

MMedAgentBench: A Multi-Stage Benchmark for Evaluating LLM-based Medical Agents

Lu Sun^{1*}, Xiangyi Zhang^{4*}, Xiangyang Zhu¹, Kaiyuan Ji^{5,1}, Zijian Chen^{1,2},
Qiang Hu³, Zicheng Zhang¹, Yuan Tian^{1†}, Guangtao Zhai^{1,2}

¹Shanghai Artificial Intelligence Laboratory

²Institute of Image Communication and Network Engineering, Shanghai Jiao Tong University

³Cooperative Medianet Innovation Center, Shanghai Jiao Tong University

⁴SJTU-Paris Elite Institute of Technology, Shanghai Jiao Tong University

⁵School of Information and Electronic Engineering, East China Normal University

Correspondence: tianyuan168326@outlook.com

Abstract

Large language models (LLMs) have demonstrated remarkable potential in medical AI through their natural language understanding and reasoning capabilities. However, real clinical practice involves multi-stage, dynamically evolving decision-making where physicians continuously adjust pathways based on updated examination results—far more complex than single-turn question answering. Existing medical benchmarks focus on static scenarios, failing to evaluate LLMs in realistic clinical environments. We present **MMedAgentBench**, a comprehensive benchmark for evaluating LLM-based medical agents in multi-stage clinical decision-making. Unlike previous benchmarks, MMedAgentBench captures the temporal complexity of clinical workflows by simulating authentic patient journeys across multiple diagnostic and treatment stages. The benchmark comprises a high-quality dataset of over 150 de-identified clinical cases from tertiary hospitals, encompassing 500+ clinical stages, with each case meticulously structured into sequential stages and annotated by board-certified physicians. Our evaluation framework simultaneously assesses diagnostic accuracy and resource efficiency, requiring agents to plan examinations and control costs while navigating progressive information disclosure. This enables systematic evaluation of medical agents across dynamic decision-making, cost control, and multi-stage reasoning dimensions, providing a realistic testbed for advancing medical AI toward clinical deployment.

1 Introduction

With the rapid advancement of large language model (LLM) technology, application of AI in the medical field has gradually become a focal point. LLMs can process and understand vast amounts

medical data. Particularly when faced with diverse and extensive medical information, they have potential in medical decision support through their natural language processing and reasoning capabilities (Thirunavukarasu et al., 2023; Singhal et al., 2023; Ji et al., 2026). In medicine, doctors must analyze extensive data—including patient histories, symptom presentations, and test results—while making accurate diagnostic and treatment decisions within limited timeframes. This poses a significant challenge for manual processing. LLMs, however, leverage their robust information retrieval and reasoning capabilities to extract relevant insights from vast medical literature and case data, thereby accelerating the diagnostic process and enhancing accuracy.

Particularly during disease diagnosis, LLM-based agent models can provide effective assistance throughout the multi-stage, dynamic decision-making process. By simulating a doctor’s thought process, these agent models adapt subsequent decision paths based on current examination results, enabling more precise decisions in complex, evolving diagnostic scenarios (Li et al., 2025). Doctors must repeatedly adjust diagnostic hypotheses based on varying symptoms and test outcomes to narrow down the most probable conditions. Therefore, LLM-based agents which are capable of flexibly adapting to changes and progressively refining pathways based on each round of test results are particularly crucial in medical clinical decision-making. They can not only process medical texts but also dynamically adjust diagnostic strategies, helping doctor make the most appropriate treatment decisions at every stage.

However, most current benchmarks focus on static, single-task processing, overlooking the complexity of multi-stage decision-making inherent in medical practice. These methods fail to authentically reflect how doctors dynamically adjust examinations and treatment plans based on evolving

*Equal contribution.

†Corresponding author.

patient conditions and test results, thereby limiting the applicability of current standards for assessing diagnostic systems.

To address this, we propose an innovative evaluation framework built upon a high-quality dataset comprising over 150 clinical cases and 500+ diagnostic stages. Unlike static approaches, our framework simulates the iterative reasoning paths of doctors, allowing diagnostic strategies to dynamically adapt based on intermediate findings. This methodology provides a comprehensive assessment of LLM agents in complex clinical scenarios, delivering more precise evaluations to advance their application in real-world decision support.

