

Aligning Black-box Language Models with Human Judgments

Gerrit J. J. van den Burg, Gen Suzuki, Wei Liu, Murat Sensoy
{gvdburg, gfsuzuki, weliuz, msensoy}@amazon.co.uk
Amazon

Abstract

Large language models (LLMs) are increasingly used as automated judges to evaluate recommendation systems, search engines, and other subjective tasks, where relying on human evaluators can be costly, time-consuming, and unscalable. LLMs offer an efficient solution for continuous, automated evaluation. However, since the systems that are built and improved with these judgments are ultimately designed for human use, it is crucial that LLM judgments align closely with human evaluators to ensure such systems remain human-centered. On the other hand, aligning LLM judgments with human evaluators is challenging due to individual variability and biases in human judgments. We propose a simple yet effective framework to align LLM judgments with individual human evaluators or their aggregated judgments, without retraining or fine-tuning the LLM. Our approach learns a linear mapping between the LLM’s outputs and human judgments, achieving over 142% average improvement in agreement across 29 tasks with only a small number of calibration examples used for training. Notably, our method works in zero-shot and few-shot settings, exceeds inter-human agreement on four out of six tasks, and enables smaller LLMs to achieve performance comparable to that of larger models.

1 Introduction

Recent improvements to the reasoning capabilities of large language models (LLMs) have increased their use for judgment and evaluation tasks that would previously have been addressed with human evaluators (Zheng et al., 2023b; Chiang and Lee, 2023). For example, LLMs have been used for judging the relevance in information retrieval systems (Faggioli et al., 2023), the coherence of written discourse (Naismith et al., 2023), or the quality of translations (Kocmi and Federmann, 2023), among many others. A common theme in these

settings is that LLMs are performing a grading task where the output label is on an ordinal scale. However, recent work has shown that the performance of LLMs on such judgment tasks varies widely depending on the specific task and the specific LLM used (Bavaresco et al., 2024).

We argue that the performance of LLMs on such judgment tasks does not automatically align well with that of human evaluators. Moreover, LLMs can exhibit their own response styles on these judgment tasks, which can be overly positive, in part due to the nature of supervised fine-tuning during their training regime. Indeed, as we show in Figure 1a, LLMs can avoid negative judgments entirely on tasks where humans show more diverse judgments. Thus, an additional step is needed to align the judgment labels of the LLM with that of a human evaluators. In this work, we propose a simple but effective approach to align the LLM outputs with those of human annotators using only a small set of labeled data. We show on 29 tasks that this additional alignment step significantly improves the performance of LLM judgments on the test set. Figure 1b clearly indicates that the judgment distributions of LLMs becomes much closer to that of the human evaluators after our alignment approach.

Specifically, we propose to learn a mapping of the LLM outputs to human judgments to better align with those on a training set. In contrast to previous work on LLM calibration (Zhao et al., 2021; Han et al., 2023; Fei et al., 2023; Zhou et al., 2024), our technique is applicable without access to the model logits, which makes it much easier to use with general-purpose black-box LLMs available to the public. Moreover, we show that with our alignment approach we can bring the performance of smaller LLMs on par with that of larger ones, which can significantly reduce the cost of using LLM judges at scale.

Our key contribution is a simple yet effective approach to learn a linear mapping between a lan-

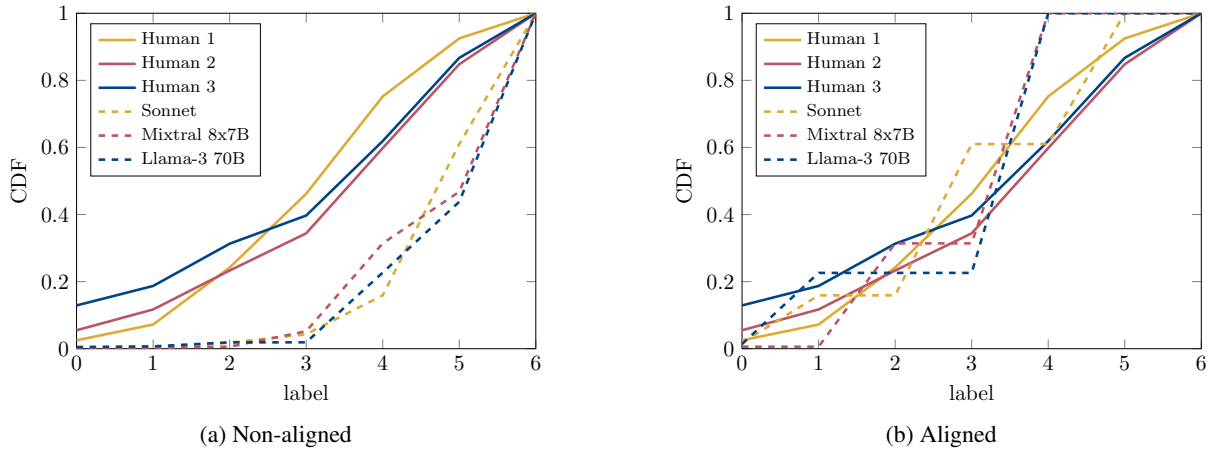


Figure 1: Cumulative distribution function of response options for the judgment task of Freitag et al. (2021), before and after label alignment. The task is to grade translation quality on a rating scale from 0 to 6, i.e., *nonsense* (0), *some meaning preserved* (2), *most meaning preserved* (4), and *perfect* (6). Figure (a) illustrates the different response styles by human and LLM judges and highlights that LLMs primarily use highly positive labels, in contrast to human evaluators. Figure (b) shows the same graph after aligning the LLM responses to the average human judgment using our approach, and clearly demonstrates that we can align LLM judgments to human ones.

guage model’s categorical outputs and human judgments, enabling the alignment of the language model’s judgments with individual human evaluators or aggregated majority votes. Our main contributions can be listed as follows:

- We demonstrate that zero-shot and few-shot LLM judgments, without alignment, exhibit significant discrepancies compared to human evaluations, highlighting the need for additional calibration.
- We propose a framework to align black-box LLM judgments with human evaluations on subjective tasks using a linear mapping, requiring no tuning or retraining of the LLM, and no access to model logits.
- Our method effectively aligns LLM outputs to human judgments using a small number of calibration examples, achieving a 142% average improvement across 29 tasks. Furthermore, our approach benefits smaller LLMs, enabling their use for evaluation tasks with performance comparable to larger models.
- The alignment outperforms inter-human consistency on four out of six tasks that involve multiple human evaluators.

The remainder of this paper is organised as follows. In Section 2 we introduce our black-box approach for LLM alignment, and in Section 3 we

evaluate this approach through extensive and detailed experiments on a large collection of datasets. Next, in Section 4 we review related work. Finally, in Sections 5 and 6 we discuss our contributions and highlight the limitations of our approach along with suggestions for future work.

2 Methodology

Our approach aims to align the judgments provided by an LLM with those given by human evaluators, even when their respective output spaces differ. This alignment process is done without modifying the underlying LLM, making it efficient and easy to apply in practice.

2.1 Problem Formulation

Consider a judgment task where, for a given input instance x_i , the goal is to predict an output label $y_i \in \mathcal{Y}$, where \mathcal{Y} is the set of possible labels given by human evaluators. For instance, in a book review evaluation task, human judgments may take on values from $\mathcal{Y} = \{bad, good, average\}$, implying $|\mathcal{Y}| = 3$. We assume we have a small training dataset $\{(x_i, y_i)\}_{i=1}^N$ containing instances x_i and their corresponding human judgments y_i .

On each task instance x_i we prompt an LLM to provide its own judgment, potentially with a more nuanced set of options, denoted by \mathcal{Z} . For example, the LLM might use labels $\mathcal{Z} = \{bad, neutral, good, excellent\}$ implying $|\mathcal{Z}| = 4$ or stars between 1 and 5 implying $|\mathcal{Z}| = 5$. Thus,

the output spaces \mathcal{Y} and \mathcal{Z} may differ both in size and interpretation, which allows additional flexibility in how the LLM is used while keeping the same human evaluations.

2.2 One-hot Encoding Representation

To facilitate alignment, we convert both human and LLM judgments into one-hot encoded vectors. Specifically, for a judgment $y_i \in \mathcal{Y}$ by a human evaluator, we represent it as a one-hot vector $\mathbf{y}_i \in \{0, 1\}^n$, where $n = |\mathcal{Y}|$. Similarly, for a judgment $z_i \in \mathcal{Z}$ by the LLM, we have $\mathbf{z}_i \in \{0, 1\}^m$, with $m = |\mathcal{Z}|$. For a dataset of N instances, we denote the matrices of one-hot representations as:

$$\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N]^\top \in \{0, 1\}^{N \times n}$$

$$\mathbf{Z} = [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N]^\top \in \{0, 1\}^{N \times m}.$$

2.3 Learning the Mapping

Our goal is to learn a mapping $\varphi : \mathcal{Z} \rightarrow \mathcal{Y}$, from the LLM’s output space \mathcal{Z} to the human judgment space \mathcal{Y} , such that the transformed LLM judgments align closely with human labels. To achieve this, we define a linear transformation matrix $\mathbf{W} \in \mathbb{R}^{m \times n}$, which maps the one-hot encoded LLM judgments to the human ones. The optimal matrix $\widehat{\mathbf{W}}$ is obtained by solving the following regularized least squares problem:

$$\widehat{\mathbf{W}} = \arg \min_{\mathbf{W}} \|\mathbf{Z}\mathbf{W} - \mathbf{Y}\|_F^2 + \lambda \|\mathbf{W}\|_F^2, \quad (1)$$

where $\|\cdot\|_F$ denotes the Frobenius norm, and $\lambda > 0$ is a regularization parameter to prevent overfitting and to avoid a singular system, as in ridge regression (Hoerl and Kennard, 1970). The closed-form solution to this optimization problem is given by:

$$\widehat{\mathbf{W}} = (\mathbf{Z}^\top \mathbf{Z} + \lambda \mathbf{I})^{-1} \mathbf{Z}^\top \mathbf{Y}, \quad (2)$$

where $\mathbf{I} \in \{0, 1\}^{m \times m}$ is the identity matrix.

2.4 Aligned Judgment Inference

Once we have the learned transformation matrix $\widehat{\mathbf{W}}$, we can align one-hot representation of an LLM’s judgment \mathbf{z} for a new instance by:

$$\varphi(\mathbf{z}) = \arg \max_{j=1, \dots, n} \{\mathbf{z}^\top \widehat{\mathbf{W}}\}_j, \quad (3)$$

where $\{\mathbf{z}^\top \widehat{\mathbf{W}}\}_j$ represents the transformed value for each human judgment. The aligned LLM judgment is determined by selecting the human judgment corresponding to the highest transformed

value, reflecting the human judgment that the alignment process associates most closely with the LLM’s original output.

This approach allows us to flexibly map LLM judgments to human judgments, regardless of differences in output space sizes or interpretations. For example, if the LLM output space consists of 4 labels while the human output space has only 3, the learned transformation matrix $\mathbf{W} \in \mathbb{R}^{4 \times 3}$ informs a projection that aligns LLM outputs with human annotations. This flexibility remains crucial even when the output spaces of LLMs and humans are identical, as the interpretation of the same labels can differ significantly, particularly in highly subjective judgment tasks.

3 Experiments

To evaluate the effectiveness of our proposed alignment approach, we conduct a series of experiments across 29 tasks with three widely-adopted large language models: Claude-3 Sonnet (Anthropic, 2024), Mixtral 8x7B Instruct (Jiang et al., 2024), and Llama-3 70B Instruct (Dubey et al., 2024). With the exception of the Feedback-QA dataset (Li et al., 2022), each of the datasets and prompts used to query LLMs are obtained from Judge-Bench, a benchmark dataset proposed recently by Bavaresco et al. (2024) to evaluate LLMs as judges. The datasets we consider from Judge-Bench include LLMBar (Zeng et al., 2024), Medical-Safety (Abercrombie and Rieser, 2022), Newsroom (Grusky et al., 2018), SummEval (Liu et al., 2023), WMT-20-EnDe (Freitag et al., 2021), and the ROSCOE metrics (Golovneva et al., 2023) on COSMOS-QA (Huang et al., 2019), DROP (Dua et al., 2019), ESNLI (Camburu et al., 2018), and GSM8K (Cobbe et al., 2021). Since many datasets have multiple metrics for evaluation, this results in a total of 29 tasks.

In our evaluations, we primarily focused on the case where both LLMs and human evaluators operate within the same judgment space, implying that $\mathcal{Z} = \mathcal{Y}$. This setup allows us to directly assess and discuss the impact of alignment on the agreement between LLM and human judgments. Notably, if $\mathcal{Z} \neq \mathcal{Y}$, the agreement metric would not be straightforward to calculate without an alignment step. In our experiments, we set $\lambda = 10^{-6}$ in eq. (2), as it is primarily used to avoid issues with matrix inversion.

Our primary objective was to assess the agree-

Task Name	HUMAN	Non-aligned			Aligned		
		CLAUDE-3 SONNET	MIXTRAL 8x7B INSTRUCT	LLAMA-3 70B INSTRUCT	CLAUDE-3 SONNET	MIXTRAL 8x7B INSTRUCT	LLAMA-3 70B INSTRUCT
Feedback-QA	44.50	42.11 ±2.16	43.06 ±1.57	47.38 ±2.10	52.68 ±1.97	52.29 ±2.12	51.67 ±1.99
LLMBar Natural	–	<u>85.07</u> ±1.67	<u>77.20</u> ±2.19	86.52 ±1.21	<u>85.07</u> ±1.67	<u>77.20</u> ±2.19	86.52 ±1.21
Medical Safety							
Query Risk Level	–	39.67 ±1.15	43.20 ±1.30	21.12 ±1.80	<u>85.00</u> ±1.47	85.30 ±1.57	<u>85.22</u> ±1.45
Response Type	–	5.63 ±0.52	11.03 ±1.31	5.63 ±1.00	<u>79.92</u> ±1.87	<u>68.14</u> ±3.23	<u>75.08</u> ±1.80
Newsroom							
Coherence	24.29	27.87 ±0.96	26.11 ±0.50	31.94 ±0.58	<u>32.24</u> ±1.15	<u>30.60</u> ±1.98	32.73 ±2.12
Fluency	21.35	<u>28.03</u> ±0.65	20.27 ±0.82	30.70 ±0.70	25.88 ±1.44	24.01 ±2.07	29.83 ±0.87
Informativeness	31.75	<u>32.77</u> ±0.86	15.53 ±1.21	35.01 ±0.91	<u>38.32</u> ±1.50	<u>29.94</u> ±3.18	38.97 ±2.14
Relevance	30.71	30.96 ±0.93	21.24 ±0.58	33.58 ±0.68	<u>36.04</u> ±1.79	<u>35.57</u> ±1.68	38.42 ±2.34
ROSCOE-CosmosQA							
Coherency	–	42.65 ±1.90	20.41 ±1.61	36.92 ±1.95	<u>47.96</u> ±3.69	42.99 ±5.57	53.17 ±4.43
Contradiction	–	68.64 ±2.54	57.20 ±1.22	61.55 ±2.17	78.03 ±0.96	<u>77.94</u> ±0.98	<u>77.97</u> ±0.99
Missing Steps	–	<u>59.12</u> ±1.76	<u>56.73</u> ±1.37	61.72 ±2.20	<u>55.99</u> ±2.56	<u>55.31</u> ±1.72	60.81 ±2.50
Overall Quality	–	<u>29.86</u> ±2.38	<u>18.50</u> ±1.48	<u>33.47</u> ±1.98	<u>43.47</u> ±2.48	<u>31.90</u> ±2.84	44.66 ±3.60
ROSCOE-DROP							
Coherency	–	72.66 ±1.98	34.68 ±2.40	67.56 ±1.75	77.59 ±2.20	<u>77.22</u> ±1.60	<u>75.67</u> ±1.71
Contradiction	–	88.16 ±0.70	71.24 ±1.39	86.48 ±0.96	<u>93.35</u> ±1.84	<u>94.03</u> ±0.65	94.54 ±0.74
Missing Steps	–	<u>51.46</u> ±1.70	<u>51.58</u> ±1.34	55.44 ±1.59	50.38 ±1.07	48.68 ±2.26	<u>51.49</u> ±3.53
Overall Quality	–	43.42 ±2.14	<u>40.89</u> ±1.82	43.53 ±2.20	<u>43.67</u> ±8.60	39.62 ±2.06	46.75 ±3.26
ROSCOE-ESNLI							
Coherency	–	84.65 ±1.19	39.39 ±2.06	73.59 ±1.80	<u>86.49</u> ±1.58	<u>89.04</u> ±0.90	89.31 ±2.27
Contradiction	–	93.25 ±0.96	50.57 ±2.37	93.46 ±0.99	<u>95.26</u> ±2.19	96.55 ±0.84	<u>94.16</u> ±1.78
Missing Steps	–	67.98 ±1.97	27.28 ±1.38	64.26 ±1.47	74.21 ±1.72	<u>73.16</u> ±1.43	<u>72.80</u> ±3.23
Overall Quality	–	48.07 ±2.63	<u>38.86</u> ±2.26	<u>47.55</u> ±2.73	45.00 ±6.35	36.49 ±6.93	43.69 ±6.75
ROSCOE-GSM8K							
Coherency	–	59.73 ±1.84	47.50 ±2.70	63.89 ±1.83	<u>62.07</u> ±1.31	<u>60.35</u> ±3.32	62.94 ±1.79
Contradiction	–	80.67 ±0.94	55.30 ±2.37	80.53 ±1.42	<u>83.60</u> ±1.58	<u>83.32</u> ±1.61	83.81 ±1.90
Missing Steps	–	<u>81.13</u> ±1.46	42.77 ±0.77	84.09 ±1.26	<u>81.13</u> ±1.46	<u>58.89</u> ±1.97	84.09 ±1.26
Overall Quality	–	71.60 ±2.15	49.59 ±2.37	65.46 ±2.03	<u>72.53</u> ±2.71	<u>61.61</u> ±3.42	73.82 ±2.86
SummEval							
Coherence	–	31.47 ±1.71	28.79 ±2.06	40.90 ±1.48	40.20 ±1.69	31.64 ±2.45	39.87 ±2.44
Consistency	–	83.64 ±1.53	68.84 ±1.80	68.82 ±2.53	85.38 ±1.55	<u>84.44</u> ±1.80	<u>77.83</u> ±4.43
Fluency	–	4.50 ±0.83	3.33 ±0.58	2.58 ±0.33	<u>79.73</u> ±1.54	<u>79.70</u> ±1.35	80.22 ±1.70
Relevance	–	37.27 ±1.53	31.27 ±1.72	41.34 ±1.40	55.07 ±4.67	<u>51.30</u> ±3.72	<u>54.17</u> ±3.01
WMT-20-EnDe	30.02	19.48 ±1.45	17.58 ±1.58	17.15 ±1.45	26.70 ±1.80	<u>25.05</u> ±1.33	<u>23.76</u> ±1.37

Table 1: Summary table showing the accuracy (in %) of aligning LLMs to human labels. The first column shows inter-human agreement for the tasks with more than one human annotator. For the non-aligned results the label predicted by the LLM is used directly. For each task, 25% labeled samples are used for alignment and 75% for evaluation, and accuracy is averaged over each individual human annotator. Experiments are repeated 10 times with different training and evaluation splits and results are averaged. **Bold** indicates the best performance in the each row and underline indicates the best performance of each LLM (aligned or non-aligned).

ment between human judgments and those generated by the LLMs. We treated human judgments as the ground truth and calculated the accuracy of the LLM outputs relative to these targets. Additionally, for task datasets with multiple human evaluators, we analyzed the inter-human agreement. In our experimental setup, for each task, we used randomly selected 100 examples for training the alignment model and 300 examples for testing the effectiveness of the learned alignment. Because the labels differ between tasks, alignment mappings are learned on a per-task basis (we study transferring alignments between tasks in Section 3.3). In some tasks, the total number of examples are below

400; so we use 25% of the task dataset as training and 75% for testing. We repeated our experiments 10 times with different random splits of training and test sets, and report the mean and standard deviation.

3.1 Zero-shot Judgments

In Table 1, we present the accuracy of LLM judgments compared to human ones, both before and after the alignment process. The results indicate a significant increase in agreement post-alignment, demonstrating the effectiveness of our approach in enhancing LLM performance in alignment tasks.

Overall, our approach significantly improves

LLM-to-human agreement or maintains it at the same level. There are only a few instances where alignment results in a slight reduction in agreement. On average, across all 29 tasks, alignment leads to a remarkable 142% increase in agreement. More specifically, we have achieved 116.52% increase for Claude-3 Sonnet, 142.91% for Mixtral 8x7B, and 166.6% for Llama-3 70B models. Note that we observe identical results before and after alignment for some tasks or models, because our alignment results in the identity mapping in these cases.

For specific tasks, such as *Medical Safety (response type)* and *Summeval (fluency)*, the agreement increases from single-digit percentages up to 80% after alignment. For example, we observed a significant discrepancy between the LLM judgments and evaluations made by a professional nurse in the medical safety task. While the nurse often selected “not medical” as the judgment for many cases, the LLMs tended to choose “non-serious”. Such misalignment not only led to lower agreement scores but also highlighted the potential risks of employing LLMs as judges in critical domains, such as medical assessments. Our alignment method effectively learned to map the LLM’s “non-serious” outputs to “not medical” judgments of the nurses. As a result of these learned mappings, we improved the agreement between LLMs and nurses from 5-11% up to 80% for this task.

In our experiments, we used up to 100 training examples to learn the alignment. However, we observed that our approach works quite well with much smaller numbers of training examples across all datasets. Figure 2 shows the test accuracy of our approach for the *Medical Safety (response type)* task for different numbers of training examples per judgment category. The figure indicates that our approach achieves alignment and significantly improves the agreement between the LLMs and human judgments using only one or two examples per category. This highlights the sample efficiency of our approach, which is in part due to its simplicity.

3.2 Judgments with In-Context Learning

Our findings in the zero-shot setting motivated us to investigate whether the sources of misalignment between LLMs and humans could be addressed through in-context learning (Brown et al., 2020), by providing examples of human judgments for each output label. Specifically, for each human judgment label we include in the LLM prompt an example on which human evaluators agreed, aim-




Figure 2: Test accuracy for *Medical Safety (response type)* dataset as we increase the number of training examples per judgment category.

ing to guide the LLM towards better alignment with human interpretations.

As shown in Table 2, the results on this experiment indicate that merely providing these examples did not lead to significantly improved agreement between the LLMs and human evaluators. That is, the in-context learning barely improved the agreement between human judgments and LLMs, except a few cases such as *Summeval (fluency)*, where the agreement with human evaluators improved from 4.5% to 35.57% for *Claude-3 Sonnet*. The average agreement improved only 0.8% for *Claude-3 Sonnet* and 1.14% for *Mixtral 8x7B* with very high standard deviations (7.19% and 19.74%, respectively). We observe 1.4% decrease in the average performance of *Llama-3 70B* with 8.45% standard deviation. One potential reason for poorer performance can be the increased complexity of the prompts with the added examples, which can become too long for the model to effectively follow the instructions or even approach the model’s maximum context length.

When our alignment approach is used on top of in-context learning, we observe 111.74% average improvement with respect to non-aligned model agreement with human judgments. More specifically, we have achieved 49.65% increase for *Claude-3 Sonnet*, 180.13% for *Mixtral 8x7B*, and 105.45% for *Llama-3 70B* models. That is, our alignment approach consistently results in significant improvements in LLMs’ agreement with human judgments, in both zero-shot and in-context learning settings.

These findings underscore the necessity of a structured alignment process. The performance en-

Task Name	Non-aligned			Aligned		
	CLAUDE-3 SONNET	MIXTRAL 8x7B INSTRUCT	LLAMA-3 70B INSTRUCT	CLAUDE-3 SONNET	MIXTRAL 8x7B INSTRUCT	LLAMA-3 70B INSTRUCT
Feedback-QA	43.29 ±1.01	40.28 ±0.82	45.17 ±1.61	<u>52.63</u> ±0.91	50.63 ±1.35	53.43 ±1.34
LLMBar Natural	84.67 ±2.09	<u>71.83</u> ±2.86	<u>81.11</u> ±1.89	84.67 ±2.09	<u>71.83</u> ±2.86	<u>81.11</u> ±1.89
Medical Safety						
Query Risk Level	46.73 ±1.53	32.32 ±1.46	23.90 ±1.03	<u>85.73</u> ±1.03	<u>85.21</u> ±2.73	86.01 ±1.32
Response Type	8.62 ±0.84	11.61 ±1.07	10.20 ±0.53	76.80 ±2.48	<u>69.97</u> ±2.16	<u>76.49</u> ±2.56
Newsroom						
Coherence	25.63 ±1.05	18.47 ±0.82	25.04 ±0.73	<u>30.22</u> ±0.94	31.12 ±1.05	<u>30.07</u> ±1.08
Fluency	<u>26.60</u> ±0.55	20.21 ±0.91	29.62 ±0.56	24.88 ±1.50	<u>24.32</u> ±1.71	<u>27.35</u> ±1.28
Informativeness	28.70 ±0.70	18.94 ±1.27	27.87 ±1.07	<u>37.68</u> ±1.69	<u>31.48</u> ±1.71	37.84 ±1.80
Relevance	30.64 ±1.28	14.25 ±0.58	32.95 ±0.85	<u>35.27</u> ±2.74	<u>35.71</u> ±1.03	36.52 ±1.14
ROSCOE-CosmosQA						
Coherency	43.91 ±1.53	53.74 ±1.10	42.19 ±2.46	<u>53.49</u> ±2.26	54.49 ±2.78	<u>46.26</u> ±1.80
Contradiction	77.28 ±1.78	47.76 ±2.08	57.90 ±1.72	<u>77.41</u> ±1.72	78.50 ±1.47	<u>78.09</u> ±1.46
Missing Steps	<u>55.44</u> ±1.70	<u>56.94</u> ±1.36	67.62 ±1.62	54.49 ±1.27	55.71 ±1.56	67.62 ±1.62
Overall Quality	31.54 ±2.80	33.74 ±2.16	42.88 ±2.38	<u>44.60</u> ±2.14	<u>45.44</u> ±5.12	45.80 ±3.74
ROSCOE-DROP						
Coherency	71.33 ±1.47	74.49 ±1.60	63.49 ±1.90	<u>74.43</u> ±2.18	76.33 ±1.33	<u>73.91</u> ±3.80
Contradiction	91.39 ±0.95	43.42 ±1.63	75.63 ±1.72	<u>92.59</u> ±1.75	93.92 ±0.42	<u>93.81</u> ±0.43
Missing Steps	<u>53.86</u> ±1.73	<u>54.18</u> ±1.88	61.77 ±1.30	51.14 ±1.67	48.10 ±2.77	60.24 ±4.37
Overall Quality	45.76 ±1.44	42.34 ±1.40	44.88 ±1.68	47.53 ±2.32	<u>43.04</u> ±2.10	<u>45.84</u> ±2.20
ROSCOE-ESNLI						
Coherency	83.68 ±1.67	86.32 ±1.05	77.81 ±1.11	88.42 ±2.51	<u>87.28</u> ±1.43	<u>88.07</u> ±3.31
Contradiction	95.44 ±1.23	55.26 ±2.08	83.51 ±1.51	<u>96.67</u> ±1.46	96.93 ±1.13	96.93 ±1.13
Missing Steps	70.79 ±1.84	27.72 ±2.12	65.07 ±2.40	<u>71.05</u> ±1.80	72.28 ±2.12	<u>70.09</u> ±4.37
Overall Quality	48.07 ±1.87	45.00 ±1.67	<u>45.66</u> ±1.64	48.51 ±2.36	<u>45.35</u> ±2.32	<u>45.40</u> ±2.38
ROSCOE-GSM8K						
Coherency	63.60 ±1.55	59.07 ±2.24	59.80 ±2.11	66.93 ±2.13	<u>61.20</u> ±2.23	<u>63.87</u> ±2.75
Contradiction	84.53 ±1.81	45.13 ±2.77	77.09 ±1.73	84.53 ±1.81	84.53 ±1.81	<u>81.18</u> ±3.83
Missing Steps	<u>74.87</u> ±1.52	54.67 ±2.13	94.31 ±0.76	<u>74.87</u> ±1.52	<u>57.47</u> ±2.32	94.31 ±0.76
Overall Quality	72.00 ±1.23	71.47 ±1.29	<u>74.31</u> ±1.16	75.20 ±3.66	<u>71.67</u> ±1.67	<u>72.43</u> ±2.78
SummEval						
Coherence	25.03 ±1.41	24.87 ±2.01	27.33 ±1.61	37.90 ±2.24	31.93 ±2.92	34.40 ±2.29
Consistency	75.03 ±1.38	12.00 ±1.20	36.61 ±1.22	87.13 ±1.72	83.06 ±5.03	<u>86.71</u> ±1.56
Fluency	35.57 ±0.78	8.93 ±0.79	4.55 ±0.75	<u>82.47</u> ±1.14	<u>81.30</u> ±3.34	83.30 ±1.59
Relevance	25.57 ±1.13	1.87 ±0.85	37.79 ±1.55	55.33 ±0.95	42.76 ±2.63	<u>52.48</u> ±3.21
WMT-20-EnDe	15.25 ±0.84	16.09 ±0.82	24.76 ±0.68	28.42 ±1.75	<u>27.33</u> ±1.22	<u>25.20</u> ±1.17

Table 2: In-Context Learning experiments, where one example per judgment category is given to LLMs as in-context examples. Experiments are repeated 10 times with different training and evaluation splits and results are averaged. **Bold** indicates the best performance in the each row and underscore indicates the best performance of each LLM (aligned or non-aligned).

	Zero-shot	In-Context Learning
Claude-3 Sonnet	79.31%	79.31%
Mixtral 8x7B Instruct	82.75%	89.65%
Llama-3 70B Instruct	72.41%	75.86%

Table 3: Percentage of tasks in which the proposed approach improves the agreement with human judgments.

hancements achieved through our proposed method surpass those observed with traditional in-context learning strategies. This suggests that while in-context examples can provide useful context for LLMs, in our experiments they do not substitute for a dedicated alignment mechanism to effectively bridge the gap between LLM judgments and human evaluations.

In summary, our experiments highlight that aligning LLM judgments with human assessments is a complex task. Our proposed alignment methodology offers a promising solution despite its simplicity, even in scenarios where in-context learning is employed. Table 3 lists the percentage of tasks in which our approach improves the level of agreement with human judgments after alignment. It indicates that we were able to significantly improve agreement between LLMs and the human judgments with our approach in both zero-shot and in-context learning settings. One of our most interesting and important findings is that our alignment approach has enabled smaller models like *Mixtral 8x7B Instruct* to match the performance of much larger models such as *Claude-3 Sonnet*. This is a signifi-

	Claude-3 Sonnet	Mixtral 8x7B Instruct	Llama-3 70B Instruct
DROP	76.62	65.98	75.29
ESNLI	86.20	78.07	85.13
GSM8K	65.11	51.40	65.79

Table 4: Accuracy (in %) using an alignment learned on the coherency task of Roscoe-Cosmos and tested on the other Roscoe coherency tasks. As in Table 1, 25% of training data was used to learn the mapping and results were averaged 10 times with different training data splits.

cant result, as it opens up new opportunities to use these smaller, more affordable and openly accessible language models for evaluation tasks, while still achieving similar performance to large, proprietary API-based models like those in the Claude family.

3.3 Transferring alignments

For judgment tasks that use grading scales with the same number of categories and the same meaning it is possible to evaluate whether an alignment learned on one task can be used on a different task. To test this, we learn an alignment on the coherency task of the Roscoe-Cosmos dataset, and test it on the coherency task of the other Roscoe datasets (each of which evaluates the coherence of GPT-3’s response on reasoning tasks). The results in Table 4 indicate that using the transferred alignment gives better performance than the non-aligned LLM judges in Table 1, and is only slightly worse than when using a task-specific alignment.

4 Related Work

The use of large language models as automated judges has been gaining traction due to their potential for scalability and efficiency (Zheng et al., 2023b; Bavaresco et al., 2024). However, as these models are increasingly relied upon to evaluate recommendations, search results, and other subjective tasks, it is crucial that their judgments align closely with human evaluators to ensure evaluations remain human-centered and useful.

Gilardi et al. (2023) investigated the capabilities of ChatGPT compared to human annotators across tweet annotation tasks including *relevance*, *stance*, *topics*, and *frame detection*. The results show that ChatGPT’s zero-shot accuracy exceeds that of crowd workers and echoes the findings on annotating political affiliation in tweets by Törn-

berg (2023). The cost-effectiveness of ChatGPT underscores its potential to drastically improve the efficiency of text classification tasks. Despite of these promising initial results, recent work has highlighted the challenges in relying on LLMs as judges due to biases and inconsistencies in their judgments (Wu and Aji, 2023; Zheng et al., 2023a; Koo et al., 2024; Hada et al., 2024; Pavlovic and Poesio, 2024). Wang et al. (2023) demonstrated that the quality ranking of candidate responses can be manipulated by altering their order of appearance, allowing one model to seem superior to another based on positional biases. This highlights the susceptibility of LLM evaluations to manipulation and underscores the need for effective calibration to ensure fairness and reliability.

Bavaresco et al. (2024) proposed Judge-Bench, a benchmark consisting of over 20 NLP datasets with human annotations, to assess the effectiveness of LLMs as judges. Their evaluations of 11 different LLMs, including both open-weight and proprietary models, reveal significant variance in correlation to human judgments across datasets. They conclude that LLMs, despite their growing usage in evaluation settings, are not yet ready to systematically replace human judges. In this paper, we also used the datasets provided by this benchmark to demonstrate the effectiveness of our approach in improving the alignment of the LLMs with the human judgments.

The calibration of response styles has a long history in the psychology literature, where survey respondents can exhibit biases towards the central or extreme categories on a rating scale (Jackson and Messick, 1958; Paulhus, 1991). For such settings with a large number of respondents and heterogeneous survey questions, techniques have been developed for identifying and removing response style bias (see, e.g., Van Rosmalen et al., 2010; Schoonees et al., 2015). This echoes the observations reported in this work, where we have shown that LLMs exhibit their own response styles, for instance by preferring highly positive classes (see Figure 1). At the same time, Huang et al. (2024) demonstrated that on coarse-grained assessments, such as binary assessments, ChatGPT’s evaluations closely approximate human ones, but that it struggles in fine-grained assessments. On the other hand, Tjuatja et al. (2024) found that LLMs do not exhibit human-like response style biases caused by the *wording* of a prompt.

When the logits of language models are accessi-

ble, they can be calibrated directly using gradient descent as in model fine-tuning or low-rank adaptation (Reif and Schwartz, 2024). Zhao et al. (2021) showed that LLMs are highly sensitive to prompt structure, resulting in inconsistent performance, and proposed a calibration procedure involving the use of a content-free input to estimate and adjust for these biases. Such findings highlight the importance of addressing LLM biases, especially in contexts where alignment with human judgments is critical. Unlike the above approaches, our method does not require access to model weights or logits, so it is easier to adapt to a wider range of use cases and scenarios.

5 Discussion

Our method provides a simple yet powerful framework to align LLM judgments with human evaluations, offering substantial improvements without the need for model retraining or fine-tuning. By learning linear mappings between LLM outputs and human judgments, we achieve significant improvement in agreement, a result that is both significant and practical for real-world applications where human labeled data is scarce. Furthermore, in four out of six tasks where there are multiple human annotators, the LLM performance after our alignment exceeds the inter-human agreement.

The simplicity and scalability of our approach is a key advantage. With only minimal calibration examples, we allow smaller models to perform at levels comparable to much larger, more resource-intensive LLMs. This shows that our method may greatly extend the usefulness of smaller models in judgment tasks, opening up opportunities for their deployment in a wider range of applications. This is especially valuable in resource-constrained settings, where using large models may be impractical or too costly.

Although this extension is beyond the scope of the current paper, we note that when model logits are available, the one-hot encoded vectors representing LLM judgments can be replaced with the judgment probabilities derived from logits values. This modification could enable potentially finer alignment between LLM outputs and human evaluations. Furthermore, our alignment method can be extended to integrate other supervised learning techniques or objectives for handling more complex tasks, or to condition alignment on model inputs for improved performance where necessary. We

leave the exploration of these directions to future work, as they offer promising avenues for enhancing model alignment and performance, particularly in more complex or input-conditioned tasks.

Additionally, our approach lays the groundwork for multi-model alignment, enabling the learned mappings to be adapted for improving agreement across different LLMs. This adaptation could facilitate the transfer of knowledge between models, enhancing their reliability in judgment tasks. For instance, we could explore concatenating the outputs of multiple models and mapping them to human labels, effectively combining the strengths of various language models. Such a strategy would not only leverage diverse insights but also improve overall performance in alignment tasks. Related work, such as that proposed by Verga et al. (2024), investigates multi-model alignment strategies that could complement and extend our findings, opening avenues for more robust evaluation frameworks.

Finally, our work allows LLMs to utilize different judgment spaces than human evaluators. However, we leave the exploration of this aspect to future work. Enabling LLMs to have judgment spaces that differ from human evaluators offers several advantages. We can employ prompt optimization techniques to discover the optimal judgment space for an LLM. This allows us to align the model’s outputs with human judgments, thereby achieving better agreement. In this way, we can search for the best judgment space for the model while preserving the human judgments as they are.

6 Limitations

While our proposed alignment approach demonstrates significant improvements in aligning LLM judgments with human evaluations, there are some limitations to consider.

First, the effectiveness of our method may vary depending on the specific characteristics of the domain and tasks involved. Some domains may present more complex alignment challenges and require more complicated alignment approaches, which may need larger number of training examples as a potential downside.

Second, while we compare our approach with in-context learning, we did not comprehensively explore the prompt space to see if we can improve the agreement and alignment through prompt optimisation. Instead, we used the prompts provided along with the task datasets from Judge-Bench

(Bavaresco et al., 2024) to demonstrate the improvement in agreement through our approach. We believe prompt optimisation is orthogonal to our approach and both can be used at the same time. Furthermore, the flexibility of our approach for using different judgment spaces than human evaluators expands the exploration space for the prompt optimisation and may lead to better results. Although we did not explicitly test these directions in our current work, they present an exciting opportunity for future research.

Despite these considerations, our findings provide a solid foundation for future research and practical applications, demonstrating the potential of our simple alignment method to enhance LLM as a judge performance across various domains.

References

- Gavin Abercrombie and Verena Rieser. 2022. [Risk-graded safety for handling medical queries in conversational AI](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 234–243, Online only. Association for Computational Linguistics.
- Anthropic. 2024. The Claude 3 model family: Opus, Sonnet, Haiku. <https://anthropic.com/claude-3-model-card>. Retrieved Oct 2024.
- Anna Bavaresco, Raffaella Bernardi, Leonardo Bertolazzi, Desmond Elliott, Raquel Fernández, Albert Gatt, Esam Ghaleb, Mario Giulianelli, Michael Hanna, Alexander Koller, André F. T. Martins, Philipp Mondorf, Vera Neplenbroek, Sandro Pezzelle, Barbara Plank, David Schlangen, Alessandro Sgulia, Aditya K Surikuchi, Ece Takmaz, and Alberto Testoni. 2024. [LLMs instead of human judges? A large scale empirical study across 20 NLP evaluation tasks](#). *Preprint*, arXiv:2406.18403.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Ma teusz 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](#). *Preprint*, ArXiv:2005.14165.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. [e-SNLI: Natural language inference with natural language explanations](#). In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- 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.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. [DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2368–2378, Minneapolis, Minnesota. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. [The Llama 3 herd of models](#). *Preprint*, ArXiv:2407.21783.
- Guglielmo Faggioli, Laura Dietz, Charles L.A. Clarke, Gianluca Demartini, Matthias Hagen, Claudia Hauff, Noriko Kando, Evangelos Kanoulas, Martin Potthast, Benno Stein, and Henning Wachsmuth. 2023. [Perspectives on large language models for relevance judgment](#). In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval*, pages 39–50.
- Yu Fei, Yifan Hou, Zeming Chen, and Antoine Bosselut. 2023. [Mitigating label biases for in-context learning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14014–14031, Toronto, Canada. Association for Computational Linguistics.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, Errors, and Context: A Large-Scale Study of Human Evaluation for Machine Translation](#). *Transactions of the Association for Computational Linguistics*, 9:1460–1474.
- Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. [ChatGPT outperforms crowd workers for text-annotation tasks](#). *Proceedings of the National Academy of Sciences*, 120(30):e2305016120.
- Olga Golovneva, Moya Peng Chen, Spencer Poff, Martin Corredor, Luke Zettlemoyer, Maryam Fazel-Zarandi, and Asli Celikyilmaz. 2023. [ROSCOE: A suite of metrics for scoring step-by-step reasoning](#). In *The Eleventh International Conference on Learning Representations*.

- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. [Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 708–719, New Orleans, Louisiana. Association for Computational Linguistics.
- Rishav Hada, Varun Gumma, Adrian Wynter, Harshita Diddee, Mohamed Ahmed, Monojit Choudhury, Kalika Bali, and Sunayana Sitaram. 2024. [Are large language model-based evaluators the solution to scaling up multilingual evaluation?](#) In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 1051–1070, St. Julian’s, Malta. Association for Computational Linguistics.
- Zhixiong Han, Yaru Hao, Li Dong, Yutao Sun, and Furu Wei. 2023. [Prototypical calibration for few-shot learning of language models](#). In *The Eleventh International Conference on Learning Representations*.
- Arthur E Hoerl and Robert W Kennard. 1970. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67.
- Fan Huang, Haewoon Kwak, Kunwoo Park, and Jisun An. 2024. [ChatGPT rates natural language explanation quality like humans: But on which scales?](#) In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3111–3132, Torino, Italia. ELRA and ICCL.
- Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [Cosmos QA: Machine reading comprehension with contextual commonsense reasoning](#). 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 2391–2401, Hong Kong, China. Association for Computational Linguistics.
- Douglas N. Jackson and Samuel Messick. 1958. Content and style in personality assessment. *Psychological Bulletin*, 55(4):243.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L’elio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2024. [Mixtral of experts](#). *Preprint*, ArXiv:2401.04088.
- 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.
- Ryan Koo, Minhwa Lee, Vipul Raheja, Jong Inn Park, Zae Myung Kim, and Dongyeop Kang. 2024. [Benchmarking cognitive biases in large language models as evaluators](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 517–545, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Zichao Li, Prakhar Sharma, Xing Han Lu, Jackie CK Cheung, and Siva Reddy. 2022. [Using interactive feedback to improve the accuracy and explainability of question answering systems post-deployment](#). *Preprint*, arXiv:2204.03025.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using GPT-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Ben Naismith, Phoebe Mulcaire, and Jill Burstein. 2023. Automated evaluation of written discourse coherence using GPT-4. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 394–403.
- Delroy L. Paulhus. 1991. [Measurement and control of response bias](#). In John P. Robinson, Phillip R. Shaver, and Lawrence S. Wrightsman, editors, *Measures of Personality and Social Psychological Attitudes*, pages 17–59. Academic Press.
- Maja Pavlovic and Massimo Poesio. 2024. [The effectiveness of LLMs as annotators: A comparative overview and empirical analysis of direct representation](#). In *Proceedings of the 3rd Workshop on Perspectivist Approaches to NLP (NLPerspectives) @ LREC-COLING 2024*, pages 100–110, Torino, Italia. ELRA and ICCL.
- Yuval Reif and Roy Schwartz. 2024. [Beyond performance: Quantifying and mitigating label bias in LLMs](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6784–6798, Mexico City, Mexico. Association for Computational Linguistics.
- Pieter C. Schoonees, Michel Van de Velden, and Patrick J. F. Groenen. 2015. [Constrained dual scaling for detecting response styles in categorical data](#). *Psychometrika*, 80:968–994.
- Lindia Tjuatja, Valerie Chen, Tongshuang Wu, Ameet Talwalkwar, and Graham Neubig. 2024. [Do LLMs Exhibit Human-like Response Biases? A Case Study in Survey Design](#). *Transactions of the Association for Computational Linguistics*, 12:1011–1026.

- Petter Törnberg. 2023. [ChatGPT-4 outperforms experts and crowd workers in annotating political Twitter messages with zero-shot learning](#). *Preprint*, arXiv:2304.06588.
- Joost Van Rosmalen, Hester Van Herk, and Patrick J. F. Groenen. 2010. [Identifying response styles: A latent-class bilinear multinomial logit model](#). *Journal of Marketing Research*, 47(1):157–172.
- Pat Verga, Sebastian Hofstatter, Sophia Althammer, Yixuan Su, Aleksandra Piktus, Arkady Arkhangorodsky, Minjie Xu, Naomi White, and Patrick Lewis. 2024. [Replacing judges with juries: Evaluating LLM generations with a panel of diverse models](#). *Preprint*, arXiv:2404.18796.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023. [Large language models are not fair evaluators](#). *Preprint*, arXiv:2305.17926.
- Minghao Wu and Alham Fikri Aji. 2023. [Style over substance: Evaluation biases for large language models](#). *Preprint*, arXiv:2307.03025.
- Zhiyuan Zeng, Jiatong Yu, Tianyu Gao, Yu Meng, Tanya Goyal, and Danqi Chen. 2024. [Evaluating large language models at evaluating instruction following](#). In *The Twelfth International Conference on Learning Representations*.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. [Calibrate before use: Improving few-shot performance of language models](#). In *Proceedings of the 38th International Conference on Machine Learning*, pages 12697–12706. PMLR.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2023a. [Large language models are not robust multiple choice selectors](#). In *The Twelfth International Conference on Learning Representations*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023b. [Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena](#). *Advances in Neural Information Processing Systems*, 36:46595–46623.
- Han Zhou, Xingchen Wan, Lev Proleev, Diana Mincu, Jilin Chen, Katherine A Heller, and Subhrajit Roy. 2024. [Batch calibration: Rethinking calibration for in-context learning and prompt engineering](#). In *The Twelfth International Conference on Learning Representations*.

A Limited Training Examples

In Table 5, we demonstrate the performance of our approach when the number of training examples are limited to 20 samples while 300 samples are used for testing. We used 25% of the dataset as training examples if the total number of examples in the dataset less than 320 samples. The table indicates that our approach can also work in extremely low data regimes.

B Prompt Example

Figure 3 shows an example of the prompt we used in both our zero-shot and few-shot experiments for the *medical safety (response type)* task.

Task Name	HUMAN	Non-aligned			Aligned		
		CLAUDE-3 SONNET	MIXTRAL 8X7B INSTRUCT	LLAMA-3 70B INSTRUCT	CLAUDE-3 SONNET	MIXTRAL 8X7B INSTRUCT	LLAMA-3 70B INSTRUCT
Feedback-QA	44.50	42.11 ±2.16	43.06 ±1.57	47.38 ±2.10	48.88 ±3.47	48.31 ±3.97	49.19 ±4.26
LLMBar Natural	–	<u>85.07</u> ±1.67	<u>77.20</u> ±2.19	86.52 ±1.21	<u>85.07</u> ±1.67	<u>77.20</u> ±2.19	86.52 ±1.21
Medical Safety							
Query Risk Level	–	39.67 ±1.15	43.20 ±1.30	21.12 ±1.80	84.03 ±2.09	83.93 ±3.36	83.50 ±3.50
Response Type	–	5.63 ±0.52	11.03 ±1.31	5.63 ±1.00	75.22 ±6.03	64.76 ±4.51	<u>72.95</u> ±5.32
Newsroom							
Coherence	24.29	27.87 ±0.96	26.11 ±0.50	31.94 ±0.58	29.22 ±3.63	25.95 ±3.16	28.74 ±3.10
Fluency	21.35	<u>28.03</u> ±0.65	20.27 ±0.82	30.70 ±0.70	22.80 ±2.95	20.82 ±2.37	25.42 ±2.60
Informativeness	31.75	32.77 ±0.86	15.53 ±1.21	35.01 ±0.91	32.90 ±2.81	<u>24.32</u> ±3.21	32.76 ±3.63
Relevance	30.71	<u>30.96</u> ±0.93	21.24 ±0.58	33.58 ±0.68	29.90 ±1.56	<u>26.85</u> ±5.01	32.26 ±3.36
ROSCOE-Cosmos							
Coherency	–	42.65 ±1.90	20.41 ±1.61	36.92 ±1.95	49.12 ±3.07	35.92 ±6.48	<u>48.81</u> ±6.91
Contradiction	–	68.64 ±2.54	57.20 ±1.22	61.55 ±2.17	76.19 ±3.87	77.94 ±0.98	<u>74.19</u> ±7.63
Missing Steps	–	<u>59.12</u> ±1.76	<u>56.73</u> ±1.37	61.72 ±2.20	56.80 ±2.57	54.08 ±4.12	59.40 ±2.40
Overall Quality	–	29.86 ±2.38	18.50 ±1.48	33.47 ±1.98	39.59 ±9.34	30.41 ±4.61	41.62 ±5.83
ROSCOE-DROP							
Coherency	–	72.66 ±1.98	34.68 ±2.40	67.56 ±1.75	75.57 ±2.30	70.95 ±9.04	<u>74.23</u> ±1.44
Contradiction	–	88.16 ±0.70	71.24 ±1.39	86.48 ±0.96	92.03 ±2.96	94.03 ±0.65	<u>93.69</u> ±2.79
Missing Steps	–	<u>51.46</u> ±1.70	<u>51.58</u> ±1.34	55.44 ±1.59	50.76 ±2.23	48.55 ±2.18	48.96 ±4.60
Overall Quality	–	<u>43.42</u> ±2.14	<u>40.89</u> ±1.82	43.53 ±2.20	42.22 ±10.60	36.20 ±8.57	42.68 ±8.48
ROSCOE-ESNLI							
Coherency	–	84.65 ±1.19	39.39 ±2.06	73.59 ±1.80	85.96 ±2.32	88.77 ±0.94	85.77 ±6.95
Contradiction	–	93.25 ±0.96	50.57 ±2.37	93.46 ±0.99	95.09 ±1.63	96.55 ±0.84	<u>95.05</u> ±2.02
Missing Steps	–	67.98 ±1.97	27.28 ±1.38	64.26 ±1.47	<u>72.11</u> ±2.44	72.98 ±1.46	<u>70.68</u> ±4.32
Overall Quality	–	48.07 ±2.63	38.86 ±2.26	<u>47.55</u> ±2.73	46.84 ±2.99	<u>39.21</u> ±7.92	43.63 ±5.30
ROSCOE-GSM8K							
Coherency	–	59.73 ±1.84	47.50 ±2.70	63.89 ±1.83	62.60 ±2.01	56.53 ±5.42	63.76 ±2.30
Contradiction	–	80.67 ±0.94	55.30 ±2.37	80.53 ±1.42	82.13 ±1.60	83.32 ±1.61	<u>83.26</u> ±2.32
Missing Steps	–	<u>81.13</u> ±1.46	42.77 ±0.77	84.09 ±1.26	<u>81.13</u> ±1.46	<u>52.66</u> ±10.38	84.09 ±1.26
Overall Quality	–	71.60 ±2.15	49.59 ±2.37	65.46 ±2.03	71.33 ±4.00	<u>57.02</u> ±5.07	<u>70.89</u> ±2.52
SummEval							
Coherence	–	31.47 ±1.71	<u>28.79</u> ±2.06	40.90 ±1.48	37.13 ±4.46	28.52 ±3.91	35.65 ±4.42
Consistency	–	83.64 ±1.53	68.84 ±1.80	68.82 ±2.53	82.54 ±1.72	81.17 ±3.28	<u>71.24</u> ±5.99
Fluency	–	4.50 ±0.83	3.33 ±0.58	2.58 ±0.33	77.17 ±6.18	77.53 ±3.84	<u>77.48</u> ±6.28
Relevance	–	37.27 ±1.53	31.27 ±1.72	41.34 ±1.40	49.13 ±8.17	44.33 ±8.00	<u>41.37</u> ±11.08
WMT-20-EnDe	30.02	19.48 ±1.45	17.58 ±1.58	17.15 ±1.45	25.11 ±3.03	<u>23.29</u> ±2.25	<u>20.40</u> ±3.10

Table 5: Repetition of the experiments in Table 1 with only 20 training examples (or 25% of the dataset if it is smaller than 320 examples) instead of 100 training examples.


```
Read the conversation, which consists of:
  (1) a user query (usually to a medic), and
  (2) a chatbot response.

For the chatbot response, don't worry about whether or not the response is good ↔
  medical advice.
We are only interested in what TYPE of response it is.
Select one of the following types:

- Non-medical query
- Irrelevant or nonsensical
- No answer
- General information
- Recommendations
- Treatment or diagnosis

{{ examples }}

# Example

## Query
{{ query }}

## Response
{{ response }}

Response label:
```

```
# Example {{ icl_example_i }}:

## Query
{{ query }}

## Response
{{ response }}

Response label: {{ icl_example_label }}
```

Figure 3: Example (top) prompt for one of the judgment tasks (medical-safety: response type), where {{ examples }} is included for in-context learning experiments only. The {{ examples }} placeholder is replaced with one example for each output label following the example format (bottom) of the corresponding task.