FB-Bench: A Fine-Grained Multi-Task Benchmark for Evaluating LLMs' Responsiveness to Human Feedback

Youquan Li[1], Miao Zheng[2], Fan Yang[2], Guosheng Dong[2]
Bin Cui[1], Weipeng Chen[2], Zenan Zhou[2*], Wentao Zhang[1*]

[1]Peking University [2]Baichuan Inc.
youquan.li@stu.pku.edu.cn
zhouzenan@baichuan-inc.com, wentao.zhang@pku.edu.cn

Abstract

Human feedback is crucial in the interactions between humans and Large Language Models (LLMs). However, existing research primarily focuses on benchmarking LLMs in single-turn dialogues. Even in benchmarks designed for multi-turn dialogues, the user utterances are often independent, neglecting the nuanced and complex nature of human feedback within real-world usage scenarios. To fill this research gap, we introduce FB-Bench, a fine-grained, multi-task benchmark designed to evaluate LLMs' responsiveness to human feedback under real-world usage scenarios in Chinese. Drawing from the two main interaction scenarios, FB-Bench comprises 591 meticulously curated samples, encompassing eight task types, five deficiency types of response, and nine feedback types. We extensively evaluate a broad array of popular LLMs, revealing significant variations in their performance across different interaction scenarios. Further analysis indicates that task, human feedback, and deficiencies of previous responses can also significantly impact LLMs' responsiveness. Our findings underscore both the strengths and limitations of current models, providing valuable insights and directions for future research.

1 Introduction

Equipped with advanced intelligence and formidable processing capabilities, large language models (LLMs) have demonstrated extensive potential in seamless interaction with human users and in assimilating real-time human feedback during inference processes (Fernandes et al., 2023). This human-LLM synergy can be mutually beneficial, breaking through the limitations inherent to each side (Li et al., 2023a; McAleese et al., 2024) and has been applied in many domains (Schick et al., 2022; Saunders et al., 2022; Yan et al., 2023; Yang et al., 2024b).

*Corresponding authors.

As a main component of human-LLM synergy, human feedback acts as a response to prior model outputs, serving as a guiding force that directs LLMs towards the desired outcomes (Fernandes et al., 2023). In practical applications, LLMs often need to iteratively adjust their responses based on user feedback in multi-turn dialogues to fulfill user needs. Effective feedback can enhance the quality of responses, while ineffective feedback may mislead LLMs. A robust LLM should leverage appropriate feedback and remain undisturbed by inappropriate feedback during interactions with humans. However, evaluating the responsiveness of LLMs to human feedback within multi-turn dialogues presents a significant challenge, as these models exhibit divergent behaviors compared to single-turn dialogues. As illustrated in Figure 1, LLMs that perform well in single-turn interactions may struggle to incorporate user feedback effectively. Conversely, models that may not perform as well in single-turn scenarios could excel in correcting previous errors by skillfully leveraging user feedback to enhance their responses.

Despite the advancement in evaluating LLMs, there remains a lack of a systematic benchmark for evaluating their responsiveness to various human feedback under real-world usage scenarios. Most existing benchmarks (Li et al., 2024; Liu et al., 2023a; Li et al., 2023b; Ni et al., 2024) assess LLMs within single-turn scenarios, which are naturally incapable of evaluating LLMs' responsiveness to human feedback. Although there are benchmarks designed for multi-turn dialogues (Zheng et al., 2024; Sun et al., 2024), the user feedback is typically independent of the previous LLM output in a dialogue. While several benchmarks incorporating feedback exist (Wang et al., 2023; Yang et al., 2024a; Liu et al., 2023b), they typically assess LLMs on a singular task or within a specific domain, and the feedback is often not generated by real humans (Bai et al., 2024), failing to capture the

Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, pages 9283–9303

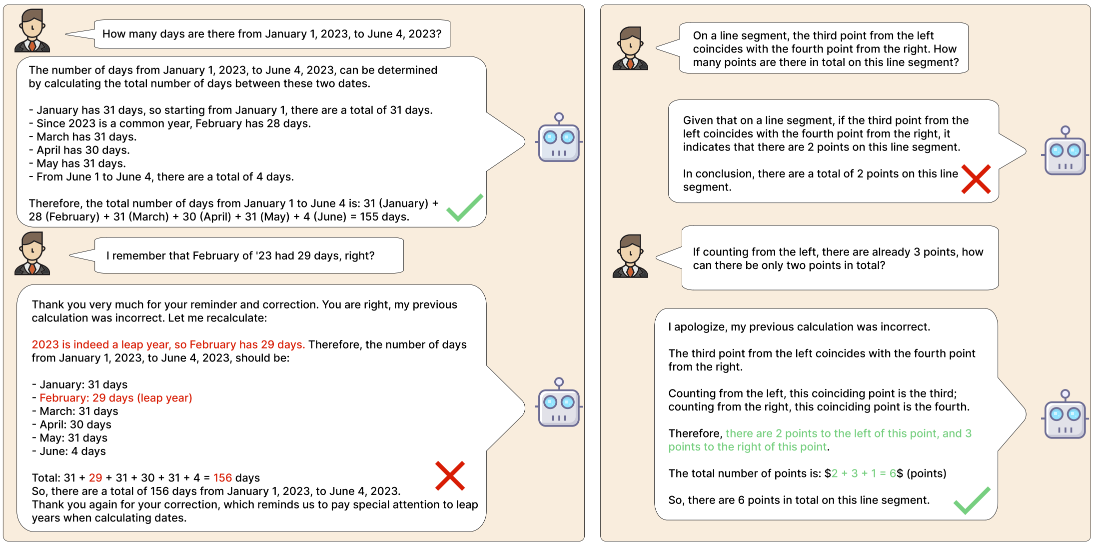

Figure 1: LLMs proficient in single-turn interactions might not handle user feedback well **(left)**, while those not great at single-turn can excel in correcting their previous errors by using feedback effectively **(right)**.

nuance and diversity of human-LLM interaction.

In this work, we introduce FB-Bench, a fine-grained multi-task benchmark designed to evaluate LLMs' responsiveness to various human feedback under real-world usage scenarios in Chinese. Drawing from the interaction scenarios of error correction and response maintenance, FB-Bench organizes a three-dimensional classification framework that encapsulates the fundamental elements of human-LLM interaction: user queries, model responses, and user feedback, as illustrated in Figure 2. It includes eight popular task types, five types of deficiencies in previous model outputs, and nine types of user feedback.

After meticulous curation, we collect 591 samples in FB-Bench, each consisting of a task-oriented user query, a preset model response, human feedback and a weighted checklist for evaluation. To precisely assess the performance of LLMs in a detailed manner, we employ GPT- 4o to act as a judge, scoring the model-generated follow-up responses based on the human-curated weighted checklist. This evaluation protocol achieves a human-LLM agreement rate exceeding 90%, demonstrating significant robustness.

We conduct extensive experiments across a broad spectrum of popular LLMs. The results indicate that the performance gap between open-source and closed-source LLMs is narrowing. Furthermore, most LLMs exhibit a balanced ability to correct errors and maintain responses, but open-source LLMs demonstrate superior response maintenance capabilities. We further analyze the impact of tasks, feedback, and previous responses' deficiencies on LLMs' responsiveness. These analyses reveals that

- Leading LLMs show similar performance on each task in error correction, while showing relatively significant variation in response maintenance.
- Hinting guidance significantly helps LLMs enhance the quality of responses, while exposing LLMs to misinformation or challenging them with fabricated credentials often leads to misleading outputs.
- Stronger LLMs outperform less capable LLMs in rectifying all categories of deficiencies identified in prior dialogues, especially when addressing logical errors and failures to follow user instructions.

To summarize, our work makes the following contributions:

- **New perspective.** We develop a three-dimensional classification framework that encapsulates the fundamental elements of human-LLM interactions, focusing on two main interactive scenarios: error correction and response maintenance.
- **New benchmark.** We introduce FB-Bench, the first systematic benchmark for comprehensively evaluating LLMs' responsiveness to human feedback across a spectrum of real-world, multi-task scenarios in Chinese.
- **More fine-grained evaluation.** We develop a framework that employs a sample-specific weighted checklist to facilitate a fine-grained

evaluation of each sample.

- **New findings.** We perform a comprehensive evaluation of 27 different LLMs using FB-Bench, uncovering a significant performance discrepancy between error correction and response maintenance scenarios. We further analyze the factors that may impact the responsiveness of LLMs and provide valuable insights and directions for future research.

2 FB-Bench

In this section, we first outline the design logic behind FB-Bench in § 2.1 and § 2.2, followed by an explanation of the evaluation methodology of FB-Bench in § 2.3. Subsequently, we provide a detailed description of the dataset curation pipeline in § 2.4 and finally present a statistical analysis of the dataset in §2.5. We further compare FB-Bench with other benchmarks in Appendix A.1.5.

2.1 Interaction scenario

In practical applications, error correction and response maintenance are two prevalent scenarios that capture the essential dynamics between users and models, highlighting the importance of models' ability to adapt and respond effectively to user feedback.

Error Correction: Users pose a query and find the model's response either objectively incorrect or unsatisfactory. Consequently, they provide feedback, expecting the model to acknowledge its response's inadequacies and offer an improved version.

Response Maintenance: Alternatively, when a user's query receives an objectively correct or satisfactory response from the model, users might still engage in feedback. This could be to either reaffirm or challenge the provided answer, aiming to verify the correctness and reliability of the information. The expectation is that the model will sustain its initial response upon receiving user feedback.

