

Language Models can Evaluate Themselves via Probability Discrepancy

Anonymous ACL submission

Abstract

In this paper, we initiate our discussion by demonstrating how Large Language Models (LLMs), when tasked with responding to queries, display a more even probability distribution in their answers if they are more adept, as opposed to their less skilled counterparts. Expanding on this foundational insight, we propose a new self-evaluation method ProbDiff for assessing the efficacy of various LLMs. This approach obviates the necessity for an additional evaluation model or the dependence on external, proprietary models like GPT-4 for judgment. It uniquely utilizes the LLMs being tested to compute the probability discrepancy between the initial response and its revised versions. A higher discrepancy for a given query between two LLMs indicates a relatively weaker capability. Our findings reveal that ProbDiff achieves results on par with those obtained from evaluations based on GPT-4, spanning a range of scenarios that include natural language generation (NLG) tasks such as translation, summarization, and our proposed Xiaohongshu blog writing task, and benchmarks for LLM evaluation like AlignBench, MT-Bench, and AlpacaEval, across LLMs of varying magnitudes.

1 Introduction

With the emergence of LLMs like ChatGPT, we’ve witnessed groundbreaking progress in tasks involving instruction following (Wang et al., 2023e), intent comprehension (Lu et al., 2023), and text generation (Zhao et al., 2023). As LLMs evolve at a rapid pace, it becomes crucial to develop solid evaluation frameworks to gauge their performance accurately. Traditional evaluation methods such as BLEU (Papineni et al., 2002) focus mainly on superficial text differences and often fail to align with human judgment (Liu et al., 2023b).

Critique models, trained on binary preference and language critique datasets to simulate human

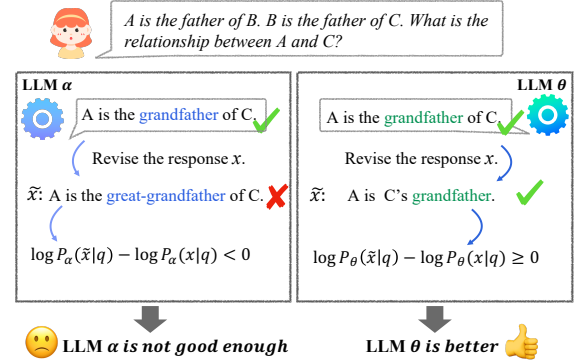


Figure 1: An overview of **ProbDiff**, wherein an LLM iteratively revises its responses, calculating the resulting probability discrepancies as a self-evaluation metric. Larger discrepancies imply decreased confidence in the generated outcomes, with greater variances indicating poorer performance.

judgment (Ouyang et al., 2022), have emerged as a core component of Learning from Human Feedback (Stiennon et al., 2020). However, critique models encounter several challenges. Firstly, they risk becoming outdated as LLMs must continually adapt to novel and intricate tasks, such as writing in specific styles like Xiaohongshu, or solving increasingly complex problems like tool use (Li et al., 2023b) and role play (Lu et al., 2024). Frozen critique models trained on past data struggle to accurately assess such unseen tasks, and it’s impractical to predefine and annotate data for all potential tasks in labeled feedback data. Secondly, there is the issue of *Reward Hacking*, where an LLM may exploit vulnerabilities in the critique models, achieving high rewards without truly fulfilling the intended objectives (Askell et al., 2021; Ramé et al., 2024), resulting in degraded performance, such as linguistically flawed (Lewis et al., 2017) or excessively verbose outputs (Singhal et al., 2023). Thirdly, with each iteration and upgrade of an LLM, there is a need to train additional critique models of matching or greater ca-

capacity, increasing the costs and resource consumption. Another popular option is to allow proprietary models like GPT-4 to score candidate responses via API (Chen et al., 2023; Fu et al., 2023; Ji et al., 2023; Kocmi and Federmann, 2023; Wang et al., 2023a; Zhang et al., 2023). However, GPT-based evaluations entail the risk of potential high costs associated with addressing data leaks (Wang et al., 2023d), and are subject to OpenAI’s terms of use¹. Overall, there currently exists no evaluation method that can fit human preferences without increasing additional training costs, maintain robustness without leaking data, and continually improve alongside advancements in model capabilities.

This paper poses a novel hypothesis: for any query q , an LLM capable of providing accurate response x tend to exhibit a more uniform probability distribution $\log p(x|q)$ compared to their less proficient counterparts. Existing studies indicate that the initial model response x often converges to local maxima of the log probability. when refining x to generate a new response \hat{x} , the relative reduction in $\log p(\hat{x}|q)$ compared to $\log p(x|q)$ is insignificant. Conversely, for LLMs incapable of producing a correct response, $\log p(x|q)$ typically demonstrates higher variance. As a result, revised answers tend to exhibit much lower log probabilities than the original response. We empirically validate this hypothesis and observe its consistency across various LLMs and tasks. Figure 2 and 3 illustrate the visualization of the underlying hypothesis and Section 3.1 showcases the empirical assessment. Exploiting this insight, we introduce ProbDiff, a self-evaluation technique applicable to any LLM across tasks. Given a query q , ProbDiff first prompts a candidate LLM to generate a response x , then asks the LLM to revise x based on q , producing a refined response \hat{x} . Finally, ProbDiff quantifies the probability discrepancy $\log p(x|q) - \log p(\hat{x}|q)$ as the evaluation metric. When comparing two candidate LLMs on q , a larger probability discrepancy indicates a lower proficiency in handling the instruction.

