

# MedBench: A Large-Scale Chinese Benchmark for Evaluating Medical Large Language Models

Yan Cai<sup>1</sup>, Linlin Wang<sup>1,2\*</sup>, Ye Wang<sup>1</sup>, Gerard de Melo<sup>3,4</sup>, Ya Zhang<sup>2,5</sup>, Yanfeng Wang<sup>2,5</sup>, Liang He<sup>1</sup>

<sup>1</sup>East China Normal University

<sup>2</sup>Shanghai Artificial Intelligence Laboratory

<sup>3</sup>Hasso Plattner Institute

<sup>4</sup>University of Potsdam

<sup>5</sup>Shanghai Jiao Tong University

{51255901104,yewang}@stu.ecnu.edu.cn, {llwang,lhe}@cs.ecnu.edu.cn, gdm@demelo.org, {ya\_zhang,wangyanfeng}@sjtu.edu.cn

## Abstract

The emergence of various medical large language models (LLMs) in the medical domain has highlighted the need for unified evaluation standards, as manual evaluation of LLMs proves to be time-consuming and labor-intensive. To address this issue, we introduce MedBench, a comprehensive benchmark for the Chinese medical domain, comprising 40,041 questions sourced from authentic examination exercises and medical reports of diverse branches of medicine. In particular, this benchmark is composed of four key components: the Chinese Medical Licensing Examination, the Resident Standardization Training Examination, the Doctor In-Charge Qualification Examination, and real-world clinic cases encompassing examinations, diagnoses, and treatments. MedBench replicates the educational progression and clinical practice experiences of doctors in Mainland China, thereby establishing itself as a credible benchmark for assessing the mastery of knowledge and reasoning abilities in medical language learning models. We perform extensive experiments and conduct an in-depth analysis from diverse perspectives, which culminate in the following findings: (1) Chinese medical LLMs underperform on this benchmark, highlighting the need for significant advances in clinical knowledge and diagnostic precision. (2) Several general-domain LLMs surprisingly possess considerable medical knowledge. These findings elucidate both the capabilities and limitations of LLMs within the context of MedBench, with the ultimate goal of aiding the medical research community.

## Introduction

The advent of large language models (LLMs) has demonstrated substantial potential for diverse real-world applications, thanks to their remarkable language understanding capabilities. In the medical domain, a notable number of Chinese medical LLMs have successively emerged, including HuaTuo (Wang et al. 2023), BianQue (Chen et al. 2023), and DoctorGLM (Xiong et al. 2023), to better assist doctors in diverse tasks ranging from clinical diagnosis to disease prevention (Singhal et al. 2023). This underscores an urgent

\*Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

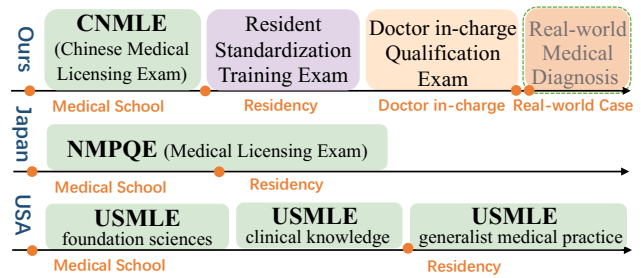


Figure 1: Comparison of procedures in different countries.

need for a standardized medical benchmark, capable of offering reliable and authoritative evaluations for such LLMs. Assessing the potential and inherent limitations of medical LLMs from diverse perspectives continues to present considerable challenges (Singhal et al. 2023; Chang et al. 2023).

The primary cause of this issue lies in the pronounced discrepancy between the existing benchmarks and the practical realities of medicine, leading to an urgent need for advances in evaluation standards. Widely used medical benchmarks such as MedQA (Jin et al. 2021) and MedMCQA (Pal, Umaphathi, and Sankarasubbu 2022) typically encompass publicly available medical question-answering datasets gathered from textbooks, research papers, and board exams. However, in light of recent research employing these open-access datasets for model training (Han et al. 2023), a process potentially resulting in data contamination with regard to the evaluation (Nori et al. 2023), it is increasingly clear that the prevailing benchmarks exhibit considerable limitations. A recent study has made advances by introducing a human evaluation framework along with MultiMedQA (Singhal et al. 2023), a multi-faceted benchmark that combines existing QA datasets with a novel set of online medical questions. Nonetheless, given the disparities among typologically diverse languages, we are unable to rely on this benchmark for assessing Chinese large language models, in particular considering the unique clinical standards and procedures across different countries illustrated in Figure 1.