Our contributions are threefold:

- We propose **MMedAgentBench**¹, a benchmark comprising over 150 clinical cases and more than 500 diagnostic stages, designed to systematically evaluate LLM-based medical agents in multi-stage clinical decision-making.
- This benchmark is an LLM-based agent evaluation framework that simulates a doctor’s reasoning and judgment during multi-stage decision-making. The LLM plays a central role here, handling dynamic decision-making and outcome updates at each stage.
- We established multiple evaluation dimensions within the benchmark framework, including diagnostic accuracy, cost-effectiveness, and multi-stage decision-making, to provide a more comprehensive assessment of LLM applications in medical diagnosis.

2 Related Work

A large body of work has focused on evaluating medical AI systems through benchmark datasets. Early efforts such as MedQA (Jin et al., 2021) and MedMCQA (Pal et al., 2022) introduced large-scale multiple-choice examinations to assess factual medical knowledge, while benchmarks like PubMedQA (Jin et al., 2019) and emrQA (Pampari et al., 2018) targeted question answering from biomedical literature or clinical notes. More recently, comprehensive resources such as MultifacetEval (Zhou et al., 2024) expanded the scope to cover broader clinical reasoning tasks. Beyond task accuracy, recent evaluation studies have considered multidimensional

model capabilities (Zhang et al., 2025a,b) and task utility under resource constraints (Tian et al., 2024, 2025b). However, these benchmarks are primarily designed around single-turn question answering or isolated decision-making tasks, which fail to capture the temporal dependencies and sequential nature of real-world clinical workflows.

With the rapid development of large language models (LLMs), increasing research efforts have extended medical agents from static question answering to interactive and workflow-oriented scenarios. Several representative works are particularly noteworthy in this direction.

First, MedAgentBench (Jiang et al., 2025) introduced a virtual electronic health record (EHR) environment for comprehensive evaluation of medical LLM agents. The benchmark covers 300 clinical tasks across 10 types of medical scenarios and supports interactive simulations through FHIR-compliant interfaces. Unlike traditional static QA, MedAgentBench emphasizes tool usage, information retrieval, and task planning capabilities. Second, τ bench (Yao et al., 2024) highlights the interaction ability of language agents with users and tools in real-world domains. It simulates dynamic dialogues and tool invocation processes, while also proposing new reliability metrics to assess an agent’s consistency and robustness across multiple interactions. Finally, AgentClinic (Schmidgall et al., 2026) is specifically designed for medical settings, supporting doctor–patient dialogues, multimodal information acquisition and tool usage. However, it still lacks explicit control over multi-stage patient journeys, making it less aligned with the sequential nature of real-world clinical workflows.

3 MMedAgentBench: A Multi-Stage Medical Agent Benchmark

We present MMedAgentBench, a comprehensive benchmark designed to evaluate medical AI agents in realistic multi-stage clinical workflows, shown as Figure 1. Unlike existing medical benchmarks that focus on single-turn question-answering or isolated diagnostic tasks, MMedAgentBench captures the temporal complexity and resource constraints inherent in real-world clinical decision-making. Our benchmark simulates complete patient journeys through multiple clinical stages, where agents must balance diagnostic accuracy with examination costs while maintaining coherent reasoning across stages.

¹The project is available at <https://github.com/ZhangXiangyi0523/MMedAgentBench>.




Figure 1: Overview of the MMedAgentBench. The benchmark models real-world clinical decision-making as a sequential, cost-aware reasoning process. Each case unfolds through multiple diagnostic stages, where the agent iteratively selects examinations, interprets results, and updates its diagnostic hypotheses. The evaluation metrics consider both diagnostic accuracy and resource costs, aiming to promote more balanced clinical decision-making.

3.1 Multi-Stage Clinical Workflow Design

Real clinical practice rarely involves single-point decisions. Patients typically progress through multiple stages of care, each with distinct objectives and available resources. We model this reality by organizing evaluation scenarios into sequential stages $S = \{s_1, s_2, \dots, s_n\}$ that mirror actual hospital workflows. Each stage s_i represents a distinct clinical phase such as initial assessment, follow-up examination, or pre-discharge evaluation. The agent must navigate through these stages sequentially, making decisions based on available information at each stage while considering the accumulated evidence from previous stages.