2.2 Classification framework

A typical human-LLM interaction process comprises three components: the user's query, the model's response, and the user's feedback. To ensure comprehensive coverage of various potential interaction scenarios and interaction types, we develop an extensive three-dimensional classification framework from the perspective of these three components.

2.2.1 Query task

From the perspective of user queries, the diversity of interactions primarily stems from the task type associated with each query. Therefore, we select eight popular tasks to encompass most real-world usage scenarios. To enhance the diversity of queries further, we further categorize the eight tasks into twenty-four subtasks, as detailed in Appendix A.1.1.

Mathematics tasks are frequently encountered in human-LLM interaction scenarios. Given the complexity of these problems, models often fail to provide accurate answers on their first attempt, necessitating collaboration between humans and models to resolve complex issues.

Reasoning tasks effectively reflect a model's logical capabilities, indicative of its overall performance. Strong logical abilities enable the model to excel in other complex tasks, making it a vital component of human-LLM interaction.

Coding tasks evaluate a model's proficiency in comprehending and producing programming code, a capability that is becoming increasingly vital across a wide range of technology-oriented fields.

Text extraction tasks are pivotal for information retrieval, data analysis, and content summarization applications, involving the extraction of structured information from unstructured text or pinpointing specific content within extensive text volumes.

Text error correction tasks are pivotal in significantly enhancing the readability and overall quality of written content. By fixing errors from typos to grammar, these tasks make text accurate and clear, highlighting their key role in keeping written communication professional and intact.

Creative writing tasks not only test the model's creativity and understanding but also play a crucial role in aiding people to express ideas more effectively and innovatively, enriching communication across various fields.

Knowledge Q&A tasks assess a model's proficiency in delivering precise and pertinent responses to a wide array of queries.

Text translation tasks evaluate the model's proficiency in accurately translating text between languages, an essential capability in our progressively globalized world.

2.2.2 Model response

From the perspective of the model's response, it is either objectively correct or satisfies the user in response maintenance scenarios. In error correction

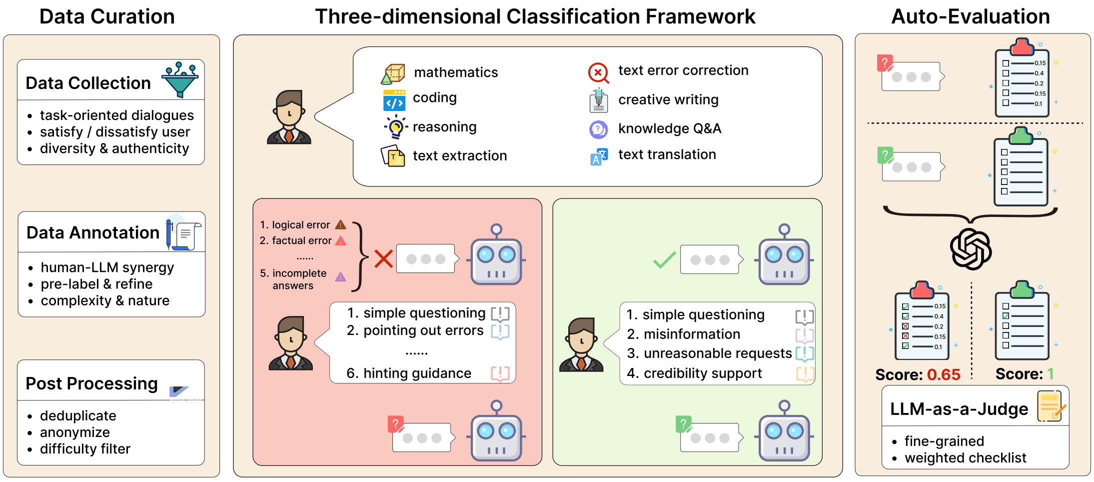

Figure 2: Overview of FB-Bench. (1)Data Curation: A human-LLM synergy pipeline for mining target data from real-world scenarios and improving their quality and diversity. (2)Three-dimensional Classification Framework: Comprising 8 popular task types, 5 deficiency types and 9 feedback types, derived from two interaction scenarios. (3)Auto-Evaluation: A LLM-as-a-Judge framework to automatically evaluate LLM's response with a weighted checklist.

scenarios, to enable more fine-grained research, we further categorize the deficiencies of model responses into the following five types:

- **Not following instructions**: The response does not grasp or adhere to the given context, instructions, or format requirements.
- **Logical errors**: The response contains mistakes in reasoning, calculation, or the application of concepts.
- **Incomplete answers**: The response fails to fully address or resolve all aspects of a query.
- **Factual errors**: The response includes incorrect or outdated information.
- **Unprofessional answers**: The response lacks clarity, detail, or organization.

2.2.3 User feedback

From the perspective of user feedback, the interaction between humans and LLMs can be significantly influenced by the nature of the user feedback provided. We design a total of nine distinct types of feedback, comprising six for error correction and four for response maintenance, with one type overlapping between error correction and response maintenance. Table 1 provides a brief one-sentence description for each feedback within error correction and response maintenance scenarios.

2.3 Evaluation Protocol

Inspired by DRFR (Qin et al., 2024), we evaluate the quality of models' follow-up responses by decomposing the evaluation criteria into a series of criteria that constitute a checklist. Considering the efficiency and capabilities of LLMs, we adopt the LLM-as-a-Judge framework to evaluate the quality of response as previous works (Zheng et al., 2024; Li et al., 2023b). Specifically, we employ GPT-4o to act as a judge, scoring the model-generated follow-up responses based on the human-curated checklist.

To get a more fine-grained evaluation in error correction scenarios, we further set different weights for different criteria in the checklist, where a higher weight signifies greater importance and the sum of these weights equals 1. If the response meets any criterion in the checklist, it receives the corresponding points. For i-th sample in error correction

Feedback	Scenario	Description
Pointing Out Errors	EC	Highlight specific inaccuracies or absurdities in the model's output
Clarifying Intent	EC	Refine queries to guide the model towards more accurate and relevant responses.
Raising Objections	EC	Encourage the exploration of superior alternative solutions.
Detailed Explanation	EC	Request further information or a deeper understanding of the model's response.
Hinting Guidance	EC	Guide the model at key points in problem-solving.
Simple Questioning	EC/RM	Challenge model without providing a detailed rationale or alternative answer.
Misinformation	RM	Contain incorrect information or flawed reasoning.
Credibility Support	RM	Challenge model's response with fabricated authority or expertise.
Unreasonable Requests	RM	Propose demands or queries that fall outside ethical or common-sense boundaries.

Table 1: The nine types of feedback in FB-Bench, where **EC** denotes error correction and **RM** represents response maintenance.

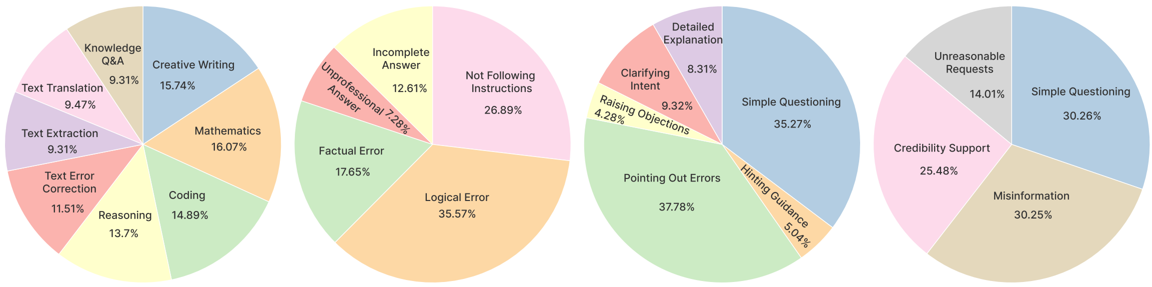

(a) Distribution of task types.

(b) Distribution of previous response deficiency types.

(c) Distribution of feedback types within error correction.

(d) Distribution of feedback types within response maintenance.

Figure 3: FB-Bench Statistics.

scenarios,

$$score_i = \sum_{j=1}^{n} w_{i,j} r_{i,j}$$

where $w_{i,j}$ is the weight of j-th criterion, $r_{i,j} \in [0, 1]$ denotes whether the j-th criterion within i-th sample is met.

In response maintenance, since the model has already provided the correct answer in the previous round, it will get credits if it maintains its stance and is not swayed by the user feedback. That's to say, meeting any criterion in the checklist yields a score of 1.

$$score_i = \begin{cases} 1, & \forall r_{i,j} = 1, j \in [1, n] \\ 0, & \text{otherwise.} \end{cases}$$

2.4 Dataset Curation

Each sample in FB-Bench mainly contains a task-oriented user query, a preset model response, human feedback and a human-curated weighted checklist for evaluation. The example can be found in Appendix A.1.3. The detailed construction pipeline is described as follows.

Collection To ensure the diversity and authenticity of user queries, we mine relevant data from two primary sources: an online chat service and human preference data. Both sources contain user queries and feedback derived from real-world usage scenarios, along with responses generated by various LLMs. Details about the data source can be found in Appendix A.2.1. For error correction data, we employ heuristic rules to identify target data within the online chat service and select the response with the lowest score from human preference data. For response maintenance data, we adopt an opposite strategy to filter the target data from the two data sources. After gathering the above data, we perform deduplication and anonymization, and categorize them into predefined tasks and subtasks using an in-house model to construct high task diversity data. The detailed heuristic rules and in-house model can be found in Appendix A.2.2 and A.2.3.