Hence, several Chinese benchmarks have been proposed,

<p>【few_shot_examples】</p> <p>问题: 女性, 58岁, 9个月来巩膜及皮肤发黄, 伴皮肤瘙痒来诊。体检: 巩膜及皮肤明显黄染, 肝肋下4cm, 质硬无压痛, 脾肋下5cm, 考虑为原发性胆汁性肝硬化。原发性胆汁性肝硬化的主要病变在 ( )。</p> <p>(Question: A 58-year-old female. Has had jaundice of the sclera and skin for 9 months, accompanied by pruritus. Physical examination: obvious yellow staining of the sclera and skin, liver palpable 4 cm below the costal margin, hard texture without tenderness, spleen palpable 5 cm below the costal margin. Considered to be primary biliary cirrhosis. The main lesions of primary biliary cirrhosis are located in the ( ).)</p> <p>选项(Options):  A. 门管区胆管和胆小管(Bile ducts in portal tracts)  B. 肝总管(Hepatic ducts)  C. 毛细胆管(Bile canaliculi)  D. 肝细胞(Hepatocytes)  E. 胆总管(Common bile duct)</p> <p>回答(Answer):</p>	<p>ChatGPT: A✓</p> <p>ChatGLM: A✓</p> <p>Baichuan-13B: A✓</p> <p>HuaTuo: A. 门管区胆管和胆小管✓</p> <p>ChatMed: 答案是A✓</p>	<p>【few_shot_examples】</p> <p>症状: 检查中发现右甲状腺甲状腺包块, 遂于2018-08-14我院甲状腺外科就诊 (Symptoms: Examination revealed a nodule in the right lobe of the thyroid gland, so the patient visited the thyroid surgery department of our hospital on August 14, 2018.)</p> <p>检查(Examinations):</p> <p>Ground Truth: 甲状腺彩超 (Thyroid ultrasound)</p> <p>ChatGPT: 甲状腺超声检查, 甲状腺功能检查 (Thyroid ultrasound, thyroid function test)</p> <p>ChatGLM: 根据症状和检查结果, 患者可能存在以下疾病(Based on the symptoms and examination results, the patient may have the following conditions): ...</p> <p>Baichuan-13B: 根据您的描述和提供的检查结果, 您可能患有以下几种疾病(Based on your description and provided examination results, you may have the following conditions): ...</p> <p>HuaTuo: 甲状腺彩超(Thyroid ultrasound)...</p> <p>ChatMed: 甲状腺穿刺活组织病理学报告 (Thyroid fine needle aspiration biopsy pathology report)</p>
---	--	--

Figure 2: Examples of prompts and corresponding answers. The left side shows the prompt and responses of LLMs on an example question from an exam. The right side is an example of a real-world case.

including a medical NLP-task oriented one<sup>1</sup> and others based on the Chinese National Medical Licensing Examination (CNMLE), including MLEC-QA (Li, Zhong, and Chen 2021) and CMExam (Liu et al. 2023). While these benchmarks provide valuable insights, they are not exhaustive and may fall short in enabling us to comprehensively gauge the potential of LLMs with respect to all-round medical knowledge and their practical utility in real-world clinical diagnostic scenarios. In mainland China, a unique three-stage examination process is employed, key components of which, including the Resident Standardization Training and Doctor in-charge Qualification Exams, have largely been overlooked in prior work. Furthermore, the evaluation of real-world clinical practical skills has not been adequately incorporated in prior work.

To address these gaps and align with the learning and growth trajectory of Chinese doctors, we introduce MedBench, a novel large-scale Chinese medical benchmark, which encompasses both authentic three-stage medical examinations and real-world clinical diagnosis cases. It surpasses prior benchmarks by being exclusively sourced from the latest validated exams and expert-annotated EHRs, thereby ensuring compliance with medical standards and practices. Furthermore, we conduct extensive experiments and offer detailed analyses to provide diverse perspectives for evaluating clinical knowledge recall and reasoning capabilities of LLMs across a range of branches of medicine. The main findings on this benchmark are as follows:

- Chinese medical large language models underperform on this benchmark, necessitating substantial improvements in clinical knowledge and diagnostic accuracy, as well as refinement of their original in-context learning abilities.
- Several large language models designed for general-domain tasks possess substantial medical knowledge, thereby exhibiting promising potential.
- Human evaluation reveals that ChatGPT possesses rich

knowledge for clinical practices, while current Chinese medical LLMs lack high-quality conversational abilities and sufficient medical knowledge.

- ChatGPT’s unsatisfactory performance in case analysis, a question type that demands profound medical knowledge understanding and clinical reasoning abilities, reveals significant room for improvement in medicine.

In summary, this study presents a comprehensive benchmark that aligns with the practical realities of medicine in mainland China and provides profound empirical findings to reveal the medical capabilities and limitations of LLMs. Furthermore, we incorporate Item Response Theory (Birnbaum 1969) to further enhance our benchmark, with the ultimate goal of aiding the medical research community.

## The Proposed Benchmark

**Overview** For better evaluating medical foundation models, this paper proposes MedBench as a large benchmark with 40,041 exercises originating from both authentic medical examinations and real-world diagnostic and treatment cases. Specifically, we collect three-stage medical examinations that serve as a faithful reflection of the comprehensive process involved in obtaining medical licenses in mainland China, exemplifying essential medical knowledge. Moreover, we construct a number of real-world cases based on electronic health records that provide examination plans, diagnoses, and treatments based on patients’ symptoms, which can reveal the medical knowledge utilization and reasoning capabilities of LLMs in the real world.