Compared with critique models, ProbDiff utilizes the probability characteristics of LLM to evaluate the LLM’s performance, avoiding the resource waste of training a second model and addressing the issues of critique models’ inability to dynamically adjust with the improvement of

LLM capabilities and task changes. Unlike GPT-based Evaluation, ProbDiff mitigates the risks of data leakage, high API call costs, and usage restrictions of external models. We conduct experiments on conventional generation tasks such as translation and summarization, novel generation tasks like Xiaohongshu blog writing, as well as English and Chinese LLM benchmarks such as AlpacaEval (Li et al., 2023c), MT-bench (Zheng et al., 2023), and AlignBench (Liu et al., 2023a), using Qwen (Bai et al., 2023), LLAMA-2 (Touvron et al., 2023), Yi (01-ai, 2023), WizardLM (Xu et al., 2023) and Tulu (Iverson et al., 2023). The results demonstrate that ProbDiff exhibits highly consistent performance across various tasks and LLMs compared to the currently prevalent GPT-4 based Evaluation methods.

2 Related Work

LLM Evaluation. Current research primarily focuses on model-based evaluation, which involves training another reward model or utilizing external proprietary LLMs such as GPT-4 as judges. Representative methods of the former like PandaLM (Wang et al., 2023d), Shepherd (Wang et al., 2023c), and AUTO-J (Li et al., 2023a) distinguish the most superior model among various candidates and can provide a brief explanation or critique to support their evaluation. However, these models are often relatively small in scale, frozen once trained, and unable to adapt to changes in LLM capabilities, output distributions, and task variations. A more popular approach currently leverages larger and more powerful LLMs like GPT-4 to evaluate numerous natural language generation tasks, including text summarization, machine translation, and so on, showcasing remarkable performance (Liu et al., 2023b; Wang et al., 2023a; Kocmi and Federmann, 2023). However, subsequent investigations have unveiled certain issues with LLM evaluators, particularly concerning biases related to position and verbosity (Wang et al., 2023b; Zheng et al., 2023). Despite researchers adopting multi-agent strategies and making multiple calls to GPT models to mitigate bias (Zhang et al., 2023; Chan et al., 2023), they still face challenges such as data leakage and high evaluation costs. In this paper, we propose for the first time the concept of probability discrepancy to enable LLMs to conduct self-evaluation from a probabilistic perspective, without requiring train-

¹<https://openai.com/policies/terms-of-use>

ing extra models or invoking external ones. It allows for the evaluation of model effectiveness on any task using any LLM and achieves results similar to current mainstream GPT-4 evaluations.

Probability Exploitation. Probability has played a pivotal role in both detecting inconsistencies (She et al., 2023; Jia et al., 2023) and identifying machine-generated text (Tang et al., 2023; Mitchell et al., 2023; Bao et al., 2023). Zero-shot faithfulness evaluation with a foundation language model (FFLM), introduced by Jia et al. (2023), combines changes in probability based on the premise that pre-adding a text segment consistent with the output increases the likelihood of predicting the desired output. Controlling the preference of the generation model with the help of the prompt (CoP) (She et al., 2023), on the other hand, leverages the disparity in probability between a document and a coherent text prompt during inference to detect factual inconsistencies. In the realm of detecting machine-generated text, DetectGPT (Mitchell et al., 2023) posits that unlike human-authored text, model-generated text tends to cluster in local maxima of the log probability. Building upon this insight, Fast-DetectGPT (Bao et al., 2023) utilizes conditional probability curvature to reveal disparities in word choices between LLMs and human authors within a given context. Inspired by DetectGPT, our work further analyzes the local structure of the generation probability function for LLMs capable and incapable of handling specific queries. We observed that proficient LLMs exhibit a flatter distribution, whereas less capable ones exhibit a steep distribution. Consequently, we propose evaluating model performance through repeated sampling and assessing the discrepancy in log probability between samples. Through this novel probabilistic perspective, we successfully offer a new approach to LLM evaluation.

3 Method

In this section, we present our self-evaluation protocol, ProbDiff, designed to self-assess the capabilities of LLMs autonomously through logarithmic probability changes. We begin with preliminary studies to validate the prior assumptions of ProbDiff, followed by a detailed exposition of its intricacies and mechanisms.

3.1 Hypothesis Validation

The first assumption underlying ProbDiff is that when we question an LLM, the more capable it is of providing an answer, the more confident it should be in its response. Therefore, the proportion of consistent or approximate answers in repeated questioning of the same question should be higher. However, when the model lacks proficiency in a specific domain, it should exhibit considerable uncertainty in its responses. Thus, there would be a significant variance in the answers provided upon repeated questioning. To test this hypothesis, we conducted experiments using two widely used LLMs: GPT-3.5 and GPT-4. It is well-known that GPT-4 surpasses GPT-3.5 in various dimensions of capability. If our hypothesis holds true, then for the same set of questions asked multiple times to both GPT-4 and GPT-3.5, the similarity between the answers provided by GPT-4 should be higher.

We employed MT-Bench as our test dataset. Since MT-Bench is a multi-turn dataset, we utilized only the first-turn query, and requested these LLMs to provide answers three times. Specifically, upon receiving the first answer from the LLM, we continued to prompt with "Answer this question again." to collect the model's second answer, and so forth until obtaining the third answer.

	GPT-3.5-turbo		GPT-4	
	Similarity	EM	Similarity	EM
$x - x_1$	86.88	21.25	92.25	26.25
$x - x_2$	90.28	18.75	90.71	25.00

Table 1: The preliminary study aims to verify the disparities in similarity between responses to the same query from a strong model (GPT-4) and a weak model (GPT-3.5-turbo). Here, x represents the initial response, x_1 denotes the re-generated response, and x_2 is the response regenerated based on x_1 . Similarity refers to the similarity score between two responses, while EM stands for Exact Match, indicating the proportion of responses that are identical in the dataset. This study was conducted using the MT-Bench benchmark.

