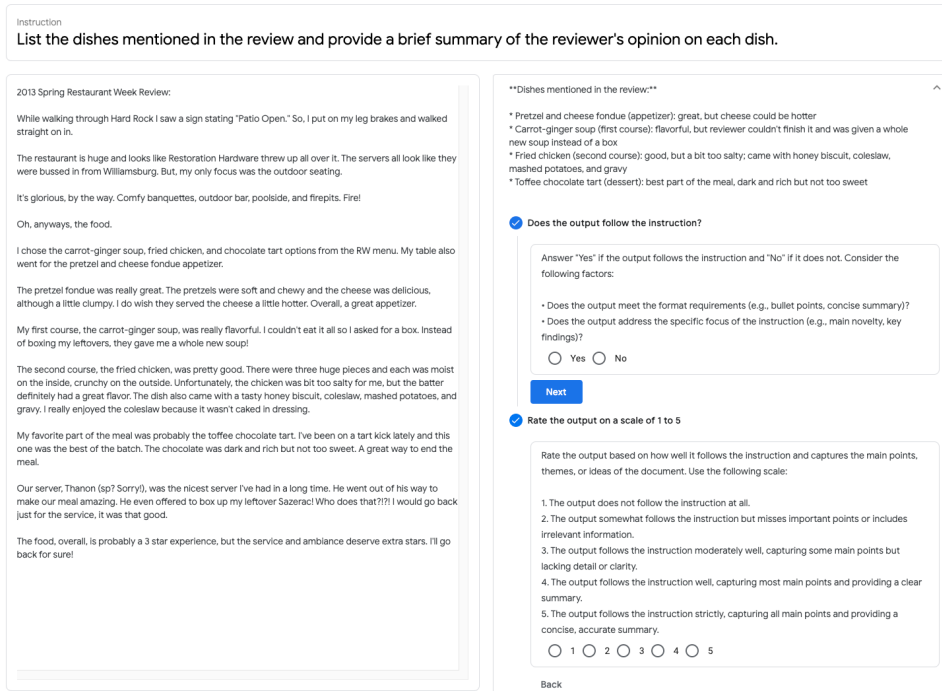


Figure 5: Example screenshot of the annotator UI.

## C.2 Annotator demographics

Table 6 presents the results of an optional questionnaire given to our annotators, aimed at understanding their background factors. Out of a total of 14 annotators, we have received responses from 7 individuals, who collectively accounted for approximately 65% of the annotation coverage for our dataset (Figure 4). This information allows us to gain a better understanding of the perspectives and experiences of our annotators, which can impact the annotation outcomes.

## D Annotator Guidelines

### D.1 Objective

The goal of this task is to evaluate the quality of summaries generated based on given instructions. You will be provided with a document, an instruction, and an output (summary). Your task is to answer two questions:

1. Does the output follow the instruction? (Yes/No), and
2. Rate the output on a scale of 1 to 5, with 1 indicating that the output does not follow the instruction at all, and 5 indicating that the output follows the instruction strictly.

### D.2 General Guidelines

**Understanding the Document** Before evaluating the output, make sure you have a clear understanding of the document. The document can be a news article, a chat conversation, an email, etc. Read the document carefully and identify the main points, themes, or ideas.

**Analyzing the Instruction** The instruction will be related to summarization. It can be general (e.g. "Summarize in 3 bullet points") or specific to the paragraph (e.g. "Summarize the main novelty of the research work concisely"). Make sure you understand the instruction and its requirements.

**Evaluating the Output** Compare the output with the document and the instruction. Check if the output follows the instruction and captures the main points, themes, or ideas of the document.

**Evaluation Criteria** For Question 1, answer "Yes" if the output follows the instruction and "No" if it does not. Consider the following factors: (i) does the output meet the format requirements (e.g. bullet points, concise summary)? and (ii) does the output address the specific focus of the instruction (e.g. main novelty, key findings)?

For Question 2, rate the output based on how well it follows the instruction and captures the main points, themes, or ideas of the document. Use the

Proficiency		Education		Age range		Hours reading English per week	
Native	1/7	Graduate	6/7	18–24	3/7	0–5	1/7
Near native	1/7	Undergraduate	1/7	25–34	4/7	5–10	2/7
Advanced	5/7	High School	0/7	35–44	0/7	10–15	1/7
Intermediate	0/7	Vocational Training	0/7	45–54	0/7	15–20	0/7
Beginner	0/7	No formal education	0/7	55+	0/7	20+	3/7

Table 6: Background statistics for annotators.

following scale:

1. The output does not follow the instruction at all.
2. The output somewhat follows the instruction but misses important points or includes irrelevant information.
3. The output follows the instruction moderately well, capturing some main points but lacking detail or clarity.
4. The output follows the instruction well, capturing most main points and providing a clear summary.
5. The output follows the instruction strictly, capturing all main points and providing a concise, accurate summary.