For each stage s_i , we define two catalogs: C_i^{exam} containing available diagnostic examinations and $C_i^{patient}$ containing accessible patient records. The agent can select items from these catalogs, with each selection revealing detailed clinical information. Let $X_i \subseteq C_i^{exam} \cup C_i^{patient}$ denote the set of unique items selected by the agent at stage s_i for cost calculation and historical retrieval. Formally, at stage s_i , the agent’s action space is defined as:

$$A_i = \{\text{pick}(c) : c \in C_i^{exam} \cup C_i^{patient}\} \cup \{\text{retrieve}(h) : h \in H_{i-1}\} \quad (1)$$

where $H_{i-1} = \bigcup_{j=1}^{i-1} X_j$ represents the set of his-

torical items selected in all previous stages.

The design of MMedAgentBench addresses several critical challenges in medical AI evaluation. First, temporal coherence is enforced through our stage-wise information access control, where agents cannot access future-stage information, preventing unrealistic perfect foresight scenarios. This mirrors how clinicians must make decisions based on currently available information without knowing future developments.

Second, our historical information retrieval mechanism allows agents to access previously obtained information without additional cost, simulating how clinicians reference prior test results and observations. This design tests an agent’s ability to synthesize longitudinal patient data while avoiding redundant examinations.

Finally, the transition between stages is governed by a judgment function $J : (D_i, T_i) \rightarrow \{0, 1\}$, where D_i and T_i represent the diagnosis and treatment outputs at stage s_i . The agent proceeds to stage s_{i+1} only if $J(D_i, T_i) = 1$, ensuring that stage-specific clinical objectives are met before progression. This mirrors real clinical workflows where certain milestones must be achieved before proceeding to the next phase of care. The judgment function is implemented using a large language

model to assess whether the agent’s conclusions adequately address the clinical requirements of the current stage.

3.2 Dataset Construction

3.2.1 Clinical Data Collection and Structuring

Our dataset is constructed from authentic clinical cases collected from real hospital settings, ensuring the benchmark reflects genuine medical complexity and decision-making scenarios. We collaborate with multiple tertiary hospitals to obtain de-identified patient records spanning various medical specialties, including internal medicine, surgery, pediatrics, and emergency care. Each case represents a complete patient journey from initial presentation through discharge or transfer, capturing the full spectrum of clinical decision-making.

The raw clinical data undergoes a rigorous structuring process to transform unstructured medical records into our standardized format. We employ a structured representation that separates examination reports from patient records, each organized by clinical stages. For each item x in the dataset, its JSON fields correspond to the tuple $(title_x, details_x, price_x, stage_x)$ where $title_x$ serves as the unique identifier, $details_x$ contains the complete clinical information, $price_x \in \mathbb{R}_{\geq 0}$ represents the associated clinical examination cost based on actual hospital pricing, and $stage_x \in S$ indicates the stage affiliation.

3.2.2 Expert Annotation and Quality Control

To ensure clinical accuracy and relevance, our annotation process involves multiple rounds of expert review. Board-certified physicians from relevant specialties annotate each case with clinical key points rather than exact text matches. Specifically, for stage s_i , we define the ground truth as $G_i = (P_i^{diag}, P_i^{treat})$, where $P_i^{diag} = \{p_{i,1}^{diag}, p_{i,2}^{diag}, \dots, p_{i,m}^{diag}\}$ represents diagnostic key points and $P_i^{treat} = \{p_{i,1}^{treat}, p_{i,2}^{treat}, \dots, p_{i,k}^{treat}\}$ represents treatment key points. This point-based representation allows for semantic flexibility while maintaining evaluation rigor.

For each stage s_i , annotators complete a structured form with three components: the diagnostic and treatment key-point sets P_i^{diag} and P_i^{treat} , the reference cost baseline $Cost_i^{ref}$ defined as the sum of the prices of essential examinations selected for that stage, and a stage progression criterion specifying the minimum clinical objectives required to

proceed. This structure localizes disagreements to the key-point sets, reference examination selection, or stage progression criterion.

The annotation guidelines ensure consistency across different annotators. Cases with low agreement undergo field-level review and discussion until consensus is reached; the original entries, disputed fields, and final consensus version are retained for audit. Furthermore, we implement quality control measures including random audits and cross-validation to maintain high annotation standards throughout the dataset construction process.