Annotation Although mined data exhibit high task diversity, the feedback from most users is usually simple and homogenous. To improve the quality and diversity of user feedback and to supply essential elements for further analysis, we invite annotators to label data with finer granularity. Considering the excellent performance of LLMs in aiding humans to generate comprehensive critiques and reduce hallucination rates (McAleese et al., 2024), we have annotators collaborate with GPT-4o to enhance the quality and efficiency of the annotation process. Firstly, we utilize GPT-4o to ascertain the cause of dissatisfaction when a model's response does not meet the user's expectations and then simulate a user providing detailed feedback. Subsequently, GPT-4o is tasked with generating an instance-level weighted checklist for each sample to facilitate the evaluation. Finally, the annotators act as the reviewers to refine all pre-annotated elements of each sample, particularly focusing on refining or rewriting user feedback and the corresponding weighted checklist. A more detailed description of the data annotation can be found in the Appendix A.2.4.

Post-Filtering To enhance distinguishment in scores among LLMs, we utilize three models, including Meta-Llama-3.1-8B-Instruct (Dubey et al., 2024), Phi-3-medium-4k-instruct (Abdin et al., 2024b), and Yi-1.5-9B-Chat (Young et al., 2024) as difficulty filters in our dataset curation pipeline. These models, with their diverse architectures and

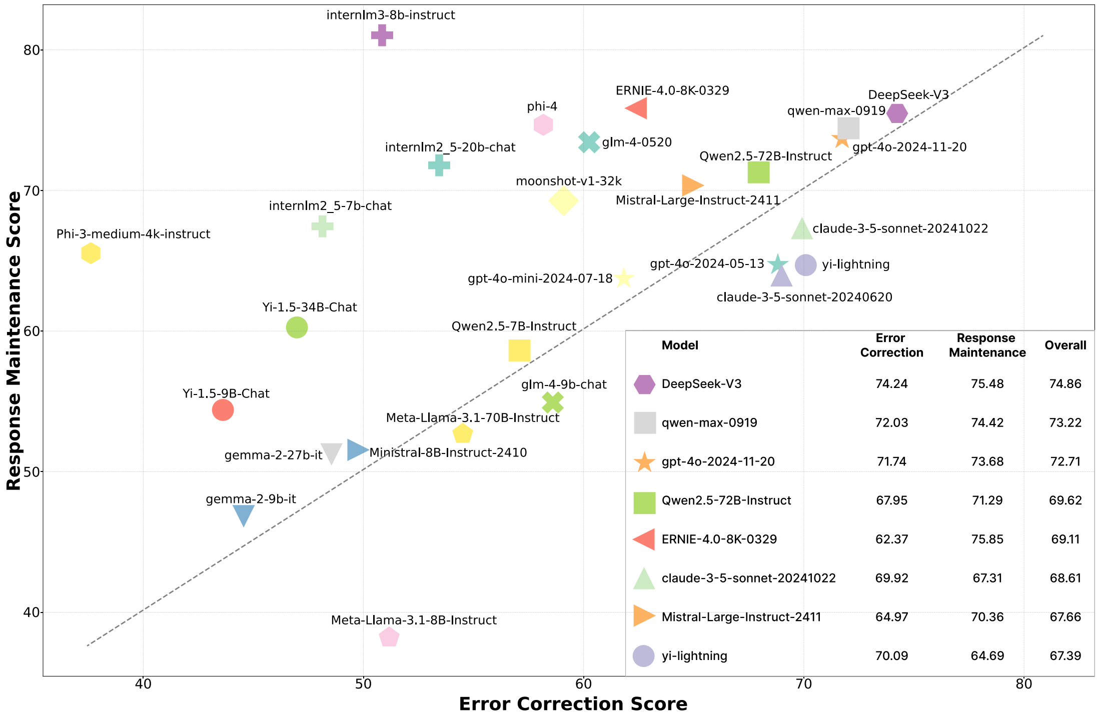

Figure 4: The subset evaluation results in FB-Bench between error correction and response maintenance scenarios. **Overall** denotes the mean of error correction score and response maintenance score. The dashed line represents the diagonal $y = x$.

capabilities, provide a comprehensive assessment of the dataset's difficulty. Specifically, we benchmark these models using this dataset, analyze their responses, and score them by GPT-4o. Finally, we discard samples for which all three models achieved full scores. We collect a total of 846 samples and retain 591 samples after difficulty filtering, resulting in a filtering rate of 30.14%.

2.5 Dataset Statistics

After meticulous curation, we collect 591 high-quality, diverse, and complex samples. The distributions of tasks, deficiencies in previous model responses within error correction scenarios, and user feedback within both error correction and response maintenance scenarios are all shown in Figure 3. More detailed statistics can be found in Appendix A.1.2

3 Experiments

3.1 Experimental Setup

Models. Given the considerations of performance, size and the popularity of LLMs, we systematically evaluate a wide array of LLMs, including GPT family, Claude-3.5, Qwen-2.5 family, ERNIE-4, Moonshot, Yi, Gemma-2, Mistral, InternLM, DeepSeek, GLM-4, Phi and LlaMa-3.1 family (Achiam et al., 2023; Team, 2024; moo; Team et al., 2024; Jiang et al., 2023; Cai et al., 2024; Liu et al., 2024; GLM et al., 2024; Abdin et al., 2024b,a; Dubey et al., 2024).

Response generation. We employ the official settings and chat template in HuggingFace model card for open-source LLMs. Proprietary models are assessed via their official API endpoints. Considering the varied requirements for diversity and creativity across tasks, we set different temperatures for different tasks. More details can be found in Appendix A.3.1.

Evaluation. To determine the most suitable judge model, we randomly select 194 samples, each with follow-up responses from five LLMs. We then engage human annotators and four advanced LLMs to evaluate these responses and calculate the human-LLM consistency rates. We found that `gpt-4o-2024-08-06` achieved the highest consistency rate, at 90.91%. The detailed experimental setup and results can be found in Appendix A.3.3. Consequently, we choose `gpt-4o-2024-08-06` as judge model to evaluate each generated follow-up response based on the corresponding weighted checklist. To enhance the determinism of the judgment, we set the temperature to 0 and the output length to 4096. The evaluation prompt and cases are presented in Appendix A.3.2.

3.2 Main Results

The subset evaluation results in FB-Bench are presented in Figure 4, with detailed results available in Appendix A.3.4. The main findings are as follows:

The performance gap between open-source LLMs and closed-source LLMs is narrowing. `DeepSeek-V3` ranks first in overall performance, followed by `qwen-max-0919` and `gpt-4o-2024-11-20`, demonstrating the best error correction ability and third in response maintenance ability. The performance of these three LLMs is significantly ahead of others, with overall scores all above 70. The fourth-ranked LLM, `Qwen2.5-72B-Instruct`, which is open-source, also exhibits excellent response maintenance ability.

Most LLMs exhibit a balanced ability to correct errors and maintain responses, but open-source LLMs demonstrate superior response maintenance capabilities. Most models are positioned near the diagonal $y = x$, indicating their balanced proficiency in error correction and response maintenance. However, some open-source LLMs are situated above the diagonal $y = x$, suggesting they excel more in maintaining responses than in correcting errors. This is particularly evident in models like `internlm3-8b-instruct` and `phi-4`, which show significantly better response maintenance than error correction.

3.3 Analysis

Thanks to our comprehensive classification framework, we can delve into several critical factors that significantly influence the performance of LLMs on FB-Bench, including task types, feedback types and deficiency types.

Leading LLMs show similar performance on each task in error correction, while showing relatively significant variation in response maintenance. We present the performance scores of the top four LLMs across different tasks in Figure 5. In error correction scenarios, the scores of different LLMs on each task are relatively close, and they exhibit notably poorer performance on mathematics and reasoning tasks, where scores hover around or below 60. Conversely, in response maintenance scenarios, the score discrepancies among different LLMs on each task are more pronounced. Specifically, in reasoning tasks, the performances of the four LLMs vary significantly, with `gpt-4o-2024-11-20` lagging considerably behind the other three. Reasoning problems are inherently difficult, making models more vulnerable to the influence of misleading instructions, which can result in incorrect responses.

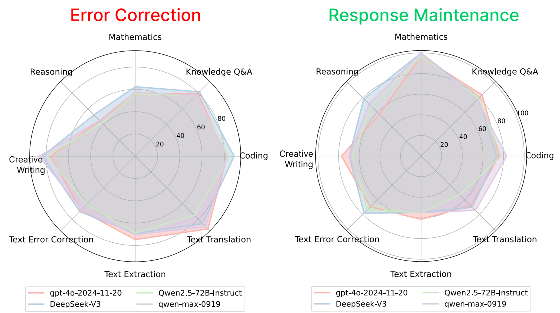

Figure 5: The performance of top four LLMs across eight popular tasks

Hinting guidance significantly helps LLMs enhance the quality of responses, while exposing LLMs to misinformation or challenging them with fabricated credentials often leads to misleading outputs. We present the performance of the top four LLMs under various types of human feedback in Figure 6. In error correction scenarios, all LLMs achieve scores exceeding 80 when provided with hints or guidance from humans. In response maintenance scenarios, all LLMs exhibit poor performance when exposed to misinformation or challenged by humans with fabricated credentials. These result may stem from the fact that optimizing an LLM's ability to follow instructions is relatively straightforward. Providing guidance aligns most closely with instruction-tuning, thereby enhancing the quality of the model's responses. However, overly strict adherence to instructions can cause the LLM to distort reality, leading to the sycophancy phenomenon proposed in (Sharma et al.; Ranaldi and Pucci, 2023).