**Construction and Statistics** Regarding the three-stage medical examinations, we collect representative exercises from the Chinese Medical Licensing Exam (CNMLE), Resident Standardization Training Exam, and Doctor in-charge Qualification Exam<sup>2</sup> across many recent years, covering 27,248, 2,841, and 8,927 questions, respectively. The examination exercises, depicted on the left of Figure 2, are

<sup>1</sup><https://github.com/michael-wzhu/PromptCBLUE>

<sup>2</sup><http://www.nmec.org.cn/>

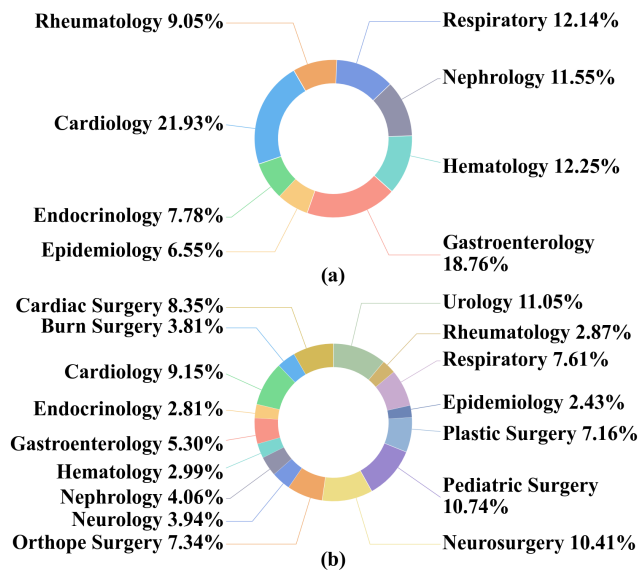


Figure 3: Branches of Medicine in MedBench. Part (a) represents the Resident Standardization Training Exam, and (b) reflects the Doctor in-charge Qualification Exam.

multiple-choice questions that can be categorized into three types:

- A1/A2/B: Single statement questions with one correct answer out of five options.
- A3/A4: A series of questions accompanied by a clinical case with one correct answer out of five options.
- Case Analysis: Given a clinical case, a series of questions are created with 6–12 options per question. Some questions may have more than one correct answer.

Figure 3 depicts the classification of branches of medicine for some examinations in MedBench. Note that we use internal medicine and surgery as representative examples for the latter two stages, considering the vast number of subcategories within these fields.

Furthermore, we collect over 2,000 real-world electronic health records, and employ experts to identify the symptoms, diagnoses, treatments, and examinations from the reports, resulting in 701 high-quality ones with an average length of 60–100 words per report (Figure 2). The annotators completed professional training and strictly following predefined annotation standards to ensure accuracy. Based on these annotated reports, a total of 1,025 question–answer pairs are eventually formed.

**Characteristics** MedBench surpasses existing benchmarks in several aspects: (1) Authenticity and Novelty. It exclusively leverages expert-annotated EHRs and authentic up-to-date medical examinations to mitigate contamination. (2) Comprehensiveness and Multi-facetedness. It is meticulously designed to align with Chinese medical standards and practices by incorporating three-stage multi-disciplinary examinations and real-world clinical cases. (3) Practicality. Human evaluation on clinical real-world cases ensures con-

gruence with practical realities of medicine, while difficulty-stratified divisions in MedBench enable rapid assessment.

## Evaluation

**Models and Evaluation Metrics** To evaluate the medical capabilities, we conduct assessments using MedBench with several representative LLMs from both the general and medical domains, including ChatGPT, ChatGLM (Zeng et al. 2023; Du et al. 2022), Baichuan-13B<sup>3</sup>, HuaTuo, and ChatMed<sup>4</sup>. Furthermore, we evaluate other Chinese medical LLMs, such as BianQue, but we have observed that it lacks the ability to deliver accurate and reasonable responses to multiple-choice questions.

For the three-stage multiple-choice examinations, we employ accuracy as the evaluation metric. When dealing with real-world cases, we combine expert-level human evaluation with the additional automatic evaluation metrics BLEU (Papineni et al. 2002) and ROUGE (Lin 2004).

**Experimental Settings** We conduct extensive experiments to evaluate the five-shot performance of LLMs, ensuring their capability to respond in a multiple-choice format. We leverage the API for ChatGPT<sup>5</sup> and opt for local deployment to facilitate evaluations for other LLMs. Furthermore, we partition the MedBench dataset based on exams, medical subdiscipline, and question types, and perform independent testing on each subset to enable a comprehensive analysis.

## Main Results

**Three-stage examination results.** In Table 1, we present a comprehensive analysis of the accuracy metrics for various LLMs across the three exams. A salient observation is that ChatGPT consistently surpasses other models, despite the latter being trained on extensive Chinese corpora or premium medical datasets. Nevertheless, our observations indicate that ChatGPT’s accuracy rate hovers around 50% for CNMLE and approximately 60% for other assessments, exposing substantial avenues for enhancement.

As depicted in Figure 5, ChatGPT exhibits subpar performance on questions pertaining to Traditional Chinese Medicine (TCM) and Chinese Western Medicine (CWM), with accuracy metrics oscillating between 40-45%. Conversely, LLMs trained on a more expansive Chinese dataset demonstrate a narrower performance disparity on TCM and CWM questions relative to ChatGPT. This suggests that a contributing factor to ChatGPT’s diminished efficacy on MedBench may be attributed to its limited exposure to Chinese data during its pretraining phase, consequently affecting its proficiency in Chinese medical knowledge.