As depicted in Table 1, the similarity between the first and second responses from GPT-3.5, compared to GPT-4, was lower by 5.37 pts. Additionally, we observed that in both GPT-3.5 and GPT-4, a substantial proportion of responses remained exactly the same across the two rounds, constituting

an exact match (EM). For instance, 21.5% of GPT-3.5’s first and second round responses were consistent, whereas this proportion increased to 26.25% in GPT-4. Comparing the first and third round responses, although the similarity score of GPT-3.5 improved, there was still a significant difference in exact match proportions compared to GPT-4, thus validating our hypothesis.

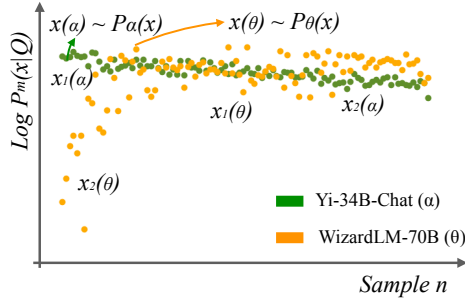


Figure 2: Log probability curves of the responses for Yi-34B-Chat and WizardLM-70B on AlpacaEval-2.0.

So, why does this phenomenon exist? We posit that it’s because stronger LLMs assign higher probabilities to each token in the answer output, resulting in greater determinism and thus smaller variance in the decoded responses. To investigate this conjecture, we conducted a second preliminary study. We selected the AlpacaEval-2.0 benchmark along with the best-performing open-source model on the benchmark, Yi-34B-Chat (Score 29.66%), and a relatively poorer-performing model, WizardLM 70B (Score 14.38%). We randomly sampled 10 queries from AlpacaEval-2.0, with each query sampled for 100 responses from each model. We plotted the log probability curves of the responses for each model, as shown in Figure 2. It is evident that the log probability curve of Yi-34B-Chat is flatter, indicating smaller variance, with nearly all response output probabilities being approximately equal.

Combining this conclusion with the observations of DetectGPT (Mitchell et al., 2023) forms the theoretical foundation of ProbDiff. DetectGPT proposes that samples x from an LLM typically reside in regions of negative curvature of the log probability function, where neighboring samples like x_1 and x_2 exhibit lower model log probability on average. In other words, the log probability of the responses initially sampled by the model tends to be at local maxima of the probability function. As illustrated in Figure 3, we depict a simulated probability function. Due to the fact that the log




Figure 3: We recognize and leverage the observation that superior LLMs typically exhibit smaller probability variances, along with the conclusion that the model-generated samples tend to reside in regions of negative curvature within the probability function. These findings serve as crucial distinctions for ProbDiff in discerning between models of varying capabilities.

probability $p(x|q)$ for a specific query q is comparatively flatter for better LLM than for worse LLM, we can infer that models with larger probability discrepancy, after revising their responses multiple times, will perform worse on the current query.

3.2 Self-evaluation Protocol ProbDiff

Response Refinement Prompt

I want you to act as a Response Rewriter. Your goal is to enhance the quality of the response given by an AI assistant to the #Given Prompt# through rewriting. But the rewritten response must be reasonable and must be understood by humans. Your rewriting cannot omit the non-text parts such as the emoji in #Given Prompt# and #Given Response#. If you think the response is already great enough, you can keep it unchanged. You should try your best not to change the length of the response. #Given Response# and #Rewritten Response# are not allowed to appear in #Rewritten Response#. #Given Prompt#: {{prompt}} #Given Response#: {{response}} #Rewritten Response#:

Based on our validated assumptions, ProbDiff employs a remarkably simple yet elegant strategy for LLM evaluation. Specifically, given any LLM α , and a query q under evaluation, we first require

α to generate a response x containing T tokens according to its decoding strategy, with the average per-token log probability $\log; p_\alpha(x|q)$ defined as:

$$\log; p_\alpha(x|q) = \frac{1}{T} \sum_{t=1}^T p_\alpha(x_t|q, x_{<t}). \quad (1)$$

Subsequently, we prompt LLM α to revise the response based on q and x , obtaining x_1 , and iterate this revision process K times to obtain the final response x_K , with the revision prompt defined as the Response Refinement Prompt. Finally, we calculate the log-probability discrepancy between x and x_K as:

$$d(\alpha, q) = \log p_\alpha(x_K|q) - \log p_\alpha(x|q). \quad (2)$$

We aim to characterize the variance of LLM α 's probability distribution for q using $d(\alpha, q)$, although employing multiple evaluations for averaging could yield more accurate estimates. However, in this paper, we present a preliminary exploration where we adopt the simplest characterization by directly computing the discrepancy between the two probabilities. Subsequently, for the evaluation set \mathbb{D} , we can compute the average $d(\alpha, q)$ for all q in the set, yielding $d(\alpha, \mathbb{D}) = \frac{1}{|\mathbb{D}|} \sum_{q \in \mathbb{D}} d(\alpha, q)$ as the probability discrepancy score of LLM α over \mathbb{D} . Ultimately, for the two LLMs, α and θ , under evaluation, we can compare $d(\alpha, \mathbb{D})$ and $d(\theta, \mathbb{D})$. A larger $d(\alpha, \mathbb{D})$ implies a higher variance in the probability distribution of the generation, indicating lower confidence in its responses and thus poorer performance under our hypothesis. Furthermore, if one wishes to statistically analyze the win/tie/lose ratios between α and θ , one can establish predefined thresholds for $d(\alpha, q)$ and $d(\theta, q)$ to determine when their performances are comparable within a certain discrepancy range. Beyond this range, a judgment can be made regarding which model performs better.

In addition, we use Response Refinement Prompt to guide the model to revise the original responses, instead of using multiple sampling to calculate log probability discrepancies. This is mainly because greedy methods are usually used to decode in math and code tasks, such as code and math tasks in MT-Bench. Multiple sampling for this decoding method is meaningless, and using prompt to guide the model to revise the answer is suitable for any decoding methods.

4 Experiment

In this section, we evaluate our ProbDiff methodology across three diverse natural language generation tasks: Translation, Summarization, and a niche task designated as Xiaohongshu blog writing. To gauge ProbDiff's efficacy, we fine-tuned the *Qwen-14B-Chat* model for each specified task. Additionally, our evaluation encompasses three benchmarks tailored for LLM assessment: MT-Bench (Zheng et al., 2023), AlpacaEval (Li et al., 2023c), and AlignBench (Liu et al., 2023a). These experiments are designed to comprehensively evaluate ProbDiff's performance across a wide array of linguistic tasks and benchmarks, showcasing its applicability in measuring LLM effectiveness

4.1 Datasets

For the **Translation** task, we selected the **WMT19** dataset (Ng et al., 2019) for Chinese-English (*Zh-En*), Czech-English (*Cs-En*), German-English (*De-En*), and Russian-English (*Ru-En*) translation tasks. We randomly chose 5,000 translation pairs for our training sets and 1,000 pairs for our test sets, ensuring an equitable distribution between source-to-English and English-to-source language pairs. Regarding the **Summarization** task, we opted for the **XSum** (Narayan et al., 2018) and **CNN_DM** (Hermann et al., 2015) datasets, randomly selecting 5,000 documents for the training set and 1,000 documents for the test set from each dataset.

We introduce an innovative **Xiaohongshu blog writing task** as a novel natural language generation task to assess the instruction-following capabilities of LLMs. For this, we compiled a new benchmark by gathering 569 blogs across various topics from the Xiaohongshu APP, a lifestyle platform where users share product recommendations or opinions. We randomly chose 469 instances for model fine-tuning and 100 instances to evaluate model performance. Detailed descriptions and example instances of the new dataset are available in Appendix A.

To evaluate the alignment capability of LLMs across diverse dimensions, we undertook experiments with AlignBench (Liu et al., 2023a), a comprehensive dataset designed to assess LLMs' alignment capabilities in Chinese. It categorizes use cases into eight principal domains: Fundamental Language Abilities, Chinese Advanced Understanding, Open-ended Questions, Writing Ability,

	Cs-En	De-En	Ru-En	Zh-En	Cnn_dm	Xsum	Blog
Qwen	1.0	0.3	0.6	3.0	8.4	6.0	55.0
Qwen _{ft}	91.0	93.4	92.5	91.0	63.0	69.4	57.0

Table 2: *Confidence* (%) obtained through ProbDiff in the translation, summarization, and Xiaohongshu blog writing tasks. “Qwen” represents the Qwen-14B-Chat model with different parameter sizes, and “Qwen_{ft}” denotes the fine-tuned model based on Qwen-14B-Chat with different parameter sizes. The best performance in each column is bold.

	Overall	Pro.	Math.	Fund.	Logic.	Writ.	Chi.	Role.	Open.
Qwen	30.3	21.8	24.3	34.8	26.3	61.3	12.1	39.7	39.5
Qwen _{ft}	32.5	19.4	39.6	30.4	47.4	38.7	17.2	33.6	28.9

Table 3: *Confidence* (%) obtained through the ProbDiff in AlignBench. “Pro.” denotes Professional Knowledge, “Math.” denotes Mathematics, “Fund.” denotes Fundamental Language Ability, “Logic.” denotes Logical Reasoning, “Writ.” denotes Writing Ability, “Chi.” denotes Advanced Chinese Understanding, “Role.” denotes Task-oriented Role Play, and “Open.” denotes Open-ended Questions. The best performance in each column is bold.

Logical Reasoning, Mathematics, Task-oriented Role Play, and Professional Knowledge, encompassing 683 samples in total. We used GPT-4 to generate 50 text samples per use case and category pair, and selected 15 generated samples per pair to eliminate duplicates, culminating in a collection of 10,245 generated samples for fine-tuning Qwen-14B-Chat.

In addition to specific task fine-tuning of Qwen-14B-Chat, we also conducted experiments on models listed on the existing LLM benchmark leaderboards. We choose MT-bench² and AlpacaEval 2.0³ as our evaluation benchmarks. **MT-bench**(Zheng et al., 2023) is tailored to examine multi-turn conversation and instruction-following capabilities through 80 high-quality multi-turn questions. **MT-bench**(Zheng et al., 2023) comprises 805 single-turn questions, formulated from Alpaca Data.

4.2 Implementation Details

For all experiments except those on the MT-Bench and AlignBench datasets, we configured the temperature and top_p parameters to 0.7 and 0.8, respectively, to generate initial responses. These two benchmarks possess officially recommended generation settings, which we adhered to for relevant experiments. To refine the generated responses, we adjusted the temperature to 0.1 and top_p to

0.9.