3.2.3 Ethics and Data Governance

All clinical records underwent de-identification before annotation and benchmark construction. Only the processed records entered the annotation workflow. Medical experts reviewed the data-governance procedure and associated safeguards. Further details are provided in Appendix A.3.

3.3 Evaluation Framework

Our evaluation framework integrates both clinical accuracy and resource efficiency, reflecting the dual objectives of real-world medical practice. We assess agent performance through comprehensive metrics that balance diagnostic precision with cost-conscious decision-making.

3.3.1 Cost Tracking and Resource Utilization

Throughout this paper, cost refers to the hospital-listed cost of clinical examinations selected by the agent; LLM inference costs are not included. Healthcare resources are inherently limited and expensive. To evaluate cost efficiency, we establish a baseline cost for each stage through expert annotation. For each stage s_i , experienced physicians annotate a reasonable reference cost $Cost_i^{ref}$ that represents the expected expense for appropriate clinical decision-making at that stage. This reference cost serves as a normalization factor for evaluating the agent’s resource utilization.

For stage s_i , the actual clinical examination cost of the unique items in X_i is calculated as:

$$Cost_i = \sum_{x \in X_i \cap C_i^{exam}} price_x \quad (2)$$

Note that patient record items have zero clinical examination cost, reflecting their nature as readily available clinical information. To prevent artificial inflation of costs, we count each item’s examination cost only once per stage.

The cost efficiency for each stage is then evaluated through a cost ratio:

$$\text{CostRatio}_i = \frac{\text{Cost}_i}{\text{Cost}_i^{\text{ref}}} \quad (3)$$

For the overall evaluation, we compute the average cost ratio across all stages that the agent has experienced. Let m denote the number of stages the agent successfully completes, the average cost ratio is:

$$\text{CostRatio}_{\text{avg}} = \frac{1}{m} \sum_{i=1}^m \text{CostRatio}_i \quad (4)$$

3.3.2 Clinical Accuracy Metrics

We evaluate clinical accuracy using LLM-based semantic matching to assess coverage of key points. Let $\hat{D}_i$ and $\hat{T}_i$ denote the agent’s diagnosis and treatment outputs at stage s_i . We define a semantic matching function $M : (\text{prediction}, \text{key point}) \rightarrow \{0, 1\}$ that determines whether a key point is covered by the prediction.

For stages that the agent successfully completes, the stage-wise diagnostic and treatment accuracies are calculated as:

$$\text{Acc}_i^{\text{diag}} = \frac{1}{|P_i^{\text{diag}}|} \sum_{p \in P_i^{\text{diag}}} M(\hat{D}_i, p) \quad (5)$$

$$\text{Acc}_i^{\text{treat}} = \frac{1}{|P_i^{\text{treat}}|} \sum_{p \in P_i^{\text{treat}}} M(\hat{T}_i, p) \quad (6)$$

For stages that the agent fails to reach or complete, we assign an accuracy of 0, reflecting the failure to address those clinical scenarios. Let n be the total number of stages in the complete patient journey. The aggregate accuracies are computed as:

$$\text{Acc}_{\text{avg}}^{\text{diag}} = \frac{1}{n} \sum_{i=1}^n \text{Acc}_i^{\text{diag}}, \quad (7)$$

$$\text{Acc}_{\text{avg}}^{\text{treat}} = \frac{1}{n} \sum_{i=1}^n \text{Acc}_i^{\text{treat}} \quad (8)$$

where $\text{Acc}_i = 0$ for all $i > m$ (stages not reached by the agent). This formulation ensures that agents are penalized for incomplete patient journeys, encouraging comprehensive clinical coverage while maintaining accuracy at each completed stage.

4 Experiments

4.1 Experimental Setup

4.1.1 Benchmark Details

MMedAgentBench represents a high-quality, comprehensive medical agent benchmark dataset specifically designed for evaluating AI agents in clinical decision-making scenarios. The dataset comprises 167 carefully curated patient cases, as shown in Table 1.

Description	Value
Number of Patients	167
Number of Clinical Stages	527
Number of Examination Records	4,019
Number of Patient Records	701
Average Stages per Patient	3.16
Average Exam Records per Patient	24.07
Average Patient Records per Patient	4.20

Table 1: Summary of MMedAgentBench.