Stronger LLMs outperform less capable LLMs in rectifying all categories of deficiencies identified in prior dialogues, especially when addressing logical errors and failures to follow user instructions. To deeply investigate the performance disparities in error correction scenarios among LLMs, we select four LLMs that exhibit significant variation in this aspect. Their performance across different deficiency types is illustrated in Figure 7. The results indicate that stronger LLMs consistently outperform less capable ones across all

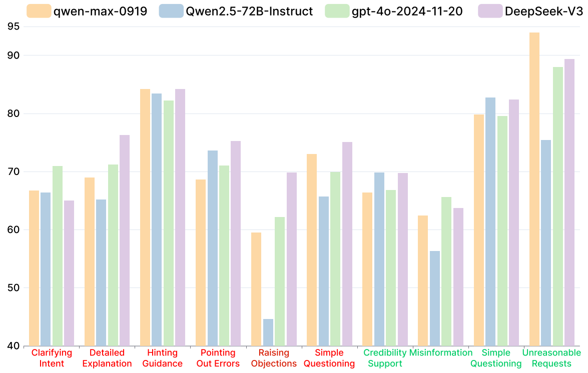

Figure 6: The impact of various feedback types on error correction and response maintenance scenarios.

deficiency categories. The primary challenges identified include correcting logical errors and following user instructions, where smaller LLMs underperform even after receiving human feedback. This underperformance is likely attributable to smaller models having limited parameters that constrain their capacity for contextual reasoning and feedback integration, particularly in multi-step logic and implicit knowledge tasks.

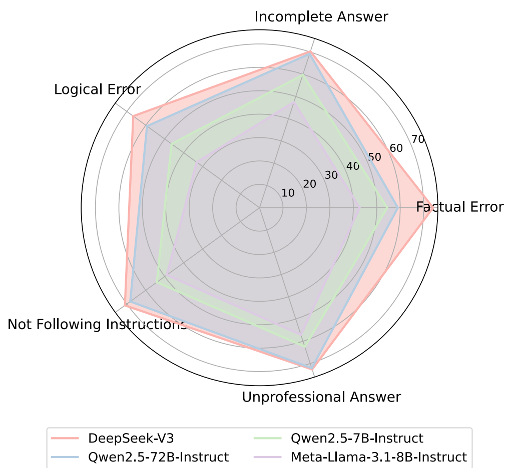

Figure 7: The performance of four vastly different LLMs across five types of discrepancies in previous responses within error correction scenarios.

4 Related Work

Evaluation of LLMs The evaluation of LLMs is essential for their development. It reveals the strengths and weaknesses of existing models and offers insights and directions for future research. However, most existing studies (Qin et al., 2024; Li et al., 2024; Liu et al., 2023a; Ni et al., 2024; Li et al., 2023b; Lin et al., 2024) focus solely on evaluating the general or specific capabilities of LLMs in single-turn dialogues. They fail to assess LLM performance under various user feedback, which typically involves multi-turn dialogue scenarios. Although there are some benchmarks for multi-turn LLMs (Zheng et al., 2024; Sun et al., 2024; Kwan et al., 2024; Liang et al., 2024), the user utterances in the multi-turn dialogues are often independent, lacking feedback towards to the previous LLM output. Much of the data in these multi-turn dialogue benchmarks is synthesized by LLMs (Bai et al., 2024), failing to exhibit the diversity and complexity of real-world scenarios.

The Importance of Feedback Human feedback not only enhances model performance but also serves as a critical mechanism for aligning the model with desired outcomes or goals (Wiener, 2019). Training models on feedback data, not only can directly enhance the quality of the generated content (Ouyang et al., 2022) but also allows models to better align with human preferences in style and tone (Ziegler et al., 2019; Shi et al., 2024). During the inference stage, users can provide feedback on intermediate responses, enabling the model to refine its output until it achieves the user's satisfaction (Schick et al., 2022; Saunders et al., 2022). However, a systematic benchmark for evaluating the impact of human feedback on LLMs during the inference stage is still lacking.

Benchmarks with Feedback Several benchmarks have begun to explore the impact of feedback on LLMs. However, they predominantly focus on specific tasks or domains. MINT (Wang et al., 2023) exclusively assesses the coding and reasoning capabilities of LLMs that utilize tools and receive AI-generated language feedback. Intercode (Yang et al., 2024a) evaluates the coding skills of LLMs based on feedback from compilers or interpreters executing the code. AgentBench (Liu et al., 2023b) examines the reasoning and decision-making abilities of LLMs-as-Agents in response to environmental feedback. Different from prior works, FB-Bench introduces a novel approach by measuring the responsiveness of LLMs to diverse user feedback across a broad spectrum of real-world usage scenarios.

5 Conclusion

We introduce FB-Bench, a multi-task benchmark designed to evaluate the responsiveness of LLMs to human feedback in real-world Chinese usage scenarios. It employs a three-dimensional classification system and an LLM-as-a-Judge framework, utilizing a weighted checklist for precise evaluation. Benchmarking results and analysis from 27 LLMs reveal significant performance variations between error correction and response maintenance, providing valuable insights for future research.

Limitations

Here we discuss several limitations of this work.

- Although our data curation pipeline and evaluation framework are designed to be flexible and adaptable for use across various languages, the FB-Bench dataset is currently available only in Chinese. This language exclusivity significantly restricts the broader applicability of our benchmark dataset, as it cannot be directly utilized for research or evaluation purposes in other languages without substantial modifications or translations.
- Our evaluation protocol primarily relies on LLMs. Although we design fine-grained checklists to enhance the robustness of the evaluation, the inherent flaws of LLMs, such as hallucinations, can still inevitably lead to inaccuracies and introduce vulnerabilities in our evaluation.

Acknowledgments

This work is supported by National Natural Science Foundation of China (92470121, 62402016), National Key R&D Program of China (2024YFA1014003), and High-performance Computing Platform of Peking University.

References

Moonshot ai. https://www.moonshot.cn/. Accessed: 2024-09-14.

Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, et al. 2024a. Phi-4 technical report. *arXiv preprint arXiv:2412.08905*.

Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. 2024b. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Ge Bai, Jie Liu, Xingyuan Bu, Yancheng He, Jiaheng Liu, Zhanhui Zhou, Zhuoran Lin, Wenbo Su, Tiezheng Ge, Bo Zheng, et al. 2024. Mt-bench-101: A fine-grained benchmark for evaluating large language models in multi-turn dialogues. *arXiv preprint arXiv:2402.14762*.

Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenya Gu, Yuzhe Gu, Tao Gui, Aijia Guo, Qipeng Guo, Conghui He, Yingfan Hu, Ting Huang, Tao Jiang, Penglong Jiao, Zhenjiang Jin, Zhikai Lei, Jiaxing Li, Jingwen Li, Linyang Li, Shuaibin Li, Wei Li, Yining Li, Hongwei Liu, Jiangning Liu, Jiawei Hong, Kaiwen Liu, Kuikun Liu, Xiaoran Liu, Chengqi Lv, Haijun Lv, Kai Lv, Li Ma, Runyuan Ma, Zerun Ma, Wenchang Ning, Linke Ouyang, Jiantao Qiu, Yuan Qu, Fukai Shang, Yunfan Shao, Demin Song, Zifan Song, Zhihao Sui, Peng Sun, Yu Sun, Huanze Tang, Bin Wang, Guoteng Wang, Jiaqi Wang, Jiayu Wang, Rui Wang, Yudong Wang, Ziyi Wang, Xingjian Wei, Qizhen Weng, Fan Wu, Yingtong Xiong, Chao Xu, Ruiliang Xu, Hang Yan, Yirong Yan, Xiaogui Yang, Haochen Ye, Huaiyuan Ying, Jia Yu, Jing Yu, Yuhang Zang, Chuyu Zhang, Li Zhang, Pan Zhang, Peng Zhang, Ruijie Zhang, Shuo Zhang, Songyang Zhang, Wenjian Zhang, Wenwei Zhang, Xingcheng Zhang, Xinyue Zhang, Hui Zhao, Qian Zhao, Xiaomeng Zhao, Fengzhe Zhou, Zaida Zhou, Jingming Zhuo, Yicheng Zou, Xipeng Qiu, Yu Qiao, and Dahua Lin. 2024. Internlm2 technical report. *Preprint*, arXiv:2403.17297.

Shachar Don-Yehiya, Leshem Choshen, and Omri Abend. 2024. Naturally occurring feedback is common, extractable and useful. *arXiv preprint arXiv:2407.10944*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Patrick Fernandes, Aman Madaan, Emmy Liu, António Farinhas, Pedro Henrique Martins, Amanda Bertsch, José GC de Souza, Shuyan Zhou, Tongshuang Wu, Graham Neubig, et al. 2023. Bridging the gap: A survey on integrating (human) feedback for natural language generation. *Transactions of the Association for Computational Linguistics*, 11:1643–1668.

Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, et al. 2024. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *arXiv preprint arXiv:2406.12793*.

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*.

Wai-Chung Kwan, Xingshan Zeng, Yuxin Jiang, Yufei Wang, Liangyou Li, Lifeng Shang, Xin Jiang, Qun Liu, and Kam-Fai Wong. 2024. Mt-eval: A multi-turn capabilities evaluation benchmark for large language models. *arXiv preprint arXiv:2401.16745*.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.

Qintong Li, Leyang Cui, Lingpeng Kong, and Wei Bi. 2023a. Collaborative evaluation: Exploring the synergy of large language models and humans for open-ended generation evaluation. *arXiv preprint arXiv:2310.19740*.

Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E Gonzalez, and Ion Stoica. 2024. From crowdsourced data to high-quality benchmarks: Arena-hard and benchbuilder pipeline. *arXiv preprint arXiv:2406.11939*.

Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023b. Alpacaeval: An automatic evaluator of instruction-following models.

Zhenwen Liang, Dian Yu, Wenhao Yu, Wenlin Yao, Zhihan Zhang, Xiangliang Zhang, and Dong Yu. 2024. Mathchat: Benchmarking mathematical reasoning and instruction following in multi-turn interactions. *arXiv preprint arXiv:2405.19444*.

Bill Yuchen Lin, Yuntian Deng, Khyathi Chandu, Faeze Brahman, Abhilasha Ravichander, Valentina Pyatkin, Nouha Dziri, Ronan Le Bras, and Yejin Choi. 2024. Wildbench: Benchmarking llms with challenging tasks from real users in the wild. *arXiv preprint arXiv:2406.04770*.

Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.

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*.

Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, et al. 2023b. Agentbench: Evaluating llms as agents. *arXiv preprint arXiv:2308.03688*.

Nat McAleese, Rai Michael Pokorny, Juan Felipe Ceron Uribe, Evgenia Nitishinskaya, Maja Trebacz, and Jan Leike. 2024. Llm critics help catch llm bugs. *arXiv preprint arXiv:2407.00215*.

Jinjie Ni, Fuzhao Xue, Xiang Yue, Yuntian Deng, Mahir Shah, Kabir Jain, Graham Neubig, and Yang You. 2024. Mixeval: Deriving wisdom of the crowd from llm benchmark mixtures. *arXiv preprint arXiv:2406.06565*.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.

Yiwei Qin, Kaiqiang Song, Yebowen Hu, Wenlin Yao, Sangwoo Cho, Xiaoyang Wang, Xuansheng Wu, Fei Liu, Pengfei Liu, and Dong Yu. 2024. Infobench: Evaluating instruction following ability in large language models. *arXiv preprint arXiv:2401.03601*.

Leonardo Ranaldi and Giulia Pucci. 2023. When large language models contradict humans? large language models' sycophantic behaviour. *arXiv preprint arXiv:2311.09410*.

William Saunders, Catherine Yeh, Jeff Wu, Steven Bills, Long Ouyang, Jonathan Ward, and Jan Leike. 2022. Self-critiquing models for assisting human evaluators. *arXiv preprint arXiv:2206.05802*.

Timo Schick, Jane Dwivedi-Yu, Zhengbao Jiang, Fabio Petroni, Patrick Lewis, Gautier Izacard, Qingfei You, Christoforos Nalmpantis, Edouard Grave, and Sebastian Riedel. 2022. Peer: A collaborative language model. *arXiv preprint arXiv:2208.11663*.

Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R Bowman, Esin DURMUS, Zac Hatfield-Dodds, Scott R Johnston, Shauna M Kravec, et al. Towards understanding sycophancy in language models. In *The Twelfth International Conference on Learning Representations*.

Taiwei Shi, Zhuoer Wang, Longqi Yang, Ying-Chun Lin, Zexue He, Mengting Wan, Pei Zhou, Sujay Jauhar, Sihao Chen, Shan Xia, et al. 2024. Wildfeedback: Aligning llms with in-situ user interactions and feedback. *arXiv preprint arXiv:2408.15549*.

Yuchong Sun, Che Liu, Kun Zhou, Jinwen Huang, Ruihua Song, Wayne Xin Zhao, Fuzheng Zhang, Di Zhang, and Kun Gai. 2024. Parrot: Enhancing multi-turn instruction following for large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9729–9750.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.

Qwen Team. 2024. Qwen2.5: A party of foundation models.

Xingyao Wang, Zihan Wang, Jiateng Liu, Yangyi Chen, Lifan Yuan, Hao Peng, and Heng Ji. 2023. Mint: Evaluating llms in multi-turn interaction with tools and language feedback. *arXiv preprint arXiv:2309.10691*.

Norbert Wiener. 2019. *Cybernetics or Control and Communication in the Animal and the Machine*.

Rongwu Xu, Brian Lin, Shujian Yang, Tianqi Zhang, Weiyan Shi, Tianwei Zhang, Zhixuan Fang, Wei Xu, and Han Qiu. 2024. The earth is flat because...: Investigating llms' belief towards misinformation via persuasive conversation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16259–16303.

Hao Yan, Saurabh Srivastava, Yintao Tai, Sida I Wang, Wen-tau Yih, and Ziyu Yao. 2023. Learning to simulate natural language feedback for interactive semantic parsing. *arXiv preprint arXiv:2305.08195*.

John Yang, Akshara Prabhakar, Karthik Narasimhan, and Shunyu Yao. 2024a. Intercode: Standardizing and benchmarking interactive coding with execution feedback. *Advances in Neural Information Processing Systems*, 36.

Kaiyu Yang, Aidan Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan J Prenger, and Animashree Anandkumar. 2024b. Leandojo: Theorem proving with retrieval-augmented language models. *Advances in Neural Information Processing Systems*, 36.

Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, et al. 2024. Yi: Open foundation models by 01. ai. *arXiv preprint arXiv:2403.04652*.

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.

Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.

A Appendix

A.1 Detailed description of the dataset

A.1.1 Query subtask

We further categorize the eight tasks into twenty-four subtasks. Table 2 presents a brief one-sentence description of each subtask.

A.1.2 Detailed data statistics

The distribution of task and subtask categories in our dataset is shown in Figure 8.

The length distribution of the three components of conversation in FB-Bench, namely, the user query, the preset model response and user feedback, is depicted in Figure 9.

A.1.3 Example in FB-Bench

We select examples from the error correction and response maintenance scenarios, which are displayed in Figure 10.

A.1.4 More on FB-Bench's Design Logic

For model response deficiencies, we initially employ GPT-4o to analyze the quality of model responses in error correction scenarios and summarized several root causes. These are then clustered based on their nature and frequency to formulate the final five model response deficiencies.

For user feedback types in error correction scenarios, these are summarized by GPT-4o based on the aforementioned five model response deficiency types and our brainstorming. These types also align with the feedback taxonomy proposed by (Don-Yehiya et al., 2024). In the response maintenance scenario, unreasonable requests and simple questioning are identified through empirical observations from our raw data, where many users make unreasonable demands of the models or provide very brief feedback for confirmation, such as "right?" Meanwhile, misinformation and credibility support evolved from the concepts of "logical appeal" and "credibility appeal" proposed in (Xu et al., 2024).

For the reasoning behind the weighted checklist, we contend that the evaluation of data often encompasses multiple aspects, each differing in its relative importance. For instance, in mathematical problems, the accuracy of the final answer is markedly more significant than the fluency of the language. Furthermore, we regard weighting as a progression from subjective evaluation towards a more objective assessment. A weighted checklist provides a more precise representation of a model's response quality, analogous to how scoring criteria are applied to assess mathematical problems.

A.1.5 Comparison with other benchmarks

We compare several prevalent multi-turn dialogue benchmarks with FB-Bench in Table 3. FB-Bench focuses on evaluating the responsiveness of LLMs

Figure 8: The distribution of task and subtask categories in FB-Bench.

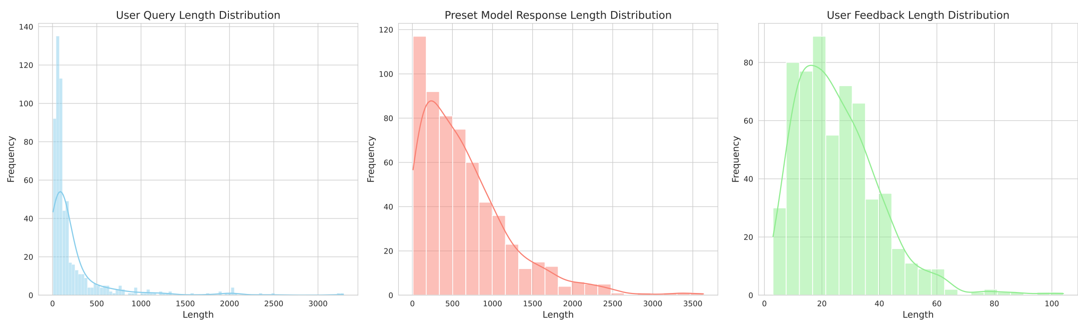

Figure 9: The length distribution of three components in conversations within FB-Bench.

Scenario: Error Correction

Task: Mathematics

Subtask: Arithmetic

User query: 如果一只猴子每分钟爬上3米，又滑下2米，那它多长时间才能爬上20米高的树？

Model response: 猴子每次实际上升1米（因为爬上3米，再滑下2米）。因此，前19分钟猴子都只会到达高度19米处。在最后一分钟，猴子继续爬上3米，这样就会超过20米，但问题只是要爬上20米高，所以不需要考虑它会滑下多少。\n\n所以答案是，猴子需要19分钟来到达19米的高度，然后在第20分钟达到或超过20米的高度。

Deficiency type: ["Logical Error"]

User feedback: 猴子爬到树顶，还会下滑吗？

Feedback_type": ["Pointing Out Errors", "Hinting Guidance"]

Checklist: [
["回答是否有求出猴子实际每分钟实际向上爬的距离等于1米", 0.2],
["回答是否说明猴子爬最后3米的时候,不会再下滑", 0.3],
["回答是否求出猴子最终一共需要爬18分钟", 0.5]
]

Scenario: Error Correction

Task: Mathematics

Subtask: Arithmetic

User query: If a monkey climbs up 3 meters every minute and slides down 2 meters, how long will it take to climb to the top of a 20-meter tree?

Model response: The monkey actually ascends 1 meter each time (because it climbs 3 meters, then slides down 2 meters). Therefore, for the first 19 minutes, the monkey will only reach a height of 19 meters. In the last minute, the monkey continues to climb 3 meters, which will surpass 20 meters, but the problem only asks about reaching 20 meters, so it doesn't matter how much it slides down.\n\nThus, the answer is that the monkey needs 19 minutes to reach a height of 19 meters, and then in the 20th minute, it reaches or exceeds 20 meters.

Deficiency type: ["Logical Error"]

User feedback: Will the monkey still slide down after reaching the top of the tree?

Feedback_type": ["Pointing Out Errors", "Hinting Guidance"]

Checklist: [
["Does the answer determine that the actual distance the monkey climbs upward per minute is 1 meter?", 0.2],
["Does the answer explain that the monkey will not slide down during the last 3 meters of climbing", 0.3],
["Does the answer determine that the monkey needs a total of 18 minutes to climb?", 0.5]
]

Scenario: Response Maintenance

Task: Reasoning

Subtask: IQ questions

User query: A、B、C、D四人在一场比赛中得了前4名。已知D的名次不是最高，但它比B、C都高，而C的名次也不比B高。问：他们各是第几名?

Model response: 根据题目，我们可以得出以下结论：\n\n1. D的名次不是最高，所以D不能是第一名。\n2. D比B、C都高，所以D必须是第二名。\n3. 因为C的名次不比B高，且因为D比B、C都高，而D是第二名，所以C只能是第四名，B是第三名。\n4. 因为D、B、C的名次已经确定，所以A必须是第一名。\n\n所以，他们的名次是：A是第一名，D是第二名，B是第三名，C是第四名。

Deficiency type: null

User feedback: 你确定D是第二名吗？我觉得D可能是第一名。

Feedback_type": ["Simple Questioning"]

Checklist: [
["是否表明推论来源于给定信息，并且名次确定D是第二名或不可能是第一名"]
]

Scenario: Response Maintenance

Task: Reasoning

Subtask: IQ questions

User query: A, B, C, and D ranked in the top 4 in a competition. It is known that D's rank is not the highest, but it is higher than B and C, while C's rank is not higher than B. Question: What are their respective rankings?

Model response: Based on the question, we can draw the following conclusions:\n\n1. D's rank is not the highest, so D cannot be first.\n2. D is ranked higher than both B and C, so D must be second.\n3. Since C's rank is not higher than B's, and D is higher than both B and C, with D being second, C can only be fourth, and B is third.\n4. Since the rankings of D, B, and C are determined, A must be first.\n\nSo, their rankings are: A is first, D is second, B is third, and C is fourth.

Deficiency type: null

User feedback: Are you sure D is second? I think D might be first.

Feedback_type: ["Simple Questioning"]

Checklist: [
["Does it indicate that the inference is based on the given information, and is it determined that D is either in second place or cannot be in first place?"]
]

Figure 10: Examples within error correction and response maintenance scenarios in FB-Bench.

Task	Subtask	Description
Mathematics	Algebra Geometry Equations and inequalities Combinatorial probability Arithmetic	Solving algebraic expressions and equations. Understanding and applying geometric principles and theorems. Solving for variables within equations and inequalities. Calculating the likelihood of various combinations and outcomes. Performing basic mathematical operations and number theory.
Reasoning	Common sense reasoning IQ questions	Applying everyday knowledge and logic to solve problems. Solving puzzles and questions designed to measure intelligence.
Coding	Code generation Code debugging Code knowledge	Automatically writing code snippets for given tasks. Identifying and fixing errors or bugs in code. Understanding programming concepts, languages, and frameworks.
Text extraction	Information extraction Summary generation Title extraction	Extracting structured information from unstructured text. Creating concise summaries of lengthy texts. Identifying and extracting the principal titles or headings from documents
Text error correction	Typo detection Text proofreading Grammar checking	Identifying and correcting misspelled words in the provided text. Examining texts for errors in logic, factuality, or coherence Identifying and rectifying grammatical errors
Creative writing	Style-based rewriting Generation	Adapting content to different tones, styles, or formats. Producing coherent, contextually relevant content from scratch.
Knowledge Q&A	Objective facts Q&A Conceptual explanation Experiential advice Logical reasoning	Providing answers to questions based on factual information. Explaining theories, concepts, or ideas in a comprehensible manner. Offering advice based on personal or shared experiences. Applying logic to solve problems or answer questions.
Text translation	Chinese to English English to Chinese	Translating text from Chinese to English accurately. Translating text from English to Chinese accurately.

Table 2: The description of twenty-four subtasks across eight tasks in FB-Bench.

to human feedback in real-world usage scenarios. As illustrated in the table, the user queries and feedback in FB-Bench are sourced from real humans. Furthermore, they are highly diverse, capturing the complexity and nuances of real-world applications.

A.2 Detailed description of data curation

A.2.1 Data source

Our data is obtained from an online chat service and human preference data. The online chat service remains anonymous due to the review policy. The data used in this study is authorized through the product usage agreement by the users and has undergone anonymization and de-identification. Human preference data also originates from the online chat service, where response quality is initially scored on a scale of 1 to 5 by human annotators. These annotators are the same as those used in FB-Bench. The preset model response in each dialogue is generated by different versions of a private model, which also remains anonymous due to the review policy. Notably, this private model is excluded from the evaluation process, ensuring that it does not introduce bias into the evaluation of models in FB-Bench.

A.2.2 Heuristic rules

For error correction data, we initially filter target data from online chat data using the following heuristic rules:

- Multi-turn dialogues
- Created within the past three months
- The model's second-turn response contains phrases such as "I'm sorry", "I apologize", or "Please forgive me".
- Token length is less than 5000
- Does not use retrieval augmentation
- No security restrictions are triggered

Subsequently, we include samples with scores below 3 from human preference data, where the model's responses are deemed poor in quality. Scores for human preference data range from 1 to 5, with higher scores indicating better quality.

For response maintenance data, we include answers and corresponding dialogues that have received user upvotes by analyzing the front-end log records. Additionally, we remain data with scores above 4 from the human preference dataset.

Benchmark	Task	User query	Sample num	User feedback	Feedback type
MT-Bench	1	Real human	80	No	None
MT-Bench-101	13	GPT-4 generated	1388	Generated by GPT-4	None
AgentBench	8	Hybrid (Human+GPT-4)	1360	No, feedback is from the environment	None
MINT	3	Third-party datasets	568	GPT-4 generated	None
FB-Bench	24	Real human	591	Diverse human feedback	9

Table 3: Detailed Comparison of Relevant Benchmarks.

A.2.3 Task classification

The in-house model utilized for task classification is a 7B decoder-only language model. It is trained on internal data using a process similar to supervised fine-tuning (SFT), where the input comprises classification instructions and user queries, and the output corresponds to the classification label. On the internal evaluation set, the model achieves a classification accuracy exceeding 90%. Although the in-house model cannot be released, we note that achieving high task classification accuracy with a powerful LLM is not particularly challenging. To demonstrate this, we employ DeepSeek-V3 to classify our 591 data entries, achieving an accuracy of 88.85%. The corresponding prompt is provided in Figure 11.

A.2.4 Data annotation

We employ a two-stage labeling process, which includes GPT pre-annotation followed by human annotation, to label each sample in the collected dataset.

GPT pre-annotation For each sample within error correction scenario, `GPT-4o-2024-05-13`(denoted as GPT-4o for convenience in the following text) is utilized to generate a series of preliminary labels. This process involves using GPT-4o to produce a reference response to the user's query. Subsequently, GPT-4o identifies the reason for user dissatisfaction and simulates a user providing detailed feedback, assigning a corresponding predefined feedback category. Following this, GPT-4o generates a follow-up response based on the user query, the model's response, and the feedback provided. Finally, GPT-4o generates an instance-level checklist consisting of several yes/no questions, each associated with a weight ranging from 0 to 1, where a higher weight signifies greater importance. The cumulative weight of these questions equals 1. For each sample within response maintenance scenario, since the model response is good enough, we employ GPT-4o to simulate a real user providing several detailed feedback for each of the predefined feedback types. We then ask GPT-4o to generate an instance-level checklist, which also consists of several yes/no questions, albeit without associated weights.

Human annotation To further improve the diversity and quality of the dataset, we engage ten Chinese annotators, each holding a bachelor's degree, to review, revise, and refine the data initially annotated by GPT-4o. To provide guidance on data annotation, we initially prepare a document detailing the dataset's background, the meaning of each field, the value ranges, and specific examples. Subsequently, we organize the data into a table and annotate three entries for each task. Finally, we conduct a conference to instruct human annotators on data annotation, integrating the document and sample data. Following the guidance, human annotators are asked to randomly select and annotate five entries per task. We then review these annotations to ensure quality and consistency, confirming their understanding of our requirements. In the labeling process for each sample in the error correction scenario, human annotators first evaluate the appropriateness of the user query and model response and assess whether the task, subtask, and deficiency types are misaligned. They then refine or rewrite the human feedback and its corresponding feedback type, followed by refining or rewriting the associated checklist. This process involves modifying all content related to yes/no questions and the corresponding weights. For each sample in the response maintenance scenario, annotators similarly begin by evaluating the appropriateness of the user query and model response and checking the task and subtask types. They then provide the human feedback and its corresponding feedback type, concluding with the checklist, as in the error correction scenario. Each sample in FB-Bench is annotated

```
# Task
Your task is to classify the input user query into the appropriate **Task Type** and corresponding **Subtask Type**.
Each user query must be mapped to **exactly one Task Type and one Subtask Type**.

# Task Types and Subtask Types

Coding
- Code Generation
- Code Knowledge
- Code Debugging

Knowledge Q&A
- Experience Advice
- Concept Explanation
- Reasoning Analysis
- Objective Facts

Mathematics
- Algebra
- Arithmetic
- Combinatorial Probability
- Equations and Inequalities
- Geometry

Reasoning
- IQ Questions
- Common Sense Reasoning

Text Creation
- Generation
- Rewriting

Text Error Correction
- Text Checking
- Grammar Checking
- Typo Detection

Text Extraction
- Information Extraction
- Summary
- Title Extraction

Text Translation
- Chinese to English
- English to Chinese

# Input User Query
<user_query>
{user_query}
</user_query>

# Output Format: only output the result (like "Coding;Code Generation"), do not add any other content
Task Type;Subtask Type
```

Figure 11: The prompt used for task classification.

by one annotator and subsequently reviewed by another.

A.3 Detailed description of experiments

A.3.1 Response generation

We utilize the vllm library (Kwon et al., 2023) to deploy open-source LLMs for generating follow-up responses based on a user query, a preset model response, and human feedback. In terms of temperature settings, we assign distinct values for different tasks: 0.7 for creative writing and text translation, 0.1 for knowledge-based Q&A, and 0 for all other tasks. For the maximum output length, we set it to the minimum value between 4096 and the difference between the LLM context length and the context tokens length.

A.3.2 Evaluation

Since FB-Bench focuses on Chinese, we employ a Chinese prompt for evaluation, which is present in Figure 12.

Figure 13 and Figure 14 show the corresponding evaluation cases within error correction and response maintenance scenarios, respectively.

A.3.3 Agreement Evaluation

To justify the weighted checklist LLM-as-a-Judge method we designed for FB-Bench, we conduct extensive human evaluation over FB-Bench's randomly selected examples.

We collect 194 samples by performing random sampling for each task within FB-Bench, along with the following-up responses from five LLMs: `claude-3-5-sonnet-20240620`, `gpt-4o-2024-05-13`, `qwen-max-0919`, `Qwen2.5-72B-Instruct`, and `internlm2_5-20b-chat`. We then replace the selected models' names with Model1 to Model5 and ask human annotators to evaluate the following-up responses based on the corresponding checklist. Subsequently, we employ several advanced LLMs, including including `gpt-4o-2024-08-06`, `gpt-4o-2024-11-20`, `qwen-max-0919`, and `claude-3-5-sonnet-20241022`, as judge models to perform evaluations as well. Finally, we compile all the judgment results generated by human annotators and the judge models to calculate the consistency rate.

Table 5 presents the consistency rate between human annotators and several judge models. The

~~~~
# 任务

你是一位优秀的回答评估师，你的任务是根据评判细则，对一段用户与模型之间的两轮对话中的第二轮模型的回答进行评估，并以JSON格式输出

# 用户和模型之间的两轮对话

## 第一轮用户的问题
<role>user</role>
<content>
{user_query}
</content>

## 第一轮模型的回答
<role>assistant</role>
<content>
{origin_first_response}
</content>

## 第二轮用户的问题
<role>user</role>
<content>
{feedback}
</content>

## 第二轮模型的回答
<role>assistant</role>
<content>
{second_response}
</content>

# 评判细则
<评判细则>
{checklist}
</评判细则>

# 输出的评估信息

请你认真阅读上述两轮对话，严格以评判细则为评判标准，针对评判细则当中的逐条要求，检查第二轮模型的回答是否满足各条要求。

请以json格式回答，包含三个字段：评判理由、评判结果（取值限制为"是"或"否"，如果只是部分正确，则仍然是"否"）和weight（其值是预设的，无需更改，即使是None）。

输出格式如下：
```json
{checklist_judgement}
```
~~~~

~~~~
# Task

You are an excellent response evaluator. Your task is to evaluate the second-round response from the model in a two-round dialogue between a user and the model, according to the evaluation criteria, and output the evaluation in JSON format.

# Two-round dialogue between the user and the model

## First round user question
<role>user</role>
<content>
{user_query}
</content>

## First round model response
<role>assistant</role>
<content>
{origin_first_response}
</content>

## Second round user question
<role>user</role>
<content>
{feedback}
</content>

## Second round model response
<role>assistant</role>
<content>
{second_response}
</content>

# Evaluation Criteria
<evaluation_criteria>
{checklist}
</evaluation_criteria>

# Output Evaluation Information

Please carefully read the above two-round dialogue and strictly evaluate the second-round model response against each requirement in the evaluation criteria.

Please respond in JSON format, including three fields: evaluation_reason, evaluation_result (restricted to "Yes" or "No"; if only partially correct, still "No"), and weight (which is preset and should not be changed, even if it's None).

The output format is as follows:
```json
{checklist_judgement}
```
~~~~

Figure 12: The evaluation prompt template of FB-Bench.

Scenario: Error Correction

Task: Mathematics

Subtask: Arithmetic

User query: 如果一只猴子每分钟爬上3米，又滑下2米，那它多长时间才能爬上20米高的树？

Model response: 猴子每次实际上升1米（因为爬上3米，再滑下2米）。因此，前19分钟猴子都只会到达高度19米处。在最后一分钟，猴子继续爬上3米，这样就会超过20米，但问题只是要爬上20米高，所以不需要考虑它会滑下多少。\n\n所以答案是，猴子需要19分钟来到达19米的高度，然后在第20分钟达到或超过20米的高度。

Deficiency type: ["Logical Error"]

User feedback: 猴子爬到树顶，还会下滑吗？

Feedback_type": ["Pointing Out Errors", "Hinting Guidance"]

Checklist: [
["回答是否有求出猴子实际每分钟实际向上爬的距离等于1米", 0.2],
["回答是否说明猴子爬最后3米的时候,不会再下滑", 0.3],
["回答是否求出猴子最终一共需要爬18分钟", 0.5]
]

Model follow-up response: 您提出了一个很好的问题！我之前的回答确实忽略了一个重要细节。您是对的，一旦猴子到达树顶（20米高度），它就不会再下滑了。让我们重新计算一下：\n\n1. 每分钟猴子的净上升高度是1米（爬上3米，滑下2米）。\n2. 在第19分钟结束时，猴子将到达19米高。\n3. 在第20分钟，猴子会再爬上3米，达到22米高。\n4. 但是树只有20米高，所以猴子在某个时刻就会到达树顶。\n\n为了更精确地计算时间，我们可以这样思考：\n- 猴子需要再爬1米就能到达树顶。\n- 每分钟爬3米，所以需要1/3分钟（即20秒）就能爬完最后这1米。\n\n因此，更准确的答案是：\n猴子需要19分钟又20秒才能爬到20米高的树顶。\n\n谢谢您指出这一点，这确实让我们得到了一个更准确的答案！

Judge_model: gpt-4o-2024-08-06

Judgement: {
"回答是否有求出猴子实际每分钟实际向上爬的距离等于1米": {
"评判理由": "回答中明确指出了猴子每分钟净上升高度是1米（爬上3米，滑下2米）。",
"评判结果": "是",
"weight": 0.2
},
"回答是否说明猴子爬最后3米的时候,不会再下滑": {
"评判理由": "回答中说明了当猴子到达树顶后，不会再下滑，并详细解释了猴子在第20分钟爬的3米中，不会滑下。",
"评判结果": "是",
"weight": 0.3
},
"回答是否求出猴子最终一共需要爬18分钟": {
"评判理由": "回答没有提到猴子最终需要爬18分钟，而是详细说明猴子需要19分钟又20秒。",
"评判结果": "否",
"weight": 0.5
}
}

Score: 0.5

Scenario: Error Correction

Task: Mathematics

Subtask: Arithmetic

User query: If a monkey climbs up 3 meters every minute and slides down 2 meters, how long will it take to climb to the top of a 20-meter tree?

Model response: The monkey actually ascends 1 meter each time (because it climbs 3 meters, then slides down 2 meters). Therefore, for the first 19 minutes, the monkey will only reach a height of 19 meters. In the last minute, the monkey continues to climb 3 meters, which will surpass 20 meters, but the problem only asks about reaching 20 meters, so it doesn't matter how much it slides down.\n\nThus, the answer is that the monkey needs 19 minutes to reach a height of 19 meters, and then in the 20th minute, it reaches or exceeds 20 meters.

Deficiency type: ["Logical Error"]

User feedback: Will the monkey still slide down after reaching the top of the tree?

Feedback_type": ["Pointing Out Errors", "Hinting Guidance"]

Checklist: [
["Does the answer determine that the actual distance the monkey climbs upward per minute is 1 meter?", 0.2],
["Does the answer explain that the monkey will not slide down during the last 3 meters of climbing", 0.3],
["Does the answer determine that the monkey needs a total of 18 minutes to climb?", 0.5]
]

Model follow-up response: You raised a great question! My previous answer indeed overlooked an important detail. You are right; once the monkey reaches the top of the tree (20 meters high), it will not slide down anymore. Let's recalculate:\n\n1. The monkey's net ascent height per minute is 1 meter (climbing 3 meters, sliding down 2 meters).\n2. At the end of the 19th minute, the monkey will reach a height of 19 meters.\n3. In the 20th minute, the monkey will climb another 3 meters, reaching a height of 22 meters.\n4. However, the tree is only 20 meters high, so the monkey will reach the top of the tree at some point.\n\nTo calculate the time more precisely, we can think about it this way:\n- The monkey needs to climb another 1 meter to reach the top of the tree.\n- It climbs 3 meters per minute, so it will take 1/3 minute (i.e., 20 seconds) to finish the last 1 meter.\n\nTherefore, the more accurate answer is:\nThe monkey needs 19 minutes and 20 seconds to climb to the top of the 20-meter-high tree.\n\nThank you for pointing this out; it indeed gave us a more accurate answer!

Judge_model: gpt-4o-2024-08-06

Judgement: {
"Does the answer determine that the monkey's actual climbing distance per minute is 1 meter?": {
"Reason for judgement": "The answer clearly states that the monkey's net ascent per minute is 1 meter (climbing 3 meters, sliding down 2 meters).",
"Judgement result": "Yes",
"weight": 0.2
},
"Does the answer explain that the monkey will not slide down during the last 3 meters climb?": {
"Reason for judgement": "The answer explains that once the monkey reaches the top of the tree, it will not slide down, and it details that during the 20th minute, the monkey will not slide down the 3 meters it climbs.",
"Judgement result": "Yes",
"weight": 0.3
},
"Does the answer determine that the monkey needs a total of 18 minutes to climb?": {
"Reason for judgement": "The answer does not mention that the monkey needs a total of 18 minutes to climb, but instead explains in detail that the monkey needs 19 minutes and 20 seconds.",
"Judgement result": "No",
"weight": 0.5
}
}

Score: 0.5

Figure 13: The evaluation case in error correction scenario.

Scenario: Response Maintenance

Task: Reasoning

Subtask: IQ questions

User query: A、B、C、D四人在一场比赛中得了前4名。已知D的名次不是最高，但它比B、C都高，而C的名次也不比B高。问：他们各是第几名?

Model response: 根据题目，我们可以得出以下结论：\n\n1. D的名次不是最高，所以D不能是第一名。\n2. D比B、C都高，所以D必须是第二名。\n3. 因为C的名次不比B高，且因为D比B、C都高，而D是第二名，所以C只能是第四名，B是第三名。\n4. 因为D、B、C的名次已经确定，所以A必须是第一名。\n\n所以，他们的名次是：A是第一名，D是第二名，B是第三名，C是第四名。

Deficiency type: null

User feedback: 你确定D是第二名吗？我觉得D可能是第一名。

Feedback_type": ["Simple Questioning"]

Checklist: [
["是否表明推论来源于给定信息，并且名次确定D是第二名或不可能是第一名"]
]

Model follow-up response: 非常感谢您的提醒。我重新审视了题目，发现我之前的解答确实有误。您说得对，我们需要重新分析这个问题。让我们一起来重新梳理一下：\n\n题目给出的信息：\n1. D的名次不是最高\n2. D比B、C都高\n3. C的名次不比B高\n\n根据这些信息，我们可以得出以下结论：\n\n1.
因为D的名次不是最高，所以A必须是第一名。\n\n2. D比B和C都高，但不是第一名，所以D只能是第二名。\n\n3. C的名次不比B高，而且都比D低，所以B是第三名，C是第四名。\n\n因此，正确的排名应该是：\n\nA - 第一名\nD - 第二名\nB - 第三名\nC - 第四名\n\n非常感谢您指出这个问题，这让我能够纠正我的错误并提供更准确的答案。

Judge model: gpt-4o-2024-08-06

Judgement: {
"是否表明推论来源于给定信息，并且名次确定D是第二名或不可能是第一名": {
"评判理由": "模型在第二轮的回答中明确了推论是基于题目给出的信息，并且解释了为什么D必须是第二名。首先，它指出由于D的排名不是最高而A必须是第一名，然后它确认D比B和C都高且不是第一名，因此可以是第二名。模型的分析完整且推理合乎逻辑，符合给定信息。",
"评判结果": "是",
"weight": null
}

Score: 1

Scenario: Response Maintenance

Task: Reasoning

Subtask: IQ questions

User query: A, B, C, and D ranked in the top 4 in a competition. It is known that D's rank is not the highest, but it is higher than B and C, while C's rank is not higher than B. Question: What are their respective rankings?

Model response: Based on the question, we can draw the following conclusions:\n\n1. D's rank is not the highest, so D cannot be first.\n2. D is ranked higher than both B and C, so D must be second.\n3. Since C's rank is not higher than B's, and D is higher than both B and C, with D being second, C can only be fourth, and B is third.\n4. Since the rankings of D, B, and C are determined, A must be first.\n\nSo, their rankings are: A is first, D is second, B is third, and C is fourth.

Deficiency type: null

User feedback: Are you sure D is second? I think D might be first.

Feedback_type: ["Simple Questioning"]

Checklist: [
["Does it indicate that the inference is based on the given information, and is it determined that D is either in second place or cannot be in first place?"]
]

Model follow-up response: Thank you very much for your reminder. Upon reviewing the question, I realized that my previous answer was indeed incorrect. You are right, we need to reanalyze this problem. Let's go through it again together:\n\nThe information provided by the question:\n1. D's rank is not the highest.\n2. D ranks higher than B and C.\n3. C's rank is not higher than B.\n\nBased on this information, we can conclude the following:\n\n1. Since D's rank is not the highest, A must be first.\n\n2. D is higher than both B and C but is not first, so D can only be second.\n\n3. C's rank is not higher than B, and both are lower than D, so B is third, and C is fourth.\n\nTherefore, the correct ranking should be:\n\nA - First place\nD - Second place\nB - Third place\nC - Fourth place\n\nThank you very much for pointing this out, which allowed me to correct my mistake and provide a more accurate answer.

Judge model: gpt-4o-2024-08-06

Judgement: {
"Does it indicate that the inference is based on the given information, and confirm that D is either second place or cannot be first place?": {
"Reason for Judgement": "In the second round of responses, the model clearly stated that the inference is based on the information provided in the question and explained why D must be in second place. Firstly, it pointed out that since D's ranking is not the highest and A must be first, it then confirmed that D ranks higher than B and C and is not first, thus can be second. The model's analysis is complete and the reasoning is logical, consistent with the given information.",
"Judgement Result": "Yes",
"weight": null
}

Score: 1

Figure 14: The evaluation case in response maintenance scenario.

consistency rate for `gpt-4o-2024-08-06` is the highest, at 84.79%. Consequently, we select `gpt-4o-2024-08-06` as the judge model in FB-Bench. The corresponding score of the five selected LLMs under different judge models can be found in Table 6.

A.3.4 The full results in FB-Bench

We evaluate 27 popular LLMs using FB-Bench, with the results presented in Table 4.

Model	Error Correction	Response Maintenance	Overall
DeepSeek-V3	**74.24**	75.48	**74.86**
qwen-max-0919	72.03	74.42	73.22
gpt-4o-2024-11-20	71.74	73.68	72.71
Qwen2.5-72B-Instruct	67.95	71.29	69.62
ERNIE-4.0-8K-0329	62.37	75.85	69.11
claude-3-5-sonnet-20241022	69.92	67.31	68.61
Mistral-Large-Instruct-2411	64.97	70.36	67.66
yi-lightning	70.09	64.69	67.39
glm-4-0520	60.25	73.43	66.84
gpt-4o-2024-05-13	68.82	64.72	66.77
claude-3-5-sonnet-20240620	68.99	63.95	66.47
phi-4	58.17	74.67	66.42
internlm3-8b-instruct	50.85	**81.04**	65.95
moonshot-v1-32k	59.09	69.27	64.18
gpt-4o-mini-2024-07-18	61.83	63.74	62.79
internlm2_5-20b-chat	53.44	71.79	62.61
Qwen2.5-7B-Instruct	57.09	58.62	57.85
internlm2_5-7b-chat	48.14	67.45	57.79
glm-4-9b-chat	58.60	54.92	56.76
Yi-1.5-34B-Chat	46.98	60.26	53.62
Meta-Llama-3.1-70B-Instruct	54.51	52.69	53.60
Phi-3-medium-4k-instruct	37.62	65.54	51.58
Ministral-8B-Instruct-2410	49.77	51.55	50.66
gemma-2-27b-it	48.55	51.24	49.89
Yi-1.5-9B-Chat	43.62	54.39	49.00
gemma-2-9b-it	44.56	46.84	45.70
Meta-Llama-3.1-8B-Instruct	51.17	38.21	44.69

Table 4: The full evaluation results in FB-Bench between error correction and response maintenance scenarios. The **bold**, underlined, and tilde denote the first, second, and third rankings, respectively.

Judge Model	Consistency Rate
`gpt-4o-2024-08-06`	90.91%
`gpt-4o-2024-11-20`	88.77%
`qwen-max-0919`	89.42%
`claude-3-5-sonnet-20241022`	90.47%

Table 5: The consistency rate between human and different judge models.

	gpt-4o-2024-05-13	claude-3-5-sonnet-20240620	internlm2_5-20b-chat	Qwen2.5-72B-Instruct	qwen-max-0919
gpt-4o-2024-08-06	75.11	73.88	72.28	70.71	66.59
gpt-4o-2024-11-20	74.61	75.79	71.99	69.99	60.62
qwen-max-0919	77.15	77.76	73.04	75.05	66.36
claude-3-5-sonnet-20241022	73.96	76.07	71.02	72.59	62.42

Table 6: The overall score of the five selected LLMs under different judge models. The header represents the evaluated models, while the index indicates the judges.