The effectiveness of an LLM is evaluated based on the confidence level of its responses. This metric quantifies the degree of probability enhancement observed when the model’s initial response undergoes refinement. The calculation is as follows:

$$confidence = \frac{1}{|\mathbb{D}|} \sum_{q \in \mathbb{D}} \phi \quad (3)$$

$$\phi = \begin{cases} 1, & d(\alpha, q) \geq \delta \\ 0, & d(\alpha, q) < \delta \end{cases} \quad (4)$$

where δ is the threshold. For NLG tasks, we set δ to -0.02 to reduce errors. For more challenging LLM’s alignment evaluation tasks such as Alignbench, MT-Bench and AlpacaEval we appropriately relax δ to -0.05 .

4.3 Results

Natural Language Generation Evaluation: We applied ProbDiff to assess *Qwen-14B-Chat*’s efficacy in tasks such as translation, summarization, and Xiaohongshu blog writing. Table 2 illustrates the *confidence* levels in the generated sentences by both the original and fine-tuned Qwen models across these tasks. Subsequent to the fine-tuning process, it was noted that the fine-tuned model exhibited increased confidence in its responses upon revision, aligning with our anticipatory hypotheses. Specifically, in the Xiaohongshu blog writing

²<https://huggingface.co/spaces/lmsys/chatbot-arena-leaderboard>

³https://tatsu-lab.github.io/alpaca_eval

	MT-Bench		AlpacaEval	
	ProbDiff	Official	ProbDiff	Official
Llama2	32.50	6.86	4.10	13.87
Yi	41.25	7.58	7.58	29.66
WizardLM	42.50	7.71	5.47	14.38
Tulu	53.75	7.89	16.27	15.98

Table 4: ProbDiff and Official results on MT-Bench and AlpacaEval. The metrics of ProbDiff is *confidence*(%), MT-Bench official metrics is model scores(0-10) judged by GPT-4, and AlpacaEval 2.0 official metrics is win rate(%).

task tailored for LLMs with a 14B parameter size, we encountered a scarcity of data for fine-tuning. Nonetheless, post-fine-tuning, there was a noticeable enhancement in response confidence, as depicted in Table 2. From these results, we deduce that fine-tuning LLMs on task-specific datasets invariably boosts the confidence of their responses, markedly surpassing the performance of the pre-fine-tuned LLM. This elevation in response confidence serves as a pertinent indicator for gauging improvements in LLM performance.

LLM’s Alignment Evaluation: To ascertain the alignment capabilities of LLMs across diverse dimensions, we conducted fine-tuning of *Qwen-14B-Chat* using documents synthesized by GPT-4, specifically targeting the AlignBench dataset. The outcomes of our evaluation are presented in Table 3.

AlignBench serves as a meticulously curated benchmark designed to assess the alignment capabilities of LLMs. This benchmark poses a more complex challenge compared to standard NLG tasks, requiring nuanced understanding and response accuracy from the model. As depicted in Table 3, post fine-tuning with task-specific data, the model demonstrates notable confidence across most evaluated alignment abilities. However, it is observed that in categories such as writing ability and open-ended questions, the model’s confidence tends to diminish, suggesting areas for further improvement and investigation.

Evaluation on other LLMs: To assess the generalization capabilities of our ProbDiff, we extended our evaluations to include several high-performing LLMs featured on established LLM benchmark leaderboards, specifically MT-Bench and AlpacaEval 2.0. We employed ProbDiff to evaluate a selection of models, including Yi-34B-

Chat (01-ai, 2023), tulu-2-dpo-70B (Iverson et al., 2023), WizardLM-70B-V1.0 (Xu et al., 2023), and Llama2-70B-chat (Touvron et al., 2023). Table 4 presents the response confidence of these LLMs alongside their official benchmarks. Notably, Yi-34B-Chat lacks official MT-Bench scores, thus we conducted its evaluation using the benchmark’s official protocol.




Figure 4: Evaluate the validity of the Qwen-14B-Chat and Qwen-14B-Chat_ft through GPT-4 in AlignBench(Align) and Xiaohongshu Blog Writing(Blog) tasks. “Align_gpt” and “Blog_gpt” represents the win rate judged by GPT-4, “Align_prob” and “Blog_prob” represents the confidence evaluate by ProbDiff. Orange histogram indicates fine-tuned Qwen-14B-Chat and blue histogram indicates Qwen-14B-Chat, respectively.

Within the MT-Bench framework, the official ranking of these LLMs is Tulu, WizardLM, Yi, and Llama2, a sequence that aligns perfectly with the model confidence hierarchy as determined by ProbDiff, as illustrated in table 4. Conversely, for AlpacaEval 2.0, the official ranking based on win rate is Yi, Tulu, WizardLM, and Llama2. Our analysis, however, positions the models as Tulu, Yi, WizardLM, and Llama2, showcasing a slight variation from the official results.

5 Discussion

5.1 Confidence vs GPT-4 scores

In this section, to substantiate our premise that enhancements in model efficacy are discernible through shifts in log probability, we juxtaposed the congruence between our evaluative outcomes and those derived from utilizing GPT-4 as an adjudicator. To juxtapose the evaluative outcomes of ProbDiff with those of GPT-4, we engaged GPT-4 to directly appraise which output from the two models prevails, eschewing the assignment of a definitive

		Overall	Pro.	Math.	Fund.	Logic.	Writ.	Chi.	Role.	Open.
GPT-4	Qwen	5.73	6.20	4.83	6.85	5.14	6.56	6.31	6.49	6.45
	Qwen _{ft}	6.62	6.19	5.62	6.49	6.93	7.80	5.86	7.73	7.71
CritiqueLLM	Qwen	5.22	5.78	4.19	6.24	4.63	6.13	5.78	6.13	6.13
	Qwen _{ft}	6.12	5.57	5.15	6.25	6.34	7.15	5.33	7.11	7.53
ProbDiff	Qwen	30.3	21.8	24.3	<u>34.8</u>	26.3	61.3	12.1	39.7	39.5
	Qwen _{ft}	32.5	19.4	39.6	30.4	47.4	38.7	17.2	33.6	28.9

Table 5: The scores(0-10) obtained through GPT-4 and CritiqueLLM of AlignBench, and the response confidence(%) obtained through ProbDiff. “Qwen” represents the Qwen-14B-Chat model, and “Qwen_{ft}” denotes the fine-tuned model based on Qwen-14B-Chat. The bold number indicates that ProbDiff is consistent with both GPT-4 and CritiqueLLM, while the underlined number indicates that ProbDiff is consistent with either GPT-4 or CritiqueLLM.

score. Figure 4 illustrates that when ProbDiff adjudicates the fine-tuned model as superior to its baseline counterpart, GPT-4’s evaluation of the outputs from these models aligns with the determinations made by ProbDiff.

As AlignBench has provided the official evaluation scripts, which can generate the scores of each model by GPT-4, or other evaluation LLM, such as CritiqueLLM(Ke et al., 2023), we also compared the results in table 5. To perform a systematic evaluation, AlignBench framed a comprehensive taxonomy of the abilities of LLMs, which inspects and summarizes their use cases into 8 main categories. We also show the details of the comparison results in table 5. We can observe that although there exist discrepancies between the results yielded by ProbDiff and those by GPT-4 in certain types of assessments, such as text composition, our ProbDiff generates consistent results with GPT-4 and CritiqueLLM in most cases.

5.2 ProbDiff vs GPT-4 and Manual Evaluation

When we ask LLMs to modify their responses, we use the log probability discrepancy to determine whether the revised responses are better. We test the consistency by measuring the degree of conflict among ProbDiff, GPT-4, and manual evaluation. We need to illustrate that from the probability perspective, the model always outputs the response with the highest probability. When we ask an LLM to modify its outputs, if some subtle changes appear, such as synonym replacement, the log probability of the revised responses will be significantly reduced. However, from the perspective of the response quality, these subtle changes will not affect the overall response quality, GPT-

4 and manual evaluation will not treat these subtle changes as an influential factor that can change the response quality. Therefore, the probability is indeed a strict indicator, leading to the low correlation coefficient among ProbDiff, GPT-4, and manual evaluation.

We employed GPT-4 and manual evaluation to check whether our ProbDiff will generate outputs that conflict with GPT-4 and manual evaluation. we aim to examine when ProbDiff yields $d(\alpha, q) > 0$, whether GPT-4 and human annotator will also confirm the quality of the response is improved. Or when ProbDiff yields $d(\alpha, q) < 0$, whether GPT-4 and human annotator will confirm the quality of the response is degraded. We conduct this experiment on AlignBench and MT-Bench, and the results are shown in table 6. As demonstrated in table 6, ProbDiff has a low conflict degree with GPT-4 and manual evaluation.

	GPT-4	Humans
AlignBench	3.84	4.49
MT-Bench	6.45	4.84

Table 6: Conflict degree(%) among ProbDiff, GPT-4, and manual evaluation. We employ 3 human annotators to report the average results.

6 Conclusion

In this paper, we propose a novel evaluation protocol ProbDiff to enable the LLMs to conduct self-evaluation based on the probability discrepancy of the generated content. Moreover, even in the absence of prior training and a benchmark for a specific domain, our ProbDiff approach can effectively assess the capabilities of LLMs and mitigate the potential risks of data leakage.

Limitations

The constraints of our methodology are outlined as follows: Firstly, while our approach is capable of ascertaining enhancements in LLM performance, it lacks the capacity to offer quantitative metrics indicating the magnitude of improvement. Secondly, variations in log probability are occasionally affected by sentence length. Should there be a notable disparity in length between the original and modified sentences, the log probability will undergo substantial changes.

References

01-ai. 2023. <https://github.com/01-ai/Yi>.

Amanda Askell, Yuntao Bai, Anna Chen, Dawn Drain, Deep Ganguli, Tom Henighan, Andy Jones, Nicholas Joseph, Ben Mann, Nova DasSarma, et al. 2021. A general language assistant as a laboratory for alignment. *arXiv preprint arXiv:2112.00861*.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Guangsheng Bao, Yanbin Zhao, Zhiyang Teng, Linyi Yang, and Yue Zhang. 2023. Fast-detectgpt: Efficient zero-shot detection of machine-generated text via conditional probability curvature. *arXiv preprint arXiv:2310.05130*.

Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2023. Chateval: Towards better llm-based evaluators through multi-agent debate. *arXiv preprint arXiv:2308.07201*.

Yi Chen, Rui Wang, Haiyun Jiang, Shuming Shi, and Ruifeng Xu. 2023. [Exploring the use of large language models for reference-free text quality evaluation: A preliminary empirical study](#).

Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. Gptscore: Evaluate as you desire. *arXiv preprint arXiv:2302.04166*.

Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *Advances in neural information processing systems*, 28.

Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. 2023. [Camels in a changing climate: Enhancing lm adaptation withulu 2](#).

Yunjie Ji, Yan Gong, Yiping Peng, Chao Ni, Peiyan Sun, Dongyu Pan, Baochang Ma, and Xiangang Li. 2023. [Exploring chatgpt’s ability to rank content: A preliminary study on consistency with human preferences](#).

Qi Jia, Siyu Ren, Yizhu Liu, and Kenny Q Zhu. 2023. Zero-shot faithfulness evaluation for text summarization with foundation language model. *arXiv preprint arXiv:2310.11648*.

Pei Ke, Bosi Wen, Zhuoer Feng, Xiao Liu, Xuanyu Lei, Jiale Cheng, Shengyuan Wang, Aohan Zeng, Yuxiao Dong, Hongning Wang, et al. 2023. Critiquellm: Scaling llm-as-critic for effective and explainable evaluation of large language model generation. *arXiv preprint arXiv:2311.18702*.

Tom Kocmi and Christian Federmann. 2023. Large language models are state-of-the-art evaluators of translation quality. *arXiv preprint arXiv:2302.14520*.

Mike Lewis, Denis Yarats, Yann N Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning for negotiation dialogues. *arXiv preprint arXiv:1706.05125*.

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*.

Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023b. Api-bank: A comprehensive benchmark for tool-augmented llms. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3102–3116.

Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023c. AlpacaEval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval.

Xiao Liu, Xuanyu Lei, Shengyuan Wang, Yue Huang, Zhuoer Feng, Bosi Wen, Jiale Cheng, Pei Ke, Yifan Xu, Weng Lam Tam, et al. 2023a. Alignbench: Benchmarking chinese alignment of large language models. *arXiv preprint arXiv:2311.18743*.

Yang Liu, Dan Iter, Yichong Xu, Shuhang Wang, Ruochen Xu, and Chenguang Zhu. 2023b. Gpteval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.

Keming Lu, Bowen Yu, Chang Zhou, and Jingren Zhou. 2024. Large language models are superpositions of all characters: Attaining arbitrary role-play via self-alignment. *arXiv preprint arXiv:2401.12474*.

Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, Chang Zhou, and Jingren Zhou. 2023. # InsTag: Instruction tagging for analyzing supervised fine-tuning of large language models. *arXiv e-prints*, pages arXiv–2308.

684	Eric Mitchell, Yoonho Lee, Alexander Khazatsky,	Jiaan Wang, Yunlong Liang, Fandong Meng, Haoxiang	739
685	Christopher D Manning, and Chelsea Finn. 2023.	Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou.	740
686	Detectgpt: Zero-shot machine-generated text detec-	2023a. Is chatgpt a good nlg evaluator? a prelimi-	741
687	tion using probability curvature. <i>arXiv preprint</i>	nary study. <i>arXiv preprint arXiv:2303.04048</i> .	742
688	<i>arXiv:2301.11305</i> .		
689	Shashi Narayan, Shay B. Cohen, and Mirella Lapata.	Peiyi Wang, Lei Li, Liang Chen, Dawei Zhu, Binghuai	743
690	2018. Don't give me the details, just the summary!	Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang	744
691	Topic-aware convolutional neural networks for ex-	Sui. 2023b. Large language models are not fair eval-	745
692	treme summarization. In <i>Proceedings of the 2018</i>	uators .	746
693	<i>Conference on Empirical Methods in Natural Lan-</i>		
694	<i>guage Processing</i> , Brussels, Belgium.	Tianlu Wang, Ping Yu, Xiaoqing Ellen Tan, Sean	747
695	Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott,	O'Brien, Ramakanth Pasunuru, Jane Dwivedi-Yu,	748
696	Michael Auli, and Sergey Edunov. 2019. Face-	Olga Golovneva, Luke Zettlemoyer, Maryam Fazel-	749
697	book fair's wmt19 news translation task submission.	Zarandi, and Asli Celikyilmaz. 2023c. Shepherd: A	750
698	<i>arXiv preprint arXiv:1907.06616</i> .	critic for language model generation. <i>arXiv preprint</i>	751
699		<i>arXiv:2308.04592</i> .	752
700	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida,	Yidong Wang, Zhuohao Yu, Zhengran Zeng, Linyi	753
701	Carroll Wainwright, Pamela Mishkin, Chong Zhang,	Yang, Cunxiang Wang, Hao Chen, Chaoya Jiang,	754
702	Sandhini Agarwal, Katarina Slama, Alex Ray, et al.	Rui Xie, Jindong Wang, Xing Xie, et al. 2023d.	755
703	2022. Training language models to follow instruc-	Pandalm: An automatic evaluation benchmark for	756
704	tions with human feedback. <i>Advances in Neural In-</i>	llm instruction tuning optimization. <i>arXiv preprint</i>	757
705		<i>arXiv:2306.05087</i> .	758
706	Kishore Papineni, Salim Roukos, Todd Ward, and Wei-	Yufei Wang, Wanjun Zhong, Liangyou Li, Fei Mi,	759
707	Jing Zhu. 2002. Bleu: a method for automatic eval-	Xingshan Zeng, Wenyong Huang, Lifeng Shang,	760
708	uation of machine translation. In <i>Proceedings of the</i>	Xin Jiang, and Qun Liu. 2023e. Aligning large lan-	761
709	<i>40th annual meeting of the Association for Comput-</i>	guage models with human: A survey. <i>arXiv preprint</i>	762
710	<i>ational Linguistics</i> , pages 311–318.	<i>arXiv:2307.12966</i> .	763
711	Alexandre Ramé, Nino Vieillard, Léonard Hussenot,	Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng,	764
712	Robert Dadashi, Geoffrey Cideron, Olivier Bachem,	Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin	765
713	and Johan Ferret. 2024. Warm: On the benefits	Jiang. 2023. Wizardlm: Empowering large lan-	766
714	of weight averaged reward models. <i>arXiv preprint</i>	guage models to follow complex instructions. <i>arXiv</i>	767
715	<i>arXiv:2401.12187</i> .	<i>preprint arXiv:2304.12244</i> .	768
716	Shuaijie She, Xiang Geng, Shujian Huang, and Jiajun	Xinghua Zhang, Bowen Yu, Haiyang Yu, Yangyu Lv,	769
717	Chen. 2023. Cop: Factual inconsistency detection	Tingwen Liu, Fei Huang, Hongbo Xu, and Yongbin	770
718	by controlling the preference. In <i>Proceedings of</i>	Li. 2023. Wider and deeper llm networks are fairer	771
719	<i>the AAAI Conference on Artificial Intelligence</i> , vol-	llm evaluators. <i>arXiv preprint arXiv:2308.01862</i> .	772
720	ume 37, pages 13556–13563.	Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xi-	773
721	Prasann Singhal, Tanya Goyal, Jiacheng Xu, and	aolei Wang, Yupeng Hou, Yingqian Min, Beichen	774
722	Greg Durrett. 2023. A long way to go: Investi-	Zhang, Junjie Zhang, Zican Dong, et al. 2023. A	775
723	gating length correlations in rlhf. <i>arXiv preprint</i>	survey of large language models. <i>arXiv preprint</i>	776
724	<i>arXiv:2310.03716</i> .	<i>arXiv:2303.18223</i> .	777
725	Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan	778
726	Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford,	Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin,	779
727	Dario Amodei, and Paul F Christiano. 2020. Learn-	Zhuohan Li, Dacheng Li, Eric. P Xing, Hao Zhang,	780
728	ing to summarize with human feedback. <i>Advances</i>	Joseph E. Gonzalez, and Ion Stoica. 2023. Judging	781
729	<i>in Neural Information Processing Systems</i> , 33:3008–	llm-as-a-judge with mt-bench and chatbot arena .	782
730			
731	Ruixiang Tang, Yu-Neng Chuang, and Xia Hu. 2023.		
732	The science of detecting llm-generated texts. <i>arXiv</i>		
733	<i>preprint arXiv:2303.07205</i> .		
734	Hugo Touvron, Louis Martin, Kevin Stone, Peter		
735	Albert, Amjad Almahairi, Yasmine Babaei, Niko-		
736	lay Bashlykov, Soumya Batra, Prajjwal Bhargava,		
737	Shruti Bhosale, et al. 2023. Llama 2: Open foun-		
738	dation and fine-tuned chat models. <i>arXiv preprint</i>		
	<i>arXiv:2307.09288</i> .		