4.1.2 Models

We conduct an extensive evaluation covering a diverse spectrum of large language models that vary in scale to ensure a comprehensive capability assessment. Among open-weight models, our benchmark includes the Qwen2.5 family (14B/72B) (Qwen Team, 2024) alongside the next-generation Qwen3 series (14B/235B-A22B) (Yang et al., 2025), the Llama series (Touvron et al., 2023; Meta AI, 2024; Grattafiori et al., 2024) encompassing Llama-3.1-70B and Llama-3.1-405B, as well as GLM-Z1-32B (GLM et al., 2024), DeepSeek-v3.1 (DeepSeek-AI, 2024), Hunyuan-A13B (Tencent Hunyuan Team, 2025), and Kimi-K2 (Team et al., 2025). For closed-weight models, we evaluate the GPT family—including GPT-3.5, GPT-4o, GPT-4.1, and GPT-5 (OpenAI, 2023a,b)—alongside the Claude-Sonnet series (3.5, 3.7, and 4) (Anthropic, 2025), Gemini-2.5-Flash (Comanici et al., 2025), Grok-4 (xAI, 2025), and Doubao-1.5-Pro (ByteDance Seed, 2025). All models are tested under unified prompting strategies and standardized evaluation protocols to guarantee fairness and comparability across results.

For the stage judgment function $J : (D_i, T_i) \rightarrow \{0, 1\}$ and the semantic matching function $M : (\text{prediction}, \text{key point}) \rightarrow \{0, 1\}$, we implemented them using Gemini-2.5-Flash and GPT-4o, respectively. The former requires high-level integration

Open-Weight Models				Closed-Weight Models			
Model	Acc ^{diag}	Acc ^{treat}	Cost _↓	Model	Acc ^{diag}	Acc ^{treat}	Cost _↓
Llama-3.1-70B	7.75	4.47	0.46	GPT-3.5	15.30	10.50	0.44
Llama-3.1-405B	16.07	8.60	0.75	GPT-4o	22.13	18.68	0.74
Qwen2.5-14B	17.47	14.04	0.66	GPT-4.1	<u>34.90</u>	<u>37.00</u>	<u>0.68</u>
Qwen2.5-72B	21.76	17.77	0.85	GPT-5	35.62	38.33	0.93
Qwen3-14B	17.66	16.28	<u>0.63</u>	Claude-Sonnet-3.5	28.10	28.10	1.45
Qwen3-235B-A22B	36.67	31.69	1.53	Claude-Sonnet-3.7	30.10	28.70	1.32
Hunyuan-A13B	13.55	18.68	0.75	Claude-Sonnet-4	31.89	30.13	1.30
GLM-Z1-32B	22.59	20.69	1.28	Gemini-2.5-Flash	32.88	28.67	1.41
DeepSeek-V3.1	25.12	23.99	0.69	Grok-4	31.35	29.11	0.98
Kimi-K2	<u>28.30</u>	<u>31.10</u>	1.14	Doubao-1.5-Pro	33.25	30.78	1.24

Table 2: Evaluation results on MMedAgentBench. Acc^{diag} and Acc^{treat} denote diagnostic and treatment accuracy respectively (in %). CostRatio measures resource utilization efficiency, where smaller values indicate better cost control relative to the expert baseline. We use **bold** to highlight the best-performing model in each category, and underline for the second-best.

and reasoning over long clinical texts, while the latter focuses solely on evaluating semantic similarity and matching accuracy.

4.2 Results

The benchmark consists of a vast array of complex medical data, ranging from simple outpatient records to intricate multi-disease co-occurrence patterns. Based on MMedAgentBench, we evaluated multiple large models’ diagnostic capability (Acc^{diag}), treatment capability (Acc^{treat}), and cost ratio (CostRatio), results are shown as Table 2. We run the benchmark for 5 times to get the average value. Even the best-performing models, Qwen3-235B-A22B and GPT-5, achieve only 36.67% and 35.62% diagnostic accuracy respectively, and 31.69% and 38.33% treatment accuracy, highlighting the substantial challenges posed by our multi-stage clinical diagnosis dataset. This relatively modest performance ceiling underscores the complexity of real-world clinical decision-making, where agents must navigate through sequential stages, synthesize evolving patient information, and make interconnected diagnostic and treