

Judging the Judges: Evaluating Alignment and Vulnerabilities in LLMs-as-Judges

Aman Singh Thakur* and Kartik Choudhary* and Venkat Srinik Ramayapally*

University of Massachusetts Amherst

{amansinghtha, kartikchoudh, vramayapally}@umass.edu

Sankaran Vaidyanathan

University of Massachusetts Amherst

sankaranv@cs.umass.edu

Dieuwke Hupkes

Meta

dieuwkehupkes@meta.com

Abstract

The *LLM-as-a-judge* paradigm offers a potential solution to scalability issues in human evaluation of large language models (LLMs), but there are still many open questions about its strengths, weaknesses, and potential biases. This study investigates thirteen models, ranging in size and family, as ‘*judge models*’ evaluating answers from nine base and instruction-tuned ‘*exam-taker models*’. We find that only the best (and largest) models show reasonable alignment with humans, though they still differ with up to 5 points from human-assigned scores. Our research highlights the need for alignment metrics beyond percent agreement, as judges with high agreement can still assign vastly different scores. We also find that smaller models and the lexical metric contains can provide a reasonable signal in ranking the exam-taker models. Further error analysis reveals vulnerabilities in judge models, such as sensitivity to prompt complexity and a bias toward leniency. Our findings show that even the best judge models differ from humans in this fairly sterile setup, indicating that caution is warranted when applying judge models in more complex scenarios.

1 Introduction

Over the last few years, large language models (LLMs) have demonstrated remarkable capabilities across various domains (Radford et al., 2019; Brown et al., 2020; Achiam et al., 2023; AI@Meta, 2024, i.a.). As more and more new LLMs with different architectures and training methods continue to be released and their capabilities expand, accurately evaluating their performance and limitations becomes increasingly challenging (Zheng et al., 2024; Ohmer et al., 2024; Benchechroun et al., 2023; Madaan et al., 2024; Li et al., 2023a).

LLM evaluation methods generally fall into one of two broad categories. Benchmarks such as

MMLU (Hendrycks et al., 2021), TruthfulQA (Lin et al., 2021), and GSM8K (Cobbe et al., 2021) assess specific capabilities, while leaderboards such as Chatbot Arena (Chiang et al., 2024) and Open LLM Leaderboard (Beeching et al., 2023) rank models based on human or automated pairwise comparisons. Both approaches face challenges in evaluating free-form text responses, as assessment can be as difficult as generation itself (see e.g. Chang et al., 2023; Bavaresco et al., 2024).

One approach to evaluating LLMs is using MCQ benchmarks like MMLU, which compare answer log-probabilities instead of assessing generated responses directly. However, this approach limits the range of measurable abilities and differs from how LLMs are used in practice. Lexical methods, such as exact match (EM) or n-gram overlap, are practical and cost-effective but prone to false negatives and often miss subtle semantic differences. These challenges are amplified for instruction-tuned chat models, which tend to produce more verbose responses (Saito et al., 2023; Renze and Guven, 2024).

For these reasons, human evaluation remains the gold standard for evaluating LLM responses.

Human evaluation is, however, expensive and often impractical, leading to the growing use of LLMs as judge models (Lin et al., 2021; Islam et al., 2023; Chiang and Lee, 2023; Liusie et al., 2024). While promising alignment with humans has been noted (Sottana et al., 2023; Zheng et al., 2024), questions about this approach remain. This work examines LLMs as judges, contrasting them with humans and automated methods. Unlike prior studies, we focus on scenarios with high human alignment to separate task ambiguity from judge model limitations. Using TriviaQA (Joshi et al., 2017), we evaluate how *judge models* of varying architectures and sizes assess *exam-taker models*.

In this work, we study the properties of LLMs as judges, comparing them with humans and auto-

*Equal Contribution

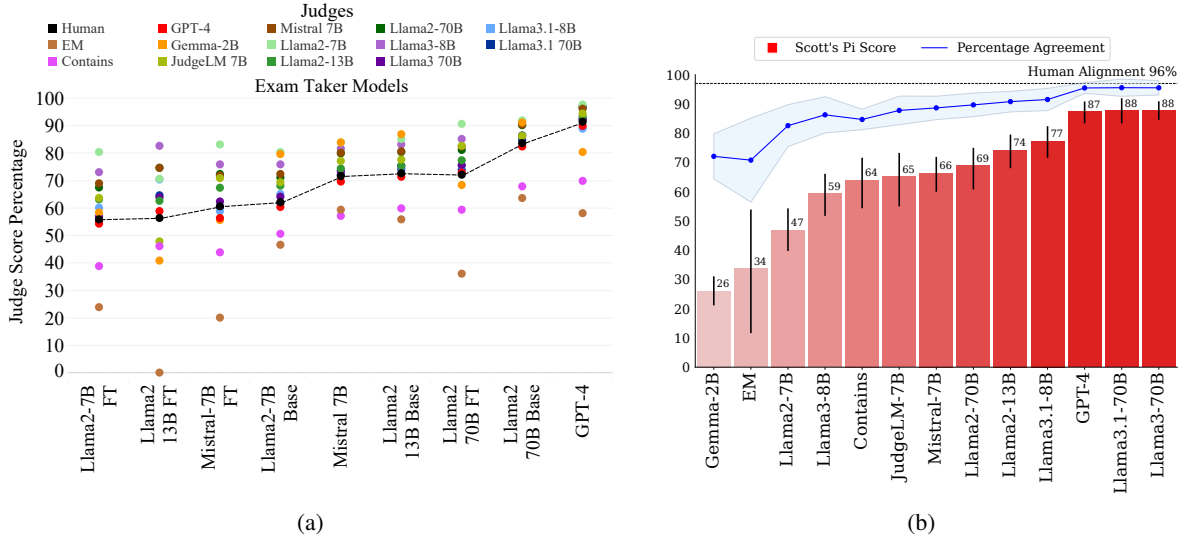


Figure 1: **Average scores assigned by judge models and alignment with human judges.** (a) Scores assigned to all exam-taker models by the various judge models. (b) Average percent agreement (blue line) and Scott’s π scores (red bars) of judge models with human judges (black line). Error bars annotate standard deviation across exam-taker models. Llama3 70B, Llama3.1 70B and GPT-4 Turbo have Scott’s π coefficient that are indicative of excellent alignment, but are still well below the human alignment score.

Exam-taker models (base & instruction-tuned)	Llama-2 (7B, 13B, 70B), Mistral 7B, GPT-4 Turbo
Judge models (instruction-tuned)	Llama-2 (7B, 13B, 70B), Llama-3 (8B, 70B), Llama-3.1 (8B, 70B), Gemma 2B, Mistral 7B, JudgeLM 7B, GPT-4 Turbo
Judge models (lexical)	Exact Match (EM), Contains

Table 1: **Exam-taker models and judge models** We consider a wide variety of exam-taker models and judge models; to get an in-depth overview of their abilities, we consider exam-taker models of various sizes & types.

mated evaluation methods. Contrary to prior work, we focus on a clean scenario in which human alignment is very high, allowing us to distinguish ambiguity and subjectivity in the task itself from potential issues with the judge models. Using the knowledge benchmark TriviaQA (Joshi et al., 2017) as our playground, we investigate how thirteen different *judge models* with varying architectures and sizes judge nine different *exam-taker models*. Our main findings are:

- **Even in clean setups, only the best models have high alignment scores.** Among the thirteen judge models, only GPT-4 Turbo, Llama-3.1; 70B, and Llama-3; 70B achieved strong alignment with humans. However, even these fall short of the human alignment coefficient (Figure 1).
- **Scott’s π distinguishes judges better than percent alignment.** In terms of percent alignment,

judges are rarely discriminable, while Scott’s π provides a more informative signal. In some cases, high percent agreement can still give scores that differ 10-20 points from the human-assigned scores (Figure 2).

- **Also Scott’s π is not all telling** While GPT-4 Turbo and Llama-3 achieve excellent alignment scores, they can differ by up to 5 points from human scores. Moreover, in discriminating between exam-taker models, their performance is comparable to cheaper alternatives like Mistral 7B and contains, which have lower alignment scores but more consistent biases (Figure 3).

Through detailed analysis (§ 5), we gain insights into judge performance. Improved alignment appear to be driven from higher recall rates and fewer false negatives. However, judge models struggle with under-specified answers and exhibit leniency, reducing evaluation consistency. They are also sen-

sitive to prompt length and quality. Surprisingly, even when asked to evaluate a verbatim match with a reference, judge models sometimes fail.

Overall, our work highlights the strengths of the LLM-as-a-judge paradigm, while cautioning against overreliance on alignment metrics, even when they are high. Through error analysis, we identify common failure cases, contributing to a deeper understanding of this emerging evaluation paradigm. With this work, our objective is to improve understanding of the emerging mainstream paradigm for evaluating LLM.

2 Related work

Various recent studies have used or considered using LLMs as judges for tasks such as evaluating story generation (Chiang and Lee, 2023), retrieval-augmented generation (Es et al., 2023), visual QA (Mañas et al., 2024), code comprehension (Zhiqiang et al., 2023), multilingual evaluation (Hada et al., 2023) and more general open-ended tasks (Zheng et al., 2024). Zhang et al. (2024) and Sottana et al. (2023) propose ways to standardise LLM evaluations and the role that judge models might play in such solutions. Several studies have demonstrated that state-of-the-art LLMs such as GPT-4 Turbo exhibit high alignment with human judgments (Sottana et al., 2023; Zheng et al., 2024), though others also illustrate that the paradigm is not yet without faults. Zeng et al. (2023) propose a benchmark for evaluating the performance of LLMs as judges, and other approaches have been proposed to improve LLM judges such that they are aligned well with humans (Shankar et al., 2024; Zhu et al., 2023).