It is imperative to underscore that Baichuan-13B and HuaTuo encounter difficulties on case analysis questions, highlighting potential domains for refinement in their logical reasoning or multi-turn dialogue competencies.

We also perform comparative evaluations with a number of recently-developed models, including MedLLaMA (Wu

<sup>3</sup><https://github.com/baichuan-inc/Baichuan-13B>.

<sup>4</sup><https://github.com/michael-wzhu/ChatMed>.

<sup>5</sup><https://chat.openai.com>.

LLM	CNMLE			Resident Standardization Training				Doctor in-Charge Qualification			
	Total	A1/A2/B	A3/A4	Total	A1/A2/B	A3/A4	Cases Analysis	Total	A1/A2/B	A3/A4	Cases Analysis
GPT-4	<b>64.88</b>	<b>63.08</b>	<b>69.03</b>	<b>75.64</b>	<b>77.08</b>	<b>75.13</b>	<b>75.00</b>	<b>68.45</b>	<b>71.91</b>	<b>68.24</b>	62.80
ChatGPT	49.57	49.40	51.85	60.59	61.30	58.72	62.96	58.75	58.04	59.73	<b>65.52</b>
ChatGLM	27.39	27.32	28.16	29.96	28.41	33.59	29.63	27.52	26.06	31.43	28.97
Baichuan-13B	30.47	30.54	29.63	34.97	37.26	32.70	22.65	29.56	31.31	31.65	17.89
HuaTuo	22.31	22.38	21.47	23.32	23.53	23.85	13.58	21.93	22.14	21.85	18.62
ChatMed	23.45	23.46	23.33	24.33	23.03	26.54	32.10	24.02	23.46	24.66	30.34

Table 1: Results on three-stage medical examinations of MedBench.

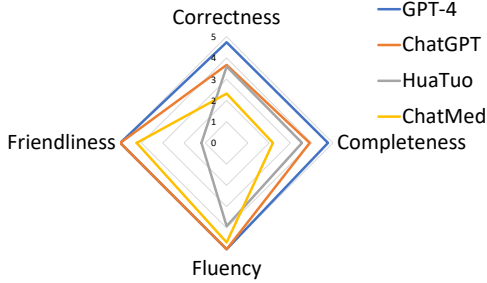


Figure 4: Human evaluation on clinical cases.

LLM	BLEU-1			ROUGE-L		
	E	T	D	E	T	D
GPT-4	<b>16.54</b>	<b>11.25</b>	<b>12.17</b>	0.95	0.82	2.66
ChatGPT	11.50	8.37	7.29	<b>15.39</b>	<b>11.30</b>	<b>10.04</b>
ChatGLM	0.77	2.12	1.90	1.79	3.62	3.29
Baichuan-13B	11.07	5.29	10.14	0.71	0.00	1.11
HuaTuo	1.56	3.77	2.32	2.43	0.52	1.06
ChatMed	4.34	4.88	5.79	1.35	0.55	0.26

Table 2: Evaluation on cases (E stands for Examinations, T stands for Treatments, and D stands for Diagnoses)

et al. 2023), Baize (Xu et al. 2023a), and ChatDoctor (Li et al. 2023). However, we have not noticed any significant enhancements in performance on our benchmark compared to these models. For example, MedLLaMA achieves an accuracy of 24.5% on the A1/A2/B type of questions, marginally outperforming ChatMed with a modest improvement of 1.5%.

**Real-world clinical case performance.** Table 2 provides the results of the assessment on real-world cases, and Figure 2 provides the prompt-response instances. We evaluated the outcomes using the BLEU and ROUGE F1-score metrics. It can be easily observed that ChatGPT and GPT-4 (OpenAI 2023) exhibit superior performance, underscoring their remarkable aptitude in medical and conversational question answering. However, it is notable that even for ChatGPT and GPT-4, the top-performers across multiple trials, the scores remain relatively moderate. This could be attributed, in part, to the fact that metrics such as BLEU and ROUGE might not holistically capture result quality. Furthermore, considerable scope remains for enhancing these LLMs’ capabilities in real-world clinical cases.

Beyond automated evaluations, the judgment of a post-

LLM	Vanilla prompt	Chain-of-Thought
ChatGLM-6B	42.69	43.37
Baichuan-13B	37.26	43.85

Table 3: A comparative assessment of Baichuan-13B and ChatGLM-6B using A1/A2 questions from the Resident Standardization Training Exam under different prompts.

graduate medical scholar was solicited to appraise the outputs delivered by the various LLMs on the real-world cases. The objective of this evaluation endeavor was to quantitatively assess the correctness, completeness, fluency, and friendliness of HuaTuo, ChatMed, ChatGPT, and GPT-4. Figure 4 provides the outcomes of this human evaluation. GPT-4 is found to consistently manifest superior performance across all delineated criteria, while ChatGPT marginally trails GPT-4, particularly on correctness and completeness. Conversely, HuaTuo’s rating in friendliness is somewhat diminished, predominantly attributable to sporadic generation of incongruous content. However, it is important to note that HuaTuo demonstrates pronounced levels of correctness and completeness, evincing profound medical knowledge, though its articulation warrants improvement. ChatMed, while laudable for its fluency and friendliness, registers suboptimal scores in correctness and completeness. Such disparities might be indicative of a potential attenuation of medical expertise during the fine-tuning phase.