A Xiaohongshu Datasets

According to the blog format and text style of Xiaohongshu App, we manually wrote 569 blogs, covering various functional topics such as life, beauty, travel, recommendations, etc., and constructed a instruction following dataset. It is characterized by the addition of rich emoji elements. The data format is shown in figure 5

USER:

写一篇去迪士尼看小熊维尼小红书文案70字

Write a 70-word Xiaohongshu blog about going to Disney to see Winnie the Pooh.

ASSISTANT:

标题: 【#迪士尼#小熊维尼】的童年梦, 你还记得吗?

正文: 🌈迪士尼, 一个让人心动的地方, 小熊维尼更是我们的童年回忆。🍯这只可爱的熊熊, 总是喜欢吃蜂蜜, 一不小心就陷入蜂蜜罐, 让人忍俊不禁😂。亲爱的, ❤️快来迪士尼, 一起找寻童年的足迹👣, 和小熊维尼一起度过一个快乐的一天吧! #迪士尼#小熊维尼

Title: [#Disney#Winnie the Pooh]'s childhood dream, do you still remember it?

Text: 🌈Disney is a heart-warming place, and Winnie the Pooh is our childhood memory. 🍯This cute bear always likes to eat honey. It accidentally falls into the honey jar, which makes people laugh 😂. Dear, ❤️ Come to Disney, look for the traces of childhood together👣, and spend a happy day with Winnie the Pooh! #Disney#winnie the pooh)

Figure 5: Xiaohongshu blog writing data.

B Threshold analysis

In this section, we analyze the impact of different thresholds on response confidence. We conducted experiments in Alignbench and set the thresholds to 0, -0.05, and -0.1 to observe the impact of the thresholds on the experimental results.

As shown in table 7, when the absolute value of the threshold becomes higher, the confidence of the model also becomes higher. In the case of Align-Bench, potential errors and biases may be present in the training data due to its generation by LLMs. This will lead to some shortcomings in the fine-tuned model itself. Therefore, it can also be seen from table 3 that when we increase the threshold appropriately, the model effect will also become better.

		Overall	Pro.	Math.	Fund.	Logic.	Writ.	Chi.	Role.	Open.
0	Qwen	17.7	10.5	6.3	11.6	18.9	44.0	6.9	24.1	26.3
	Qwen _{ft}	27.4	15.3	32.4	27.5	42.1	26.7	15.5	30.2	23.7
-0.05	Qwen	30.3	21.8	24.3	34.8	26.3	61.3	12.1	39.7	39.5
	Qwen _{ft}	32.5	19.4	39.6	30.4	47.4	38.7	17.2	33.6	28.9
-0.1	Qwen	37.5	21.0	46.8	34.8	49.5	44.0	17.2	51.7	36.8
	Qwen _{ft}	40.3	25.0	37.8	26.1	36.8	73.3	15.5	54.3	55.3

Table 7: *Confidence (%)* obtained through the ProbDiff in AlignBench with different δ . “Pro.” denotes Professional Knowledge, “Math.” denotes Mathematics, “Fund.” denotes Fundamental Language Ability, “Logic.” denotes Logical Reasoning, “Writ.” denotes Writing Ability, “Chi.” denotes Advanced Chinese Understanding, “Role.” denotes Task-oriented Role Play, and “Open.” denotes Open-ended Questions.