Despite promising results in various settings, judge models still suffer from known issues of current LLMs such as hallucinations and factual errors (Ye et al., 2023; Turpin et al., 2023) and difficulty in following complex instructions (Li et al., 2023b; He et al., 2024). Furthermore, various studies have reported challenges such as position bias (Pezeshkpour and Hruschka, 2023; Zheng et al., 2023; Wang et al., 2023), verbosity bias (Saito et al., 2023) in their preferences, confusing evaluation criteria (Hu et al., 2024), or focusing more on the style and grammar compared to factuality (Wu and Aji, 2023). Recently, Liusie et al. (2024) have shown that LLMs perform better in comparative assessment compared to absolute scoring, which can be used for reliably measuring the relative per-

formance of models (Liu et al., 2024) and creating classifiers for pairwise grading (Huang et al., 2024).

We build on previous work to investigate the strengths and weaknesses of LLMs as judges. Unlike previous studies, we focus on comparing LLM outputs with reference answers rather than pairwise comparisons on open-ended tasks. With high human alignment in this setting, we gain a clearer view of LLM performance. Furthermore, we extend previous research by considering more LLMs, both as judges and as evaluated models.

3 Methodology

To evaluate the strengths and weaknesses of the LLM-as-a-judge paradigm, we focus on a comparatively controlled setup, in which judge models assess answers of exam-taker models on the knowledge benchmark TriviaQA (Joshi et al., 2017). With this methodological design, it is possible to focus on the abilities of the judges in isolation, without having to address human disagreement and error at the same time. In this section, we elaborate the main aspects of our methodology.

Evaluation data As our testbed, we use the TriviaQA dataset (Joshi et al., 2017), consisting of 95K question-answer pairs sourced from 14 trivia and quiz league websites. Each question in the train and validation set is annotated with a list of short answers containing a minimal set of facts and evidence documents collected from Wikipedia and the Web. For our experiments, we use the validation set of the *unfiltered* partition of the benchmark, using the short answers as reference answers. We use the training set for few-shot examples.

Since experiments require manual annotation of the exam-taker model responses, we use a random sample of 400 questions from the dataset. In Appendix I, we show with a bootstrapping test that this sample size has low variance for our main result. Through experiments described in § 3, we establish that humans have high agreement on judgements of answers given to the questions in the benchmark.

Exam-taker models To understand the strengths and weaknesses of different judges, we consider answers of pre-trained (base) and instruction-tuned (chat) ‘exam-taker models’ across a wide variety of model sizes. In particular, we consider Llama-2 (Touvron et al., 2023) in 7B, 13B, and 70B parameter sizes for both base and chat versions, Mistral

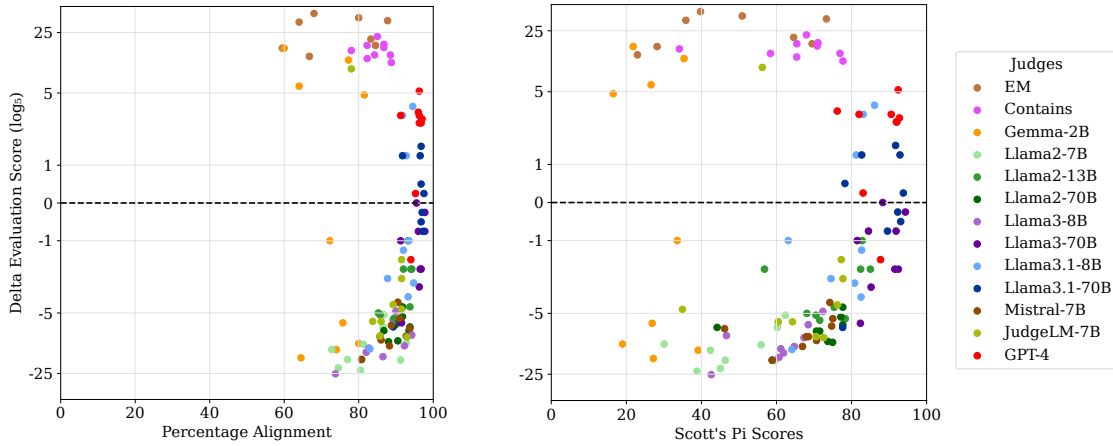


Figure 2: **Difference with human evaluation scores versus alignment metric.** The delta evaluation score is the difference between the judge and the human score; y-axes are in log scale. Percent alignment (left) shows a very skewed distribution, making it difficult to distinguish models. Scott’s π (right) provides a clearer difference between models, and is more indicative of deviation of the gold score.

7B (Jiang et al., 2023) base and chat versions, and GPT-4 Turbo¹ (Achiam et al., 2023) as the exam-taker models. The prompts for the exam-taker models contain five few-shot examples of (question, answer) pairs from the TriviaQA training set. The prompts for the instruction-tuned models additionally include a command signaling the model to answer the given question in a succinct manner similar to the provided examples. The prompts are provided in Appendix D.

Judge models To get a comprehensive view of the strengths and weaknesses of judge models across different model sizes and architectures, we use instruction-tuned versions of Llama-2 (Touvron et al., 2023) in 7B, 13B, and 70B sizes, Llama-3 (AI@Meta, 2024) in 8B and 70B sizes, Llama-3.1 (Dubey et al., 2024) in 8B and 70B sizes, Mistral 7B (Jiang et al., 2023), GPT-4 Turbo (Achiam et al., 2023), Gemma 2B (Gemma Team et al., 2024), and JudgeLM 7B (Zhu et al., 2023) as judges. To maintain parity with human and judge evaluation, judge prompts were built from human guidelines in Appendix G. The judges are instructed to respond with only a single word, “correct” or “incorrect”. An overview of all exam-taker models and judge models is shown in Table 1. For ease of reading, the judge models are depicted in a different font than the exam-taker models.

Baselines As baselines, we use two commonly used lexical evaluation techniques – exact match (EM) and contains match (contains). For EM, a response is considered correct if the response exactly matches one of the reference answers for the given question. For contains, an answer is considered correct if at least one of the reference answers is a sub-string of the response string. Both EM and contains match are computed in a case-insensitive manner.

Alignment We use two metrics to quantify alignment between judges: percent agreement and Scott’s Pi coefficient (Scott, 1955).² Percent agreement expresses a simple percentage of the samples on which two annotators agree. Scott’s Pi, denoted as Scott’s π , is an alignment metric that corrects for chance agreement between two annotators and is considered to provide a more robust measure of alignment. Details about the computation of both metrics are given in Appendix F.

Human judgements As a ground-truth assessment, we obtain human annotations for each exam-taker model answer. The inter-human alignment is calculated between three human judges using the answers to 1200 randomly sampled questions answers; the human guidelines can be found in Appendix G. We then determine collective “Human

¹Accessed via the OpenAI API between Mar 19th, 2024 and Sep 20, 2024.

²In an earlier version of this paper, we used Cohen’s kappa (Cohen, 1960) to measure alignment. It has since come to our attention that – despite its widespread use – this metric has some well-documented theoretical issues (e.g. Pontius and Millones, 2011; Chicco et al., 2021). For the interested reader, we elaborate on these issues in Appendix B.

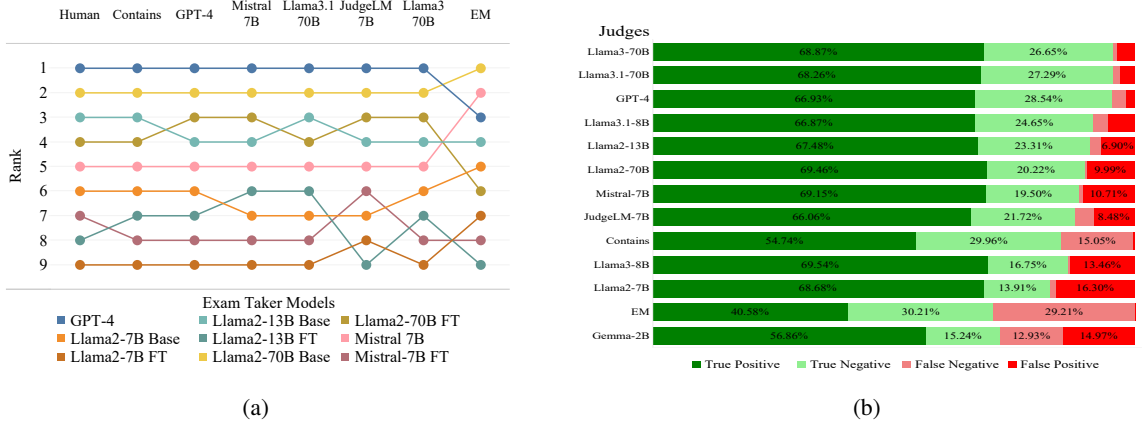


Figure 3: **Judge rankings and true/false positives and negatives.** (a) Assigned exam-taker model rankings assigned by highly human aligned judges. Contains stays closely to human-assigned rankings, as well as GPT-4 Turbo and Mistral 7B. (b) False positives and negatives across different judge models, in descending order of human alignment. Both false negatives and false positives increase as human alignment decreases, but well-aligned models tend to produce more false positives than false negatives.

Judgment” through a majority vote.

The average alignment between human evaluators and the majority vote yielded a Scott’s π of 96.2 ± 1.07 ,³ while the average percentage agreement was $98.52\% \pm 0.42\%$, exceeding the alignment previously reported in comparable studies (Zeng et al., 2024).

The details of this experiment are mentioned in Appendix A. Given this near-perfect alignment score, we consider only one human evaluator per sample for the rest of our experiments, to reduce the overall cost of human annotations. The set of questions for which we obtain human annotations is identical for each exam-taker model.

4 Results

In this section we discuss our main results, primarily focusing on the relationship between evaluations by various judge models and human evaluations (§ 4.1), and how that impacts their usability (§ 4.2). To do so, we evaluate their alignment with human judgment and assess how differently they rank the nine exam-taker models compared to humans. In Section 5, we further analyse their precision and recall to further investigate the types of errors that can be made by various judge models. Details about compute requirements and others costs for experiments are given in Appendix H.

4.1 Alignment between judge models and humans

We start by computing Scott’s π scores and percent agreement between the evaluations of each judge model and the human annotators. We show the result in Figure 1. We observe that percent alignment is high for virtually all models, with the exception of Gemma 2B and EM. Scott’s π , on the other hand, has low values for most models, though its value is in the high 80s for Llama-3 70B, Llama-3.1 70B and GPT-4 Turbo. Nevertheless, there still is a significant disparity between human judgment and judge models: the best scoring judge, Llama-3 70B, is 8 points behind human judgment. Notably, EM has the most variance in alignment, while Gemma 2B has the lowest alignment amongst all judges.

In most cases, we observe that Scott’s π and percent agreement are following the same trend, with the exception of the values for Gemma 2B and EM. Gemma 2B shows higher percent agreement compared to EM, yet it yields the lowest Scott’s π score within the ensemble. For the percent agreement of judge models, we note a 26-point difference between human judgment and EM, while Scott’s π exhibits a more substantial 64-point gap. This is also visible in the general decline of alignment scores: while Llama-3 8B has a Scott’s π score of only 59, its percent agreement is still well above 80%. Overall, Scott’s π appears to be better able of discriminating various judge models, showing more divergence across the tested judges.

³The coefficient is scaled by 100 for easier comparison with percentage alignment.

To understand how indicative the two alignment metrics are of the expected accuracy of the overall judgement of the models, we plot, for each judge model and exam-taker model, the difference between the score assigned by the judge and the score assigned by a human. In the figure, we can see that for Scott’s π values higher than 80, the evaluation scores are comparatively close to the human evaluation scores, with a difference of up to 5 points in their assigned scores (complete results table provided in [Appendix J](#)). For percent alignment, on the other hand, even judges that have more than 90% may still differ more than 10 points in their assigned score. Interestingly, the deviation from human-judgements for a single judge model can be quite different depending on the exam-taker model. In [Figure 1a](#), Gemma 2B, for instance, sometimes assigns higher scores than humans, and sometimes much lower. In the next section, we further explore this particular pattern.

4.2 Exploring consistent patterns in judge models

In the previous section, we saw that none of the judge models were as aligned with humans as humans were with each other. As shown in [Figure 2](#), even the best-aligned judge models can differ by up to 5 points from human-assigned scores. While this limits their ability to perfectly estimate exam-taker model capabilities, judge models can still provide valuable insights to *differentiate* between exam-taker models. For example, judges with consistent biases may not assign identical scores but could rank models similarly, akin to a very strict teacher.

To assess this, we compare the rankings given by each judge model to the nine exam-taker models, computing Spearman’s rank correlation coefficients ρ ([Spearman, 1904](#)) with the human ranking. The rankings are shown in [Figure 3a](#), with ρ and σ values in [Appendix L](#). Most judge models have rank correlations above 0.7, indicating they struggle to distinguish poorer models but do well with better ones. Notably, models like contains and Mistral 7B, which have divergent scores from humans, show high rank correlation (ρ of 0.99 and 0.98, respectively), performing similarly to GPT-4 Turbo and outperforming the better Llama models – though with lower significance values – indicating that identifying which models are better should not be equated to assigning them the correct score.

5 Analysis

To better understand the judge models, we conduct multiple case studies aimed at identifying common errors and vulnerabilities in the judges we investigate. Specifically, we study their precision and recall and error types (§ 5.1), their sensitivity to the instruction prompt prompt (§ 5.2), how they respond to controlled responses of specific types (§ 5.3), and the extent to which they have a *leniency bias* (§ 5.4).

5.1 Better aligned models: Precision and recall gains with error spotlights

We first examine the precision and recall of the judge models. As shown in [Figure 4a](#), both metrics increase moderately with alignment. [Figure 3b](#) reveals a similar trend, with a clearer distribution of false positives and negatives. True positives remain consistent across varying judge quality, whereas true negatives exhibit a slight decline as judge quality decreases. Notably, a reduction in judge quality leads to an increase in false positives.

Next, we analyze the errors made by judge models by manually annotating 900 outputs from Llama-7B Base, focusing on top performers GPT-4 Turbo and Llama-3;70B. We categorize error types and determine how often they are correctly judged as incorrect. The results in [Table 2](#) show that both GPT-4 Turbo and Llama-3;70B excel at identifying answers referring to incorrect entities or containing too many entities. Under-specified and incorrect answers are more challenging, with GPT-4 Turbo performing better on answers with fewer entities than Llama-3;70B.

5.2 Judge model sensitivity to prompt length and specificity

Next, we investigate how prompt length and specificity affect judge models’ inferences to determine whether their performance is influenced by *specificity* of the prompt. We use four prompt versions with varying length and specificity.

The first two prompts (Without;guidelines;V1/V2, 45 and 58 tokens) ask for an evaluation without further details. The longer prompts (Guidelines;without;examples and Guidelines;with;examples, 245 and 301 tokens) provide more elaborate guidance and examples. All prompts are listed in [Appendix M](#).

[Figure 4b](#) shows that GPT-4 Turbo, Llama-3;70B, and Llama-3.1;70B exhibit

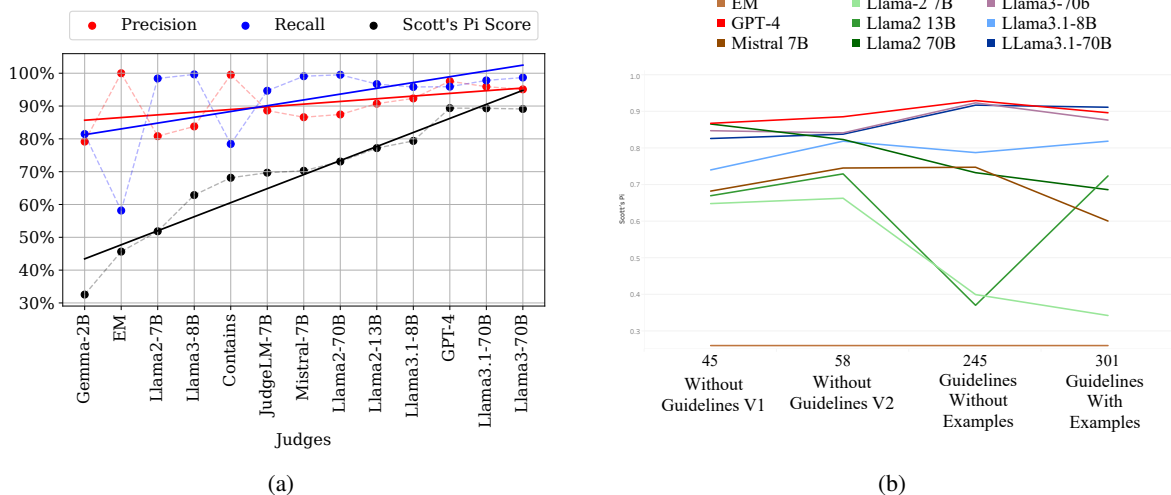


Figure 4: **Precision, recall and prompt sensitivity.** (a) Recall and precision improve with increasing human alignment ($R^2 = 0.31$ and $R^2 = 0.21$, respectively). (b) Scott’s π scores for judges across different instructions.

Error code	Explanation	Example	Proportion	GPT-4 recall	Llama-3 70B recall
Incorrect entity	Response refers to a wrong entity	Henry VII, James I, Edward VI, Mary I and Elizabeth I	86.9%	98.3%	96.6%
Under-specified	Response contains only part of the answer	Henry VII, Henry VIII, Edward, Mary, and Elizabeth	37.3%	33.9%	23.3%
Too few entities	Response contains too few entities	Henry VII, Edward VI, Mary I and James I	2.47%	80.0%	60.0%
Too many entities	Response contains too many entities	Henry VII, Henry VIII, Edward VI, Mary I, James I, and Elizabeth I	2.7%	90.1%	90.1%
Other	Response is incorrect but cannot be put into any of the above buckets	I’m sorry but I do not know the answer to that question	1.23%	20.0%	40.0%

Table 2: **Error analysis for GPT-4 and Llama-3 70B judges.** The example question is “*Excluding Lady Jane Grey, who were the five monarchs of the House of Tudor?*”, the correct answer “*Henry VII, Henry VIII, Edward VI, Mary I and Elizabeth I*” (in any order).

low variance in human agreement as prompt length and specificity increases. Top performers show high alignment with humans even with minimal instructions, while they slightly improve with more detailed prompts. In contrast, other models lose alignment with increased instructions, likely due to difficulty processing complex instructions.

In a follow-up experiment, we investigate the impact of reference order (see Appendix N). Figure 14 and Figure 15 shows that larger models maintain consistent judgments regardless of reference order, while smaller models, except Mistral ; 7B, are more sensitive to it.

5.3 Evaluating controlled responses

We conduct simple tests on the judge models by having them evaluate dummy benchmark responses. In the first test, the answer is a verbatim reference from the dataset (always correct). In the next three tests, the answers are incorrect. For the second and third tests, the dummy exam-taker model responds with “Yes”, and “Sure” respectively. In the fourth

test, the evaluated answer is a repetition of the question.

In Figure 5, we observe that while some judge models correctly identify and mark answers as correct (first test) or incorrect (next three tests), others, like Llama-2 ; 70B, incorrect evaluate many dummy answers, despite showing high human alignment on benchmark evaluations (see Figure 1b). We hypothesize that when the answers are plausible but incorrect, judges can correctly identify them as wrong by comparing them with the reference. However, when the answer is unrelated (e.g., “Yes”, and “Sure”), judge models may mistakenly mark them as correct, though further research is needed to clarify this behavior.

5.4 Leniency bias in judge models

Lastly, to get a general sense of the inherent biases or misalignment in the evaluation criteria that might be present in the judge models, we estimate if they have a positive or negative bias in their judgment. To do so, we assume that a judge as-

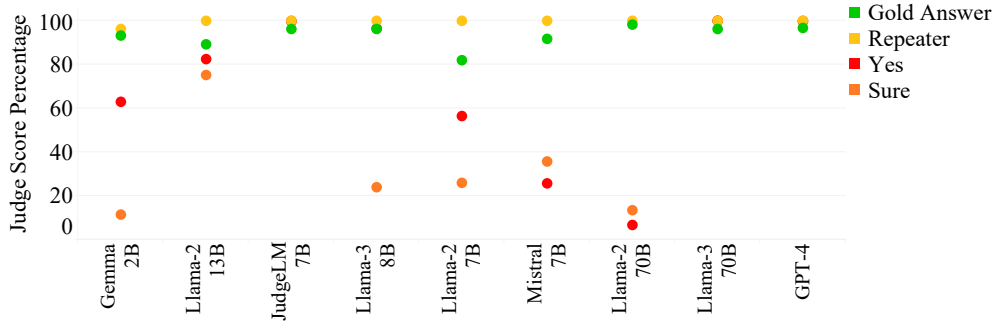


Figure 5: **Judge responses to dummy answers.** We investigate how judge models respond to dummy answers. judge models remain robust when exam-taker models produce responses identical to the prompt (‘repeater’), but are less robust when the responses are “Yes” and “Sure”. Even when the answer matches one of the reference answers verbatim (‘Gold answer’), judges do not always arrive at the correct judgement.

signs the correct judgment (i.e. same evaluation as the ground truth) with a probability of P_c and assigns the rest of the samples to be “correct” with a probability P_+ , which we call their *leniency bias*. We estimate the values of P_c and P_+ from the benchmark results⁴ and show them in Figure 16a. We observe that P_+ for most models is significantly higher than 0.5 (Figure 16b), indicating a tendency of the judge models to evaluate responses as “correct” when their evaluation criteria are not completely aligned with the provided instructions.

6 Conclusion

In this work, we conduct an extensive study of LLMs as judges, comparing them to human judges and automated evaluation methods. By focusing on a clean evaluation scenario with high inter-human agreement, we identify potential issues with the LLM-as-a-judge paradigm, separate from task ambiguity.

We find that smaller, cost-efficient models, like Mistral;7B, are less effective than larger models such as GPT-4 Turbo, Llama-3.1;70B, and Llama-3;70B, which are better aligned but still fall short of human alignment. Even with high alignment, their scores can differ by up to 5 points from human scores, highlighting the need for caution when using judges in more complex scenarios. We also note that the commonly used metric of *percent aligned* fails to differentiate between judges effectively. We suggest future work adopt the more robust Scott’s π metric for better distinction.

Next, we note that high alignment scores are not

⁴The theoretical derivation of the expressions for P_c and P_+ , as well as the empirical validation for their estimated values can be found in Appendix O.

always necessary to *discriminate* between models. While GPT-4 Turbo and Llama-3 have excellent alignment scores, simpler and more cost-efficient models, like contains, perform similarly in ranking exam-taker models, despite lower alignment scores and score deviations. For studies focused on ranking models rather than estimating exact scores, these approaches can be as suitable as more expensive ones.

Lastly, we run experiments to assess judge models’ sensitivity to prompts, precision, recall, error types, leniency, and vulnerability to dummy answers. We find that smaller models are more likely to judge positively when in doubt, that lower-alignment models lack precision, and that better models are more robust across different prompts but harder to “steer.” Some judge models are easily fooled by dummy answers like “Yes” and “Sure” and are better at detecting completely incorrect answers than partially incorrect ones.

Overall, this work contributes to LLM evaluation by assessing judges in a clearly defined framework. Our results highlight the potential of LLMs as judges but caution against blindly trusting their judgments, even when aligned with humans. We recommend computing both percent agreement and Scott’s π , paired with qualitative analysis, to avoid bias. We discuss limitations in Appendix A and plan to expand our work to more complex scenarios in the future.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. [GPT-4 technical re-](#)

- port. [arXiv preprint arXiv:2303.08774](#).
- AI@Meta. 2024. [Llama 3 model card](#).
- 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 Suglia, 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.
- Edward Beeching, Clémentine Fourier, Nathan Habib, Sheon Han, Nathan Lambert, Nazneen Rajani, Omar Sanseviero, Lewis Tunstall, and Thomas Wolf. 2023. Open llm leaderboard. https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard.
- Youssef Benchechrout, Megi Dervishi, Mark Ibrahim, Jean-Baptiste Gaya, Xavier Martinet, Grégoire Mialon, Thomas Scialom, Emmanuel Dupoux, Dieuwke Hupkes, and Pascal Vincent. 2023. [Worldsense: A synthetic benchmark for grounded reasoning in large language models](#). [arXiv preprint arXiv:2311.15930](#).
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). [Preprint](#), arXiv:2005.14165.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2023. [A survey on evaluation of large language models](#). *ACM Transactions on Intelligent Systems and Technology*.
- Cheng-Han Chiang and Hung-yi Lee. 2023. [Can large language models be an alternative to human evaluations?](#) [arXiv preprint arXiv:2305.01937](#).
- Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. 2024. [Chatbot arena: An open platform for evaluating LLMs by human preference](#). [Preprint](#), arXiv:2403.04132.
- Davide Chicco, Matthijs J. Warrens, and Giuseppe Jurman. 2021. [The matthews correlation coefficient \(mcc\) is more informative than cohen’s kappa and brier score in binary classification assessment](#). [ieee access](#), 9:78368–78381.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. [Training verifiers to solve math word problems](#). [arXiv preprint arXiv:2110.14168](#).
- J. Cohen. 1960. [A Coefficient of Agreement for Nominal Scales](#). *Educational and Psychological Measurement*, 20(1):37.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Al-lonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Young, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shao-liang Nie, Sharan Narang, Sharath Rapparthi,

Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Barambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Ibrahim Damla, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khan-delwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal Chawla, Kushal Lakhota, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng

Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabza, Manav Avalani, Manish Bhatt, Maria Tsim-poukelli, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Her-moso, Mo Metanat, Mohammad Rastegari, Mun-ish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pa-van Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratan-chandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Mah-eswari, Russ Howes, Ruty Rinott, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lind-say, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agar-wal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vitor Albiero, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaofang Wang, Xiao-jian Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. 2024. [The llama 3 herd of models](#). [Preprint](#), arXiv:2407.21783.

Shahul Es, Jithin James, Luis Espinosa-Anke, and Steven Schockaert. 2023. [RAGAS: Automated eval-uation of retrieval augmented generation](#). [Preprint](#), arXiv:2309.15217.

Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christo-pher A. Choquette-Choo, Clément Crepy, Daniel Cer,

- Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikula, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimenko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira, Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. 2024. [Gemma: Open models based on gemini research and technology](#). Preprint, arXiv:2403.08295.
- Rishav Hada, Varun Gumma, Adrian de Wynter, Harshita Diddee, Mohamed Ahmed, Monojit Choudhury, Kalika Bali, and Sunayana Sitaram. 2023. [Are large language model-based evaluators the solution to scaling up multilingual evaluation?](#) arXiv preprint arXiv:2309.07462.
- Qianyu He, Jie Zeng, Wenhao Huang, Lina Chen, Jin Xiao, Qianxi He, Xunzhe Zhou, Jiaqing Liang, and Yanghua Xiao. 2024. [Can large language models understand real-world complex instructions?](#) Proceedings of the AAAI Conference on Artificial Intelligence, 38(16):18188–18196.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). Preprint, arXiv:2009.03300.
- Xinyu Hu, Mingqi Gao, Sen Hu, Yang Zhang, Yicheng Chen, Teng Xu, and Xiaojun Wan. 2024. [Are LLM-based evaluators confusing nlg quality criteria?](#) arXiv preprint arXiv:2402.12055.
- Hui Huang, Yingqi Qu, Jing Liu, Muyun Yang, and Tiejun Zhao. 2024. [An empirical study of LLM-as-a-Judge for LLM evaluation: Fine-tuned judge models are task-specific classifiers](#). Preprint, arXiv:2403.02839.
- Pranab Islam, Anand Kannappan, Douwe Kiela, Rebecca Qian, Nino Scherrer, and Bertie Vidgen. 2023. [FinanceBench: A new benchmark for financial question answering](#). arXiv preprint arXiv:2311.11944.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7B](#). arXiv preprint arXiv:2310.06825.
- Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). arXiv preprint arXiv:1705.03551.
- Junlong Li, Shichao Sun, Weizhe Yuan, Run-Ze Fan, Hai Zhao, and Pengfei Liu. 2023a. [Generative judge for evaluating alignment](#). arXiv preprint arXiv:2310.05470.
- Shiyang Li, Jun Yan, Hai Wang, Zheng Tang, Xiang Ren, Vijay Srinivasan, and Hongxia Jin. 2023b. [Instruction-following evaluation through verbalizer manipulation](#). arXiv preprint arXiv:2307.10558.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2021. [TruthfulQA: Measuring how models mimic human falsehoods](#). arXiv preprint arXiv:2109.07958.
- Yinhong Liu, Han Zhou, Zhijiang Guo, Ehsan Shareghi, Ivan Vulic, Anna Korhonen, and Nigel Collier. 2024. [Aligning with human judgement: The role of pairwise preference in large language model evaluators](#). arXiv preprint arXiv:2403.16950.
- Adian Liusie, Potsawee Manakul, and Mark Gales. 2024. [LLM comparative assessment: Zero-shot NLG evaluation through pairwise comparisons using large language models](#). In Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers), pages 139–151, St. Julian’s, Malta. Association for Computational Linguistics.
- Lovish Madaan, Aaditya K. Singh, Rylan Schaeffer, Andrew Poulton, Sanmi Koyejo, Pontus Stenetorp, Sharan Narang, and Dieuwke Hupkes. 2024. [Quantifying variance in evaluation benchmarks](#). arXiv preprint arXiv:2406.10229.
- Oscar Mañas, Benno Krojer, and Aishwarya Agrawal. 2024. [Improving automatic vqa evaluation using large language models](#). Proceedings of the AAAI Conference on Artificial Intelligence, 38(5):4171–4179.
- Xenia Ohmer, Elia Bruni, and Dieuwke Hupkes. 2024. [From form \(s\) to meaning: Probing the semantic depths of language models using multisense consistency](#). arXiv preprint arXiv:2404.12145.
- Pouya Pezeshkpour and Estevam Hruschka. 2023. [Large language models sensitivity to the order of options in multiple-choice questions](#). arXiv preprint arXiv:2308.11483.
- Robert Gilmore Pontius and Marco Millones. 2011. Death to kappa: birth of quantity disagreement and allocation disagreement for accuracy assessment. *Int. J. Remote Sens.*, 32(15):4407–4429.

- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Matthew Renze and Erhan Guven. 2024. [The benefits of a concise chain of thought on problem-solving in large language models](#). *arXiv preprint arXiv:2401.05618*.
- Keita Saito, Akifumi Wachi, Koki Wataoka, and Youhei Akimoto. 2023. [Verbosity bias in preference labeling by large language models](#). *Preprint*, arXiv:2310.10076.
- W.A. Scott. 1955. Reliability of content analysis: The case of nominal scale coding. *The Public Opinion Quarterly*, 17:133–139.
- Shreya Shankar, JD Zamfirescu-Pereira, Björn Hartmann, Aditya G Parameswaran, and Ian Arawjo. 2024. [Who validates the validators? aligning llm-assisted evaluation of llm outputs with human preferences](#). *arXiv preprint arXiv:2404.12272*.
- Andrea Sottana, Bin Liang, Kai Zou, and Zheng Yuan. 2023. [Evaluation metrics in the era of gpt-4: reliably evaluating large language models on sequence to sequence tasks](#). *arXiv preprint arXiv:2310.13800*.
- C. Spearman. 1904. [The proof and measurement of association between two things](#). *The American Journal of Psychology*, 15(1):72–101.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel Bowman. 2023. [Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 74952–74965. Curran Associates, Inc.
- Peiyi Wang, Lei Li, Liang Chen, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023. [Large language models are not fair evaluators](#). *arXiv preprint arXiv:2305.17926*.
- Minghao Wu and Alham Fikri Aji. 2023. [Style over substance: Evaluation biases for large language models](#). *Preprint*, arXiv:2307.03025.
- Hongbin Ye, Tong Liu, Aijia Zhang, Wei Hua, and Weiqiang Jia. 2023. [Cognitive mirage: A review of hallucinations in large language models](#). *arXiv preprint arXiv:2309.06794*.
- Zhiyuan Zeng, Jiatong Yu, Tianyu Gao, Yu Meng, Tanya Goyal, and Danqi Chen. 2023. [Evaluating large language models at evaluating instruction following](#). *arXiv preprint arXiv:2310.07641*.
- Zhiyuan Zeng, Jiatong Yu, Tianyu Gao, Yu Meng, Tanya Goyal, and Danqi Chen. 2024. [Evaluating large language models at evaluating instruction following](#). *Preprint*, arXiv:2310.07641.
- Yue Zhang, Ming Zhang, Haipeng Yuan, Shichun Liu, Yongyao Shi, Tao Gui, Qi Zhang, and Xuanjing Huang. 2024. [Llmeval: A preliminary study on how to evaluate large language models](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(17):19615–19622.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2023. [On large language models’ selection bias in multi-choice questions](#). *arXiv preprint arXiv:2309.03882*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. [Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena](#). *Advances in Neural Information Processing Systems*, 36.
- Yuan Zhiqiang, Liu Junwei, Zi Qiancheng, Liu Mingwei, Peng Xin, Lou Yiling, et al. 2023. [Evaluating instruction-tuned large language models on code comprehension and generation](#). *arXiv e-prints arXiv:2308.01240*.
- Lianghui Zhu, Xinggang Wang, and Xinlong Wang. 2023. [Judgelm: Fine-tuned large language models are scalable judges](#). *Preprint*, arXiv:2310.17631.

A Limitations

In our work, we have evaluated how 11 different LLMs fare as judges in a scenario in which judgments should be relatively straight-forward, and human alignment is high. As any study, our work has several limitations as well as directions that we did not explore but would have been interesting too. In this section, we discuss both.

Simplicity of the task As mentioned in the introduction of our work, the scenario in which judges are used are typically much more complicated than the scenario that we focussed on. Specifically, judges are most often deployed in preference rankings (where two model responses are compared) or to judge complex answers that are difficult to automatically parse. In such tasks, human agreement is often low, making it challenging to judge the judges themselves. In our work, we have deliberately chosen for a simple task, in which human alignment is high. The main premise is, that if a judge does not perform well in this simple setup, caution is suggested also in more complex setups – if someone cannot do multiplication, why would they be able to solve ordinary differential equations. Given the poor understanding of which abilities of LLMs generalise in what dimensions, however, more studies are needed to understand how our results generalise to various other scenarios.

Human alignment In an earlier version of this paper, due to the high cost of human annotations, we opted to select a single model for human annotation as we iteratively modified the exam taker prompt, few-shot examples, and guidelines. We selected the Llama2 7B for this purpose with a random sample of 600 questions. As this is only a single model, it is possible that our human alignment scores are biased because of that. After, we have therefore extended our results with another 600 human-annotated examples from Llama3.1 70B.

For Llama2 7B The average alignment among human evaluators had a Scott’s π of 96.36 ± 1.46 , and the average percent agreement was $98.33\% \pm 0.76\%$. For Llama3.1 70B, we noted that the average alignment among human evaluators had Scott’s π of 95.78 ± 0.30 , % and the average percent agreement was $98.72\% \pm 0.10\%$. Given the similarity of these two numbers, we believe that these 1200 samples provide an adequate estimate. In the paper, we take the average.

Size of the judged samples As each of the nine exam-taker models requires human annotations for each sample, we restricted our analysis to 400 samples in total. This sample size also allowed us to conduct manual annotations and error analysis within 75 human hours/200 GPU hours (see [Appendix H](#)) and give reliable confidence intervals while also providing the flexibility to compare a range of models. We were not able to increase the size due to the cost, but a statistical analysis (details provided in [Appendix I](#)) illustrated that the variance because of this sample size was very low.

Selection of judges With our selection of judges, we have stuck to autoregressive judges that can be used off-the-shelf, as well as one LLM specifically trained to judge. They are – at the moment of writing – the ones that are most commonly used as LLM-judges, and we have tried to be comprehensive across size and family. Nevertheless, we acknowledge that there are other judges that we could have considered as well. As including more judges in – compared to including more exam-taker models – relatively straightforward because it requires only computational power, no manual annotation, we hope that others may evaluate their newly proposed judges using our setup as well.

Future work All in all, these differences underline how finicky using LLMs as judges can be, and with that confirm the overall conclusions of our study that much more work is needed to better understand the strengths and limitations of judge models across a wide range of scenarios and model accuracies. We consider assessing the strengths across multiple different samples and tasks, which would require many more human annotations, outside the scope of this paper and leave such experimentation for future work.

B A brief explanation of the theoretical issues with Cohen’s kappa

Cohen’s Kappa Coefficient ([Cohen, 1960](#)) is a statistic to measure inter-rater agreement for categorical responses. Cohen’s Kappa coefficient measures this agreement by computing the observed (percent) agreement between raters (p_o) and comparing it with the hypothetical probability of chance agreement (p_e), which is taken as a baseline, as follows:

$$\kappa \equiv \frac{p_o - p_e}{1 - p_e} \quad (1)$$

In this equation, the chance agreement p_o constitutes the hypothetical probability that observed agreement occurred by chance, given the observed distributions of the considered raters, under the assumption that the probabilities the raters assign to the observed labels are independent. Specifically, it is defined as:

$$\begin{aligned} p_e &= \sum_k \widehat{p}_{k12} =_{ind} \sum_k \widehat{p}_{k1} \widehat{p}_{k2} \\ &= \sum_k \frac{n_{k1}}{N} \cdot \frac{n_{k2}}{N} = \frac{1}{N^2} \sum_k n_{k1} n_{k2} \end{aligned}$$

where \widehat{p}_{k12} is the estimated probability that rater 1 and rater 2 will classify the same item as k , rewritten to $\widehat{p}_{k1} \widehat{p}_{k2}$ under the assumption that p_{k1} and p_{k2} are independent. The crux of the issue with this method of computation, is that \widehat{p}_{k1} and \widehat{p}_{k2} are estimated independently from the data. As such, the chance agreement adjusts for the observed average differences between raters, which is in fact part of what we intend to measure.

To address this issue, Scott’s Pi (Scott, 1955) instead defines the chance baseline under the assumption that the raters have the same distribution, which is estimated considering the joint distribution of rater 1 and rater 2, rather than considering them separately. It defines p_e as:

$$p_e = \sum_k \widehat{p}_k^2 = \sum_k \sum_k \left(\frac{n_{k1} + n_{k2}}{2N} \right)^2 \quad (2)$$

As such, contrary to Cohen’s Kappa, it captures differences surpassing the chance agreement if rater 1 and rater 2 were in fact equivalent. In other words, we compare against a baseline in which raters would be equivalent, and we measure how much they deviate from that.

Note that if the empirical distributions of rater 1 and rater 2 are the same, so will the values of Scott’s Pi and Cohen’s Kappa be. This also implies that for larger observed (percent) alignment values, the values for Cohen’s Kappa and Scott’s Pi will be closer.

C Model and dataset details

In Appendix C, we show the different models and datasets used in our experiments, along with version and license details.

D Model evaluation prompt templates

In Figure 6 and Figure 7, we show the prompt templates used for the base and chat exam-taker models during the question answering process.

E Judge LLM Prompt templates

In Figure 8, we show the prompt template used to guide the judge models during the evaluation process of a 400-question sample from the TriviaQA unfiltered dataset.

F Metrics for judge models

If one of the annotators is taken to be the reference, then the annotations of the other annotator can be categorized as true positives, false positives, true negatives, and false negatives, with the total number of each of them in a benchmark being represented by T_P , F_P , T_N , and F_N respectively.

Percent agreement is simply the ratio of the numbers of times two annotators agree with each other relative to the total number of annotations. This ratio can have values between 0 and 1. For the binary case, the alignment ratio ρ is given as

$$\rho = \frac{T_P + T_N}{T_P + F_P + T_N + F_N}. \quad (3)$$

Scott’s Pi, (Scott, 1955), measures the alignment of two annotators while also taking into account the possibility of agreement by pure chance. This coefficient usually has values above 0 in most real-world situations. The value of Scott’s Pi is given below where p_o is the relative observed agreement, and p_e is the hypothetical probability of chance agreement.

Prompt template for Base exam-taker models

Q: Can you name the actress who links 'The Darling Buds of May' and 'Rosemary and Thyme'?

A: Pam Ferris

Q: A neologism is a new?

A: Word/expression

Q: Who, in 2010, became the first person from outside the British Isles to win the World Snooker Championship title since Cliff Thorburn in 1980, and the first non British player to win the title since Ken Doherty in 1997?

A: Neil Robertson

Q: Which German Nazi leader flew solo from Ausberg in 1941 and landed by parachute near Glasgow on a private peace mission?

A: Hess

Q: Where would you find Narita airport?

A: Tokyo, Japan

Q: Which cartoon title character has a friend called Captain Haddock?

A:

Figure 6: Prompt template for base exam-taker models

Prompt template for Chat exam-taker models

You are a part of a question answering benchmark. Look at the following examples on how to answer the questions.

““

Q: Can you name the actress who links 'The Darling Buds of May' and 'Rosemary and Thyme'?

A: Pam Ferris

Q: A neologism is a new?

A: Word/expression

Q: Who, in 2010, became the first person from outside the British Isles to win the World Snooker Championship title since Cliff Thorburn in 1980, and the first non British player to win the title since Ken Doherty in 1997?

A: Neil Robertson

Q: Which German Nazi leader flew solo from Ausberg in 1941 and landed by parachute near Glasgow on a private peace mission?

A: Hess

Q: Where would you find Narita airport?

A: Tokyo, Japan

““

Your task is to answer the following question. Remember to be concise and only give the answer in a few words.

Q: Which cartoon title character has a friend called Captain Haddock?

A:

Figure 7: Prompt template for Chat exam-taker models

Asset	Version	License
TriviaQA	mandarjoshi/trivia_qa	apache-2.0
Llama-2 7B Base	meta-llama/Llama-2-7b-hf	llama2
Llama-2 7B Chat	meta-llama/Llama-2-7b-chat-hf	llama2
Llama-2 13B Base	meta-llama/Llama-2-13b-hf	llama2
Llama-2 13B Chat	meta-llama/Llama-2-13b-chat-hf	llama2
Llama-2 70B Base	meta-llama/Llama-2-70b-hf	llama2
Llama-2 70B Chat	meta-llama/Llama-2-70b-chat-hf	llama2
Mistral 7B Base	mistralai/Mistral-7B-v0.1	apache-2.0
Mistral 7B Chat	mistralai/Mistral-7B-Instruct-v0.2	apache-2.0
Llama-3 8B Chat	meta-llama/Meta-Llama-3-8B-Instruct	llama3
Llama-3 70B Chat	meta-llama/Meta-Llama-3-70B-Instruct	llama3
Llama-3.1 8B Chat	meta-llama/Meta-Llama-3.1-8B-Instruct	llama3.1
Llama-3.1 70B Chat	meta-llama/Meta-Llama-3.1-70B-Instruct	llama3.1
JudgeLM	BAAI/JudgeLM-7B-v1.0	Non-commercial license
GPT-4 Turbo	gpt-4-turbo-2024-04-09	N/A

Table 3: Version and license details for the different models and datasets used in experiments.

<i>Prompt template for judge models</i>
<p>Your task is to look at the following question, and based on the references provided, determine if the model’s response is correct or incorrect. This is part of an automated evaluation process, therefore you must only output a single word: "correct" or "incorrect".</p> <p>Question: Which Australian did Roger Federer defeat to win his first Wimbledon Men’s Singles title in 2003?</p> <p>References: MARK PHILIPPOUSSIS MARK PHILIPPOUSSIS</p> <p>Model Response: Mark Philippoussis</p> <p>Evaluation (correct/incorrect):</p>

Figure 8: Prompt templates for the judge models

$$\pi = \frac{p_o - p_e}{1 - p_e} = \frac{\left(\frac{TP+TN}{TP+TN+FP+FN} \right) - \left(\frac{(TP+FP)(TP+FN)+(TN+FN)(TN+FP)}{(TP+TN+FP+FN)^2} \right)}{1 - \left(\frac{(TP+FP)(TP+FN)+(TN+FN)(TN+FP)}{(TP+TN+FP+FN)^2} \right)}$$

G Human Annotation Guidelines

Human annotation guidelines

You will be given a question, a set of reference answers and the answer given by an LLM. Your task is to judge if the answer given by the LLM is correct, as if you were the LLMs teacher grading their exam. An answer should be counted as correct if it is semantically equivalent to (one of the) reference answers. In doing so, please follow the following guidelines:

- Underspecified answers (e.g. "December" instead of "December 20") should be marked incorrect.
- Answers that have more information than requested (e.g. "December 20, in Paris" instead of "December 20") should be marked correct, provided the extra information is not incorrect or contrasting the rest of the answer.
- Answers with unnecessary verbosity but correct answers should be marked correct (E.g. "Thanks for asking this question! The correct answer is: ...").

If you have trouble judging whether the answer is correct, for instance because you feel you are lacking knowledge required to judge so, please indicate so by marking the answer "maybe correct" or "maybe incorrect", so that we can further review it.

Preliminary research involved iterative refinement of human annotation guidelines to ensure consistency and reproducibility across annotators with general English semantic knowledge. CS graduate students served as annotators for this experiment. We provide the guidelines used for human evaluation below.

H Experiment costs

The costs for the different experiments described in this work belong in three categories – GPU-hours for running open-source models on one or more Nvidia A100 GPUs, OpenAI credits for making API calls to OpenAI models,⁵ and human hours for manual annotations of benchmark responses. The estimated costs for the final reported experiments are given in [Appendix K](#). In addition to this, previous unreported experiments and trials had an approximate cost of 120 GPU-hours, 100 USD in OpenAI credits, and 50 human hours, bringing the total experimental cost for this work to approximately 200 GPU-hours, USD 125 OpenAI credits, and 75 human annotation hours.

I Statistical reliability of Evaluation sample

Due to computational constraints discussed in [Appendix A](#) and [Appendix H](#), we limit our evaluation set to randomly sampled 400 questions from TriviaQA ([Joshi et al., 2017](#)). In this section, we further take 5 samples of 300 randomly selected questions from the evaluation set and calculate the mean and standard deviation of Scott’s π . From [Appendix I](#), it can be observed that even on down-sampled sets, the Scott’s π values are similar to [Figure 1b](#). Standard deviation of all the judge models from the mean Scott’s π is also minimal, barring EM lexical match.

Judge Model	Mean Scott’s π	Std Dev
Llama3-70B	0.88	0.0046
Llama3.1-70B	0.88	0.0039
Llama3.1-8B	0.78	0.0050
Llama2-13B	0.75	0.0043
Llama2-70B	0.69	0.0114
Mistral-7B	0.67	0.0108
JudgeLM-7B	0.66	0.0026
Contains	0.64	0.0087
Llama3-8B	0.60	0.0126
Llama2-7B	0.47	0.0112
EM	0.47	0.29
Gemma-2B	0.26	0.007

Table 4: Weak Scott’s π variation for the 5 down-sampled sets indicating robustness for the evaluation sample

J Judge Scores

We show the scores assigned by each judge model to each exam-taker model, visualised in [Figure 1a](#) in [Appendix K](#).

K Exam-taker model base vs chat analysis

Given the human judgments we have available, we take the opportunity to investigate the performance differences between base and their corresponding chat models. In [Appendix K](#), we show the scores assigned by various judge models to four base-chat pairs. According to the default metric EM, the base models outperform the chat models by a large margin. Interestingly, while this difference gets smaller when the answers are judged by humans (second column) or GPT-4 Turbo, there is still a substantial difference for all four pairs, suggesting that the difference is not merely an effect of the increased verbosity of the chat models. Further evidence for that hypothesis is provided by [Figure 9b](#), in which we can see that while 14% of the errors are shared between the base-chat pairs, almost another 14% of the examples get judged correctly by the base models but not by the chat models, while the opposite happens in only 2.5% of the cases.

⁵Pricing details for OpenAI models are available at <https://openai.com/api/pricing/>

Experiment	GPU-hours	OpenAI credits	Human hours
Main benchmarks	5	2	-
Main evaluations	30	8	10
Human alignment	2	-	9
Error analysis	1.5	-	5
Controlled responses	15	-	-
Leniency bias	5	5	-
Guideline bias	10	5	1
Reference bias	5	4	1
Total	73.5	24	26

Table 5: Estimated costs for the final reported experiments. GPU-hours are in equivalent Nvidia A100 hours, OpenAI credits are in USD, and human hours are time spent in manual annotation.

Judge Models	Exam taker models								
	Llama2						Mistral		GPT-4
	7B	Base 13B	70B	7B	Chat 13B	70B	Base 7B	Instruct	
Llama 3.1 8B	65.25	75.00	83.50	60.25	70.50	75.50	73.75	59.00	89.00
Llama 3.1 70B	62.00	74.25	85.00	55.50	64.75	74.00	72.25	60.50	92.25
Llama 3 8B	76.00	83.25	91.50	73.25	82.75	85.25	81.75	76.0	97.25
Llama 3 70B	64.25	75.50	86.50	57.00	64.00	75.75	73.5	62.50	92.75
Llama 2 7B	80.50	85.25	92.00	80.50	70.75	90.75	84.00	83.25	97.75
Llama 2 13B	68.25	75.50	86.50	63.25	62.75	77.50	74.50	67.50	93.5
Llama 2 70B	71.25	80.5	90.25	67.50	74.75	81.25	80.0	72.5	96.75
Mistral 7B	72.50	80.75	90.50	69.00	74.75	82.50	80.25	72.00	96.25
Gemma 2B	79.75	87.00	91.25	58.50	41	68.50	84.0	55.75	80.50
JudgeLM	69.50	77.75	86.25	63.75	48.0	82.75	77.25	71.0	94.50
GPT-4	60.50	71.50	82.50	54.50	59.0	73.0	69.75	56.50	90.0
Exact Match	46.75	56.00	63.75	24.00	0.25	36.25	59.50	20.25	58.25
Contains Match	50.75	60.00	68.00	39.00	46.25	59.50	57.25	44.00	70.00
Human Eval	62.50	72.75	83.75	56.00	56.50	72.25	71.75	60.75	91.50

Table 6: Judge model score card for every exam-taker model.

We consider two alternative hypotheses:

- i) The chat models have a worse understanding of the particular prompt format, which is tuned more to fit base models; or
- ii) The chat models have ‘unlearned’ some knowledge during their alignment training.

To disentangle these two factors, we manually analyse 400 questions for Llama-2 70B and Llama-2 70B-chat, using our earlier error codes. The results, shown in Figure 9a, suggest that, at least to some extent, the difference between base and chat models is in fact due to ‘unlearning’ of knowledge: while the number of errors is more or less equal among most categories, there is a stark difference in the incorrect entity category. Substantially more often than the base models, the chat models do answer the question with a semantically plausible but incorrect entity. In Appendix M-Appendix M, we provide examples of such cases. The results do not show any evidence to support the first hypothesis: the number of errors where the answer cannot be parsed or is just entirely incorrect does not differ between base and chat models.

L Exam-taker model ranking correlation

In Appendix L, We use the Spearman Rank correlation coefficient (Spearman, 1904) to assess the rankings of the exam-taker models. To validate these rankings, we randomly select 6 out of 9 exam-taker models across 5 samples, subsequently calculating the mean (ρ) and standard deviation (σ) of the rankings. The results reveal that the contains model exhibits the highest stability and ρ among the rankings, while the majority of judge models achieve a coefficient exceeding 0.7, indicating a strong alignment. Notably, smaller models such

as Mistral 7B perform on par with GPT-4 Turbo, highlighting the robustness of smaller models in maintaining rankings.

Judges	ρ	σ
Contains	0.99	0.02
Mistral-7B	0.98	0.03
GPT-4	0.98	0.03
Llama2-13B	0.95	0.18
JudgeLM-7B	0.95	0.05
Llama2-7B	0.94	0.04
Llama3.1-70B	0.94	0.07
Llama3-70B	0.93	0.05
Llama3.1-8B	0.89	0.10
Llama3-8B	0.86	0.07
Llama2-70B	0.84	0.13
Gemma-2B	0.71	0.20
EM	0.67	0.13

Table 8: Spearman Rank Correlation Coefficient ρ .

M Too much info confuses judges

In Figure 10-13, we report the guidelines we used for the experiments in § 5.2. The simplest prompt used is *Without Guidelines v1* (see Figure 10) where we define a sequential and structured process for the judge model. In *Without Guidelines v2* (see Figure 11), we add an additional focus on the overall task and outcome as well. For *Guidelines without examples* (see Figure 12), we provide the judge models with detailed instructions about the task at hand, along with explicit guidelines on how to evaluate the answers. Additionally, for *Guidelines with examples* (see Figure 13), we also provide examples to the judge models for further reference.

Table 7: Scores of base and chat models by various judges

Base-Chat pair	Judge models									
	EM		Contains		Human		GPT-4 Turbo		Llama-3 70B	
	Base	Chat	Base	Chat	Base	Chat	Base	Chat	Base	Chat
Llama-2 7B	46.75	24.00	50.75	39.00	62.25	56.00	60.50	54.50	64.25	57.00
Mistral 7B	59.50	20.25	57.25	44.00	71.75	60.75	69.75	56.50	73.50	62.50
Llama-2 13B	56.00	0.25	60.00	46.25	72.75	56.50	75.00	59.00	76.50	64.00
Llama-2 70B	63.75	36.25	68.00	59.50	83.75	72.25	82.50	73.00	86.50	75.75

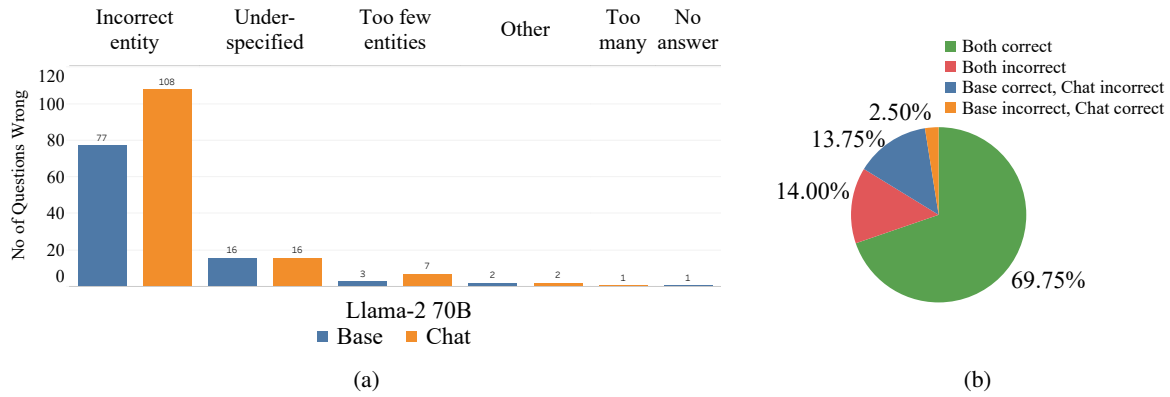


Figure 9: a) Distribution of incorrect question counts by error codes for Llama2 70B Base vs Chat exam-taker models evaluated on 400 questions. b) Pie chart showing the percentage of questions categorized by the judgment from Base and Chat models.

Question:	
Which British artist's works include 'The First Real Target'?	
References	Peter Blake, Peter Balke, Sir Peter Blake
LLama-2 70B Base	Peter Blake
LLama-2 70B Chat	Patrick Caulfield
Mistral 7B Base	David Hockney
Mistral 7B Chat	Damien Hirst

Table 9: Knowledge unlearning example 1.

Question:	
Who was the first cricketer to score 10,000 test runs?	
References	Sunil Gavaskar, Sunil Manohar Gavaskar, SM Gavaskar, Sunny gavaskar, Gavaskar
LLama-2 70B Base	Sunil Gavaskar
LLama-2 70B Chat	Sachin Tendulkar
Mistral 7B Base	Sachin Tendulkar
Mistral 7B Chat	Sachin Tendulkar was the first cricketer to score 10,000 runs in Test matches.

Table 10: Knowledge unlearning example 2

Question: ‘Uncle Harry’s Coat’ was the first garment produced by which famous jacket manufacturer, based in Simonside, Newcastle Upon Tyne?	
References	Barbour
LLama-2 70B Base	Barbour
LLama-2 70B Chat	Barbour
Mistral 7B Base	Barbour
Mistral 7B Chat	Jack Walker & Sons

Table 11: Knowledge unlearning example 3

Without Guidelines v2 (58 tokens)

Your task is to look at the following question, and based on the references provided, determine if the model's response is correct or incorrect. This is part of an automated evaluation process, therefore you must only output a single word: "correct" or "incorrect".

Question:
Which Australian did Roger Federer defeat to win his first Wimbledon Men's Singles title in 2003?

References:
MARK PHILIPPOUSSIS
MARK PHILIPPOUSSIS

Model Response:
Mark Philippoussis

Evaluation (correct/incorrect):

Figure 11: *Without Guidelines v2* prompt template for the judge models

Without Guidelines v1 (44 tokens)

Review the question and examine the references provided, then evaluate the model's response. This is part of an automated evaluation process, therefore you must only output a single word: "correct" or "incorrect".

Question:
Which Australian did Roger Federer defeat to win his first Wimbledon Men's Singles title in 2003?

References:
MARK PHILIPPOUSSIS
MARK PHILIPPOUSSIS

Model Response:
Mark Philippoussis

Evaluation (correct/incorrect):

Figure 10: *Without Guidelines v1* prompt template for the judge models

With Guidelines but no examples (245 tokens)

Review the Question and Examine the Reference: Carefully read and understand the question. Carefully read each of the reference provided for the question:
Evaluate the LLM Response based on guidelines (1-4):
1. Given the model response, only consider the part(s) of the response answering the question, and ignore any extra text.
2. When evaluating the language model's answer, evaluate True if it matches with at least one of the provided references otherwise False.
3. An answer and reference pair is considered a match if they are semantically equivalent given the context of the question. This determination should not require world knowledge from the human, and if it does, consider it not a match. The answer must contain all the information of the reference for it to be considered a match.
4. Considering there might be multiple correct answers, each answer must match with at least one reference for the response to be considered correct overall.
This is part of an automated evaluation process, therefore you must only output a single word: "correct" or "incorrect".

Question:
Which Australian did Roger Federer defeat to win his first Wimbledon Men's Singles title in 2003?

References:
MARK PHILIPPOUSSIS
MARK PHILIPPOUSSIS

Model Response:
Mark Philippoussis

Evaluation (correct/incorrect):

Figure 12: *Guidelines without examples* Prompt template for the judge models

Reference Bias example 2

Your task is to look at the following question, and based on the references provided, determine if the model's response is correct or incorrect. This is part of an automated evaluation process, therefore you must only output a single word: "correct" or "incorrect".

Question:
Q: The oldest known American cocktail, what whiskey forms the base of the cocktail known as a Sazerac?

Model Response:
RYE WHISKEY

References	References
Rye	Secale cereale
Centeio	Rye Extract
Winter rye	Ryep
Rye	Rye (botany)
Ryep	Rye
Secale cereale	Winter rye
Rye (botany)	Centeio
Rye (grain)	Rye (grain)
Rye Extract	Rye
Uramat	Uramat
	Judge Output
Judge Output	
Correct	Incorrect. The correct answer is "bourbon whiskey"

Figure 15: Example of Llama2-7B failing to identify the task by changing the order of the references.

With Guidelines and examples (301 tokens)

You will be given a question, a set of reference answers and the answer given by an LLM. Your task is to judge if the answer given by the LLM is correct, as if you were the LLMs teacher grading their exam. An answer should be counted as correct if it is semantically equivalent to (one of the) reference answers. In doing so, please follow the following guidelines (1-3):

1. Underspecified answers (e.g. "December" instead of "December 20") should be marked incorrect
2. Answers that have more information than requested (e.g. "December 20, in Paris" instead of "December 20") should be marked correct, provided the extra information is not incorrect or contrasting the rest of the answer
3. Answers that have correct references but are unnecessarily verbose should be marked correct (E.g. "Thanks for asking this question! The correct answer is: ...")

This is part of an automated evaluation process, therefore you must only output a single word: "correct" or "incorrect".

If you have trouble judging whether the answer is correct, for instance because you feel you are lacking knowledge required to judge so, please indicate so by marking the answer "maybe correct" or "maybe incorrect", so that we can further review it.

Question:

Which Australian did Roger Federer defeat to win his first Wimbledon Men's Singles title in 2003?

References:

MARK PHILIPPOUSSIS

MARK PHILIPPOUSSIS

Model Response:

Mark Philippoussis

Evaluation (correct/incorrect):

Figure 13: *Guidelines with Examples* Prompt template for the judge models

Reference Bias example 1

Your task is to look at the following question, and based on the references provided, determine if the model's response is correct or incorrect. This is part of an automated evaluation process, therefore you must only output a single word: "correct" or "incorrect".

Question:

Q: Aberdeen is known as what?

Model Response:

Granite City

References

The Granite City

The granite city

Granite City (disambiguation)

The Granite City

Granite City

Judge Output

Incorrect

References

Granite City

Granite City (disambiguation)

The granite city

The Granite City

The Granite City

Judge Output

Correct

Figure 14: Example of Llama2-7B getting confused when the order of the references are changed

N Judge models are sensitive to reference order

We investigate the judges' sensitivity to reference order by providing the same prompt, question and model response to the judge models, but shuffling the reference order in three different permutations. We compute the consistency score of the model as the percentage of questions for which it gives the same judgment all the 3 times. We observe that the model is more likely to evaluate an answer as correct if the corresponding reference appears early in the list of references (see Figure 14). The smaller judge models sometimes fail to capture all the information in the prompt, and provide judgement based on their own knowledge rather than going by the references (see Figure 15).

O Leniency Bias

As described in § 5.4, for the purpose of the leniency bias experiments, we assume that a judge assigns the correct judgment with a probability of P_c and randomly assigns the rest of the samples to be "correct" with a probability P_+ . In this section, we derive the mathematical expressions for P_c and P_+ . We assume that in the case of misalignment between the evaluation criteria of guidelines and judge models, the probability of getting an evaluation of "correct" is independent of the actual correctness of the answer (i.e. the judge model effectively flips a coin to give out its judgement). For any given benchmark and judge model, we denote the ground-truth score as s , and the true positive and true negative rates as t_P and t_N , respectively, all normalized to be between 0 and 1.

Now, based on our assumptions, the true positives, where the exam-taker model response is correct, and also correctly identified by the judge model to be correct, would be comprised of two possible cases: 1) The judge evaluates it correctly according to the given evaluation criteria with a probability of P_c ; and 2) The judge does not evaluate it according to the given criteria with a probability of $1 - P_c$, but the evaluation still happens to be correct with a probability of P_+ . With the total ratio of the correct responses being s , the true positive rate is therefore given by –

$$t_P = s[P_c + (1 - P_c)P_+] \quad (4)$$

Similarly, the true negatives, where the exam-taker model response is incorrect, and also cor-

rectly identified by the judge model to be incorrect, would also be comprised of two cases: 1) The judge evaluates it correctly according to the given evaluation criteria with a probability of P_c . 2) The judge does not evaluate it according to the given criteria with a probability of $1 - P_c$, but the evaluation still happens to be correct with a probability of $1 - P_+$. With the total ratio of the incorrect responses being $1 - s$, the true negative rate is therefore given by –

$$t_N = (1 - s)[P_c + (1 - P_c)(1 - P_+)]. \quad (5)$$

Using Equation (5), we can derive the following.

$$t_N = (1 - s)[P_c + (1 - P_c)(1 - P_+)] \quad (6)$$

$$= P_c + 1 - P_+ - P_c + P_c P_+ \quad (7)$$

$$- s P_c - s + s P_+ + s P_c - s P_c P_+ \quad (8)$$

$$= 1 - P_+ + P_c P_+ - s + s P_+ - s P_c P_+ \quad (9)$$

$$= 1 - s - P_+(1 - P_c - s + s P_c) \quad (10)$$

$$= 1 - s - P_+(1 - s)(1 - P_c) \quad (11)$$

$$\Rightarrow P_+ = \frac{1 - s - t_N}{(1 - s)(1 - P_c)} \quad (12)$$

$$= \frac{1 - \frac{t_N}{1-s}}{1 - P_c} \quad (13)$$

Substituting the value of P_+ in Equation (4), we get:

$$t_P = s[P_c + (1 - P_c)P_+] \quad (14)$$

$$= s \left[P_c + (1 - P_c) \frac{1 - \frac{t_N}{1-s}}{1 - P_c} \right] \quad (15)$$

$$= s \left[P_c + 1 - \frac{t_N}{1 - s} \right] \quad (16)$$

$$\Rightarrow \frac{t_P}{s} = P_c + 1 - \frac{t_N}{1 - s} \quad (17)$$

$$\Rightarrow P_c = \frac{t_P}{s} + \frac{t_N}{1 - s} - 1 \quad (18)$$

The values of P_c and P_+ can be estimated from observed data using the derived expressions. The estimated probabilities using this method, with human evaluation as the reference, are shown in Figure 16a.

To validate these derived values, we observe the correlation between the estimated values of P_c and Scott's Pi (π). As shown in Figure 16b, we observe that the estimated values of P_c are highly correlated to the Scott's π values for the judge models, with a Pearson correlation coefficient of 0.98.

Judge model	π	P_c	P_+
Gemma-2B	0.26	0.38	0.87
Llama2-7B	0.47	0.63	0.75
Llama3-8B	0.59	0.63	0.74
JudgeLM-7B	0.65	0.68	0.19
Mistral-7B	0.66	0.70	0.87
Llama2-70B	0.69	0.66	0.99
Llama2-13B	0.74	0.74	0.87
Llama3.1-8B	0.77	0.77	0.82
GPT-4	0.87	0.87	0.69
Llama3.1-70B	0.88	0.88	0.82
Llama3-70B	0.88	0.87	0.90

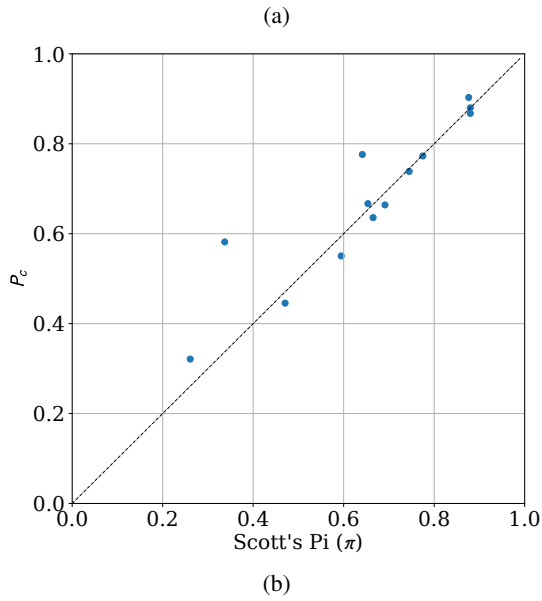


Figure 16: a) Estimated values of P_c and P_+ for different judge models. b) Pearson's correlation coefficient between π and P_c for judge models.

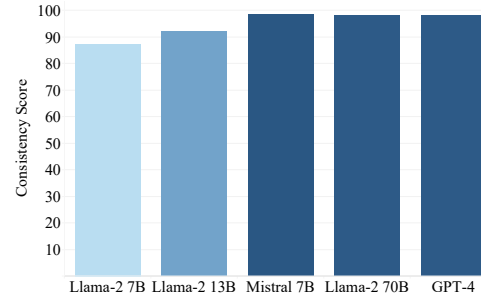


Figure 17: **Leniency bias and answer consistency.** Consistency score, defined as the percentage of questions for which the judge model gives the same judgment for three different answer orders.