## Quantitative Analysis

**Comparison on different question types.** Figure 5 provides the detailed results on CNMLE, the Resident Standardization Training Exam (denoted as “Resident”) and the Doctor in-charge Qualification Exam (“In-charge”). Clearly, ChatGPT continues to exhibit a substantial performance advantage over the other LLMs. A3/A4 questions and case analysis questions, characterized by their multi-question structure, present a substantial challenge to the conversational abilities of the models. If a given LLM fares poorly on these question types, it indicates its inherent limitations in efficiently managing multi-turn dialogue or questions necessitating intricate reasoning steps.

**Chain-of-Thought** In the experimental investigation, it is observed that prompts tailored to the LLM have the potential to enhance the reasoning capabilities of the LLM. Figure 7 illustrates the choice distribution of Baichuan-13B under both the vanilla prompt and Chain-of-Thought prompt-



Figure 5: Results for three-stage medical examinations.

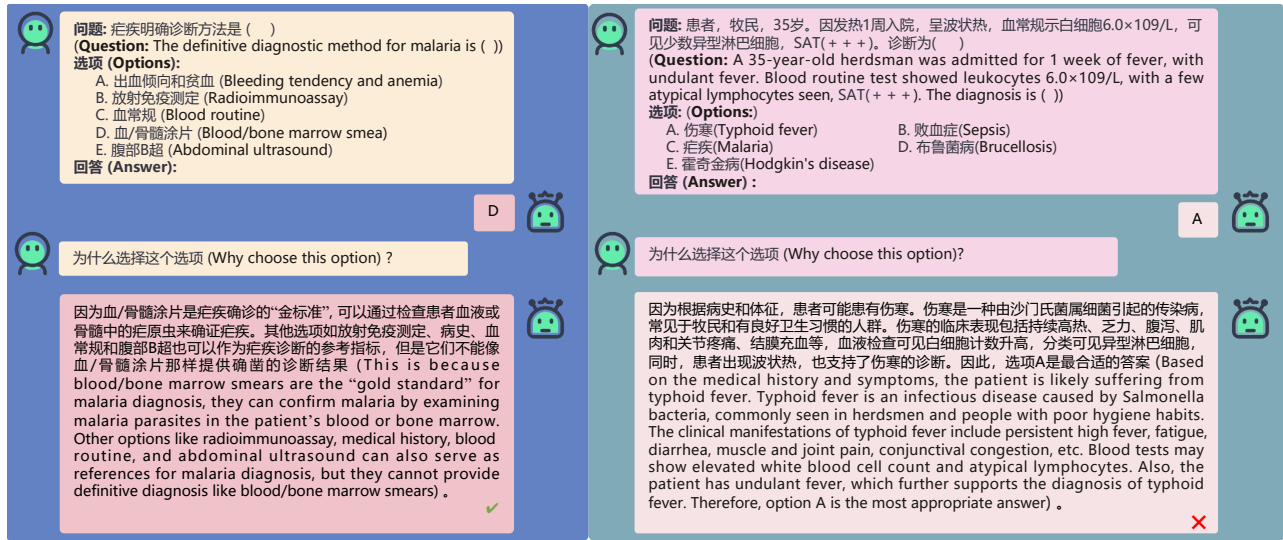


Figure 6: Examples of conversation demonstrating the reasoning capability of ChatGLM.

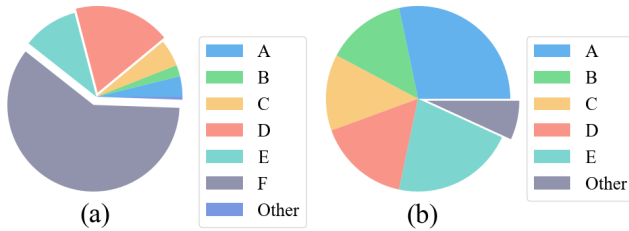


Figure 7: Distribution of choices made by Baichuan-13B under different prompts, where (a) and (b) depict the results using vanilla and Chain-of-Thought prompt, respectively.

ing. Notably, under the vanilla prompt, Baichuan-13B shows a strong inclination towards option F, although this option is not part of the valid choice set (A-E). In contrast, when using Chain-of-Thought prompting, Baichuan-13B primarily gravitates towards valid choices. Furthermore, Table 3 details the accuracy associated with both the vanilla prompt and the Chain-of-Thought one. The data suggests a significant improvement in the accuracy of Baichuan-13B when using Chain-of-Thought prompting. Parallel experiments are conducted on ChatGLM-6B, and the improvement in accuracy when using Chain-of-Thought prompting is negligible for ChatGLM-6B.

**Fewer questions, higher differentiation** During the assessment process, some shortcomings became apparent: (1) Certain LLMs manifest suboptimal inference speed on GPUs or high computational cost when interfacing via an API. Given that MedBench comprises approximately 40,000 questions, performing inference on the entirety of this dataset consecutively takes considerable time for LLMs. (2) While categorization of questions based on the type provides a rudimentary gauge of difficulty, it is important to acknowledge that questions within the same category can exhibit disparate levels of difficulty. (3) Assigning particularly

challenging questions to less capable LLMs is inadvisable, as it may culminate in uniformly diminished accuracy, making meaningful distinctions impossible.

To ameliorate these challenges, we propose methodological strategies to classify questions of analogous types according to their inherent difficulty gradients. By adopting this paradigm, we can optimize the evaluation process, allowing LLMs to undertake inference on a curtailed set of questions. As a result, this allows for a more nuanced alignment between LLMs and questions of commensurate difficulty tiers, congruent with their individual proficiencies. Specifically, we introduce an advanced evaluative framework for LLMs drawing inspiration from Item Response Theory (IRT). Our approach integrates the three-parameter logistic model (IRT-3PL), given by:

$$P(X_{ij} = 1|\theta_j) = c_i + (1 - c_i) \frac{1}{1 + e^{-a_i(\theta_j - b_i)}} \quad (1)$$

Here,  $\theta_j$  represents the proficiency of LLM  $j$ , and  $P(X_{ij} = 1|\theta_j)$  is the probability that an LLM  $j$  with proficiency  $\theta_j$  gives a correct response to question  $i$ . This equation highlights three crucial parameters for each question  $i$ : discrimination ( $a_i$ ), difficulty ( $b_i$ ), and the guessing factor ( $c_i$ ). In our approach, we assume constant values for both  $a_i$  and  $c_i$ , and then group items based on the shared difficulty metric  $b_i$ . For a specific difficulty subset, these three parameters remain unchanged. The equation captures the LLM's probability  $\theta_j$ , determined from the inference results when interacting with a specific subset.

A total of 7,335 questions were divided into 10 difficulty levels. To validate question differentiation, we analyze them using BLOOMZ-7.1B (Muennighoff et al. 2022), Qwen-7B<sup>6</sup>, ChatGLM-turbo (Zeng et al. 2023), Qwen-max (Bai et al. 2023) across various difficulty levels. Part (a) of Figure 8 displays the accuracy trends of these LLMs across

<sup>6</sup><https://github.com/QwenLM/Qwen-7B>.

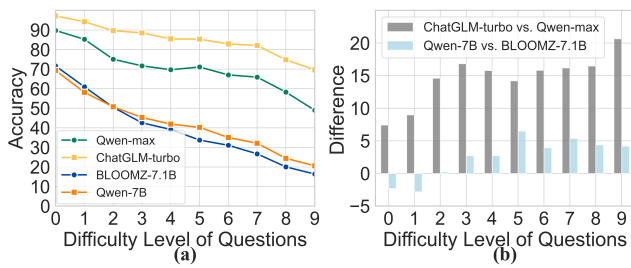


Figure 8: Evaluation of a few LLMs across varying levels of question difficulty, with Level 9 denoting the highest level of difficulty.

these levels, highlighting a decrease in accuracy with increasing difficulty. Figure 8(b) shows the difference in accuracy between BLOOMZ-7.1B and Qwen-7B, as well as between Qwen-max and ChatGLM-turbo, for different levels of question difficulty. It is noticeable that the differences between the LLMs are significant when the question difficulty is appropriate.

### Reasoning Abilities

In MedBench, there exist several typical types of reasoning:

- **Multi-condition single-hop reasoning:** A type of reasoning that requires the LLM to engage in single-hop reasoning based on a clinical scenario.
- **Statement identification:** A type of reasoning that requires the LLM to judge the correctness of multiple statements, which further tests the LLM’s knowledge base.
- **Multi-hop reasoning:** A type of reasoning with multiple questions, where the questions are related, requiring the LLM to perform multi-hop reasoning.

To ascertain the extent to which LLMs make accurate decisions based on their grasp of pertinent knowledge, we extended our inquiry, which entailed requesting the LLMs to furnish justifications for their responses, as depicted in Figure 6. Our observations reveal that LLMs can substantiate their answers when they are accurate. Conversely, in cases where erroneous responses are given, the accompanying explanations often prove illogical, which provides compelling evidence that these LLMs either lack the requisite knowledge in the domain or are incapable of rationalizing towards the correct solutions.

### Related Work

Traditionally, medical LLMs relied on classic medical QA benchmarks for evaluation. Some studies used USMLE for evaluation and achieved satisfactory results (Kung et al. 2023; Nori et al. 2023), with zero-shot GPT-4 achieving average scores of 86.65% and 86.7% on the Self-Assessment and Sample Exam parts of USMLE, respectively. However, questions from USMLE have a very distinct Western medical perspective that is notably different from that of Chinese medicine. As two distinct medical systems, the latter

includes traditional Chinese medicine, with obvious differences in concepts, diagnosis, and treatment methods. Therefore, the same symptoms may lead to different diagnoses and treatments. As a result, the USMLE cannot serve as an adequate benchmark for Chinese medical practice. MedQA, PubMedQA (Jin et al. 2019; Liévin, Hother, and Winther 2022), and MLEC-QA are also useful benchmarks for medical QA, comprising a large number of high-quality questions, but they were proposed early on, potentially leading to data contamination. Although MedMCQA is a relatively new benchmark, its questions originate primarily from Indian medical institutions. As a result, analogous to USMLE, it lacks Chinese medical content and thus cannot serve to evaluate Chinese medical question answering sufficiently well. CMExam is a novel benchmark for medical QA in Chinese, comprising the latest exam questions from CNMLE, which has been manually annotated by medical experts with assistance from ChatGPT to ensure high quality. However, CMExam only comprises questions from CNMLE, while overlooking other major medical exams in mainland China. Moreover, as a predominantly multiple choice dataset with only a small portion of fill-in-the-blank questions, CMExam lacks real-world medical data and scenarios. This constrains its ability to fully assess LLMs’ medical question answering capabilities. In fact, the optimal evaluation approach is to evaluate LLMs manually (Xu et al. 2023b). A recent work (Tu et al. 2024) employs specialist physicians to assess the quality. However, manual evaluation is a time-consuming and labor-intensive task, making it difficult to conduct on a large scale.

### Conclusion and Discussions

In this paper, we present MedBench, an exhaustive benchmark specifically designed for the domain of Chinese medical question answering. Preliminary empirical analyses underscore the suboptimal performance of Chinese medical LLMs when subjected to this benchmark, highlighting the need for improved clinical acumen and diagnostic precision. Furthermore, the adeptness of these LLMs in contextual learning requires further refinement.

During our empirical investigations, we found that certain models manifest pronounced hallucinatory behavior. As our research progresses, it is imperative to ensure data veracity while delving into the systematic evaluation of such hallucinatory phenomena. Furthermore, the appraisal of the LLM’s inferential competencies, as presented in this research, points to the need for further methodological refinements. Within the realm of clinical diagnostics, a diagnosis is typically predicated upon a plethora of corroborative evidence. In subsequent work, we plan to compile an enriched dataset, encompassing patients’ antecedent medical records and comprehensive physical examination narratives, to strengthen the evaluative framework for medical LLMs. Simultaneously, our findings underscore the efficacy of psychometric methodologies in aiding the evaluation. Moving forward, we intend to enhance and further integrate such methodologies into our model assessment paradigm.

## Acknowledgments

This work was supported by the National Innovation 2030 Major S&T Project of China (No.2022ZD0160700 & 2022ZD0160705) and the National Natural Science Foundation of China under Grant 12105105. The computations were performed in the ECNU Multifunctional Platform for Innovation (001). We also thank the financial support from Shanghai Chinafortune Co., Ltd.

## References

- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; Hui, B.; Ji, L.; Li, M.; Lin, J.; Lin, R.; Liu, D.; Liu, G.; Lu, C.; Lu, K.; Ma, J.; Men, R.; Ren, X.; Ren, X.; Tan, C.; Tan, S.; Tu, J.; Wang, P.; Wang, S.; Wang, W.; Wu, S.; Xu, B.; Xu, J.; Yang, A.; Yang, H.; Yang, J.; Yang, S.; Yao, Y.; Yu, B.; Yuan, H.; Yuan, Z.; Zhang, J.; Zhang, X.; Zhang, Y.; Zhang, Z.; Zhou, C.; Zhou, J.; Zhou, X.; and Zhu, T. 2023. Qwen Technical Report. arXiv:2309.16609.
- Birnbaum, A. 1969. Statistical theory for logistic mental test models with a prior distribution of ability. *Journal of Mathematical Psychology*, 6(2): 258–276.
- Chang, Y.; Wang, X.; Wang, J.; Wu, Y.; Yang, L.; Zhu, K.; Chen, H.; Yi, X.; Wang, C.; Wang, Y.; Ye, W.; Zhang, Y.; Chang, Y.; Yu, P. S.; Yang, Q.; and Xie, X. 2023. A Survey on Evaluation of Large Language Models. arXiv:2307.03109.
- Chen, Y.; Wang, Z.; Xing, X.; huimin zheng; Xu, Z.; Fang, K.; Wang, J.; Li, S.; Wu, J.; Liu, Q.; and Xu, X. 2023. BianQue: Balancing the Questioning and Suggestion Ability of Health LLMs with Multi-turn Health Conversations Polished by ChatGPT. arXiv:2310.15896.
- Du, Z.; Qian, Y.; Liu, X.; Ding, M.; Qiu, J.; Yang, Z.; and Tang, J. 2022. GLM: General Language Model Pretraining with Autoregressive Blank Infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 320–335.
- Han, T.; Adams, L. C.; Papaioannou, J.-M.; Grundmann, P.; Oberhauser, T.; Löser, A.; Truhn, D.; and Bressen, K. K. 2023. MedAlpaca – An Open-Source Collection of Medical Conversational AI Models and Training Data. arXiv:2304.08247.
- Jin, D.; Pan, E.; Oufattole, N.; Weng, W.-H.; Fang, H.; and Szolovits, P. 2021. What Disease Does This Patient Have? A Large-Scale Open Domain Question Answering Dataset from Medical Exams. *Applied Sciences*, 11(14).
- Jin, Q.; Dhingra, B.; Liu, Z.; Cohen, W.; and Lu, X. 2019. PubMedQA: A Dataset for Biomedical Research Question Answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Kung, T. H.; Cheatham, M.; Medenilla, A.; Sillos, C.; De Leon, L.; Elepaño, C.; Madriaga, M.; Aggabao, R.; Diaz-Candido, G.; Maningo, J.; and Tseng, V. 2023. Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models. *PLOS Digital Health*, 2(2): 1–12.
- Li, J.; Zhong, S.; and Chen, K. 2021. MLEC-QA: A Chinese multi-choice biomedical question answering dataset. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 8862–8874.
- Li, Y.; Li, Z.; Zhang, K.; Dan, R.; Jiang, S.; and Zhang, Y. 2023. ChatDoctor: A Medical Chat Model Fine-Tuned on a Large Language Model Meta-AI (LLaMA) Using Medical Domain Knowledge. *Cureus*.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Liu, J.; Zhou, P.; Hua, Y.; Chong, D.; Tian, Z.; Liu, A.; Wang, H.; You, C.; Guo, Z.; Zhu, L.; et al. 2023. Benchmarking Large Language Models on CMExam—A Comprehensive Chinese Medical Exam Dataset. *NeurIPS 2023*.
- Liévin, V.; Hother, C. E.; and Winther, O. 2022. Can large language models reason about medical questions? arXiv:2207.08143.
- Muennighoff, N.; Wang, T.; Sutawika, L.; Roberts, A.; Biderman, S.; Scao, T. L.; Bari, M. S.; Shen, S.; Yong, Z.-X.; Schoelkopf, H.; Tang, X.; Radev, D.; Aji, A. F.; Almubarak, K.; Albanie, S.; Alyafeai, Z.; Webson, A.; Raff, E.; and Raffel, C. 2022. Crosslingual Generalization through Multitask Finetuning. arXiv:2211.01786.
- Nori, H.; King, N.; McKinney, S. M.; Carignan, D.; and Horvitz, E. 2023. Capabilities of GPT-4 on Medical Challenge Problems. arXiv:2303.13375.
- OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774.
- Pal, A.; Umapathi, L. K.; and Sankarasubbu, M. 2022. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Conference on Health, Inference, and Learning*, 248–260. Proceedings of Machine Learning Research (PMLR).
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- Singhal, K.; Azizi, S.; Tu, T.; Mahdavi, S. S.; Wei, J.; Chung, H. W.; Scales, N.; Tanwani, A.; Cole-Lewis, H.; Pfohl, S.; Payne, P.; Seneviratne, M.; Gamble, P.; Kelly, C.; Babiker, A.; Schärli, N.; Chowdhery, A.; Mansfield, P.; Demner-Fushman, D.; Agüera y Arcas, B.; Webster, D.; Corrado, G. S.; Matias, Y.; Chou, K.; Gottweis, J.; Tomasev, N.; Liu, Y.; Rajkomar, A.; Barral, J.; Semturs, C.; Karthikesalingam, A.; and Natarajan, V. 2023. Large language models encode clinical knowledge. *Nature*, 620(7972): 172–180.
- Tu, T.; Palepu, A.; Schaekermann, M.; Saab, K.; Freyberg, J.; Tanno, R.; Wang, A.; Li, B.; Amin, M.; Tomasev, N.; et al. 2024. Towards Conversational Diagnostic AI. *arXiv preprint arXiv:2401.05654*.
- Wang, H.; Liu, C.; Xi, N.; Qiang, Z.; Zhao, S.; Qin, B.; and Liu, T. 2023. HuaTuo: Tuning LLaMA Model with Chinese Medical Knowledge. arXiv:2304.06975.



Wu, C.; Lin, W.; Zhang, X.; Zhang, Y.; Wang, Y.; and Xie, W. 2023. PMC-LLaMA: Towards Building Open-source Language Models for Medicine. arXiv:2304.14454.

Xiong, H.; Wang, S.; Zhu, Y.; Zhao, Z.; Liu, Y.; Huang, L.; Wang, Q.; and Shen, D. 2023. DoctorGLM: Fine-tuning your Chinese Doctor is not a Herculean Task. arXiv:2304.01097.

Xu, C.; Guo, D.; Duan, N.; and McAuley, J. 2023a. Baize: An open-source chat model with parameter-efficient tuning on self-chat data. *EMNLP 2023*.

Xu, J.; Lu, L.; Yang, S.; Liang, B.; Peng, X.; Pang, J.; Ding, J.; Shi, X.; Yang, L.; Song, H.; Li, K.; Sun, X.; and Zhang, S. 2023b. MedGPTEval: A Dataset and Benchmark to Evaluate Responses of Large Language Models in Medicine. arXiv:2305.07340.

Zeng, A.; Liu, X.; Du, Z.; Wang, Z.; Lai, H.; Ding, M.; Yang, Z.; Xu, Y.; Zheng, W.; Xia, X.; et al. 2023. Glm-130b: An open bilingual pre-trained model. *ICLR 2023*.