### D.3 FAQs

**What if the output is well-written but does not follow the instruction?** Rate the output based on how well it follows the instruction, not on its writing quality. If the output does not follow the instruction, give it a low rating.

**What if the output follows the instruction but has grammatical errors or typos?** Focus on the content and adherence to the instruction. Minor grammatical errors or typos should not significantly impact the rating unless they affect the clarity or accuracy of the summary.

**What if the output is too long or too short?** Consider whether the output meets the requirements of the instruction. If the instruction specifies a length (e.g. “Summarize in 3 bullet points”), the output should adhere to that length. If the output is too long or too short, it may not follow the instructions strictly, and you should adjust the rating accordingly.

**What if the output is accurate but not concise?** If the instruction requires a concise summary, the

output should be brief and to the point. If the output is accurate but not concise, it may not follow the instructions strictly, and you should adjust the rating accordingly.

### E Prompts

List of prompts used in different parts of the paper:

- GPT-4 prompt for generating *riSum* instructions: Figure 6.
- Self-agreement prompt: Figure 7.
- Multi-LLM agreement prompt: Figure 8.
- Constrained Softmax prompt: Figure 9.

Read the paragraph given by the user and generate a list of 3-5 instructions for human annotators. Each instruction must be in a new line.

The instructions must be related to the task of summarization. Some general examples are: Summarize in 3 bullet points. Write the main topics of the document in 2 sentences. Summarize the paragraph in not more than 20 words.

However, you can ask them to perform something specific related to the content of the paragraph. Summarize the main novelty of the research work concisely. Summarize the cleaning tips using soap and sponge in details for me so I sound like a professional. Summarize the purpose of the dialogue and then convert each person's opinion into a bullet list while keeping their orders.

Be as creative as possible, and use the information present in the paragraph to make the instructions unique.

Figure 6: Prompt given to GPT-4 for creating the instructions.

You are given a document, an instruction, and a candidate answer.  
You have to evaluate the answer based on how well it follows the instructions on a scale of 1 to 5 (larger is better), and provide a rationale.  
Carefully evaluate the various constraints that may be present in the instructions.

—

Document:  
{document}

Instruction:  
{instruction}

Answer:  
{answer}

Rating:

Figure 7: Self-agreement prompt. The bottom part under and including “—” is repeated for  $n > 1$ -shot variants.

This is a chat room with AI assistants that specialize in summarizing and question answering.  
You are given a paragraph of text, an instruction, and a candidate answer.  
You have to evaluate the answer based on how well it follows the instructions on a scale of 1 to 5.  
Carefully evaluate the various constraints that may be present in the instruction.  
After evaluation, present a brief rationale not exceeding 2-3 sentences, and your rating, to the AI assistants.  
If there is consensus among the AI assistants, the rating will be accepted.  
If there is no consensus, you should read the rationale of the other AI assistants and try to reach a consensus by either changing your rating or convincing the other assistants to change theirs.  
You will be given 3 chances to reach a consensus.  
Always try to reach a consensus.  
Remember, end your response with 'Rating:'.

Document:  
{document}

Instruction:  
{instruction}

Answer:  
{answer}

(User: Agent {aid}, please share your response.)

Agent {aid}: ... rationale ... Rating: 4.

...

Figure 8: Prompt given to the models before the consensus discussion.



Question #1

Does the output follow the instruction? Rate “Yes” if the output follows the instruction and “No” if it does not. Consider the following factors:

\* Does the output meet the format requirements (e.g., bullet points, concise summary)?

\* Does the output address the specific focus of the instruction (e.g., main novelty, key findings)?

Document:  
{ document }

Instruction:  
{ instruction }

Output:  
{ answer }

Rating:

Question #2

Rate the output on a scale of 1 to 5. Rate the output based on how well it follows the instruction and captures the main points, themes, or ideas of the document. Use the following scale:

1. The output does not follow the instruction at all.

2. The output somewhat follows the instruction but misses important points or includes irrelevant information.

3. The output follows the instruction moderately well, capturing some main points but lacking detail or clarity.

4. The output follows the instruction well, capturing most main points and providing a clear summary.

5. The output follows the instruction strictly, capturing all main points and providing a concise, accurate summary.

Document:  
{ document }

Instruction:  
{ instruction }

Output:  
{ answer }

Rating:

Figure 9: Prompts for Constrained Softmax-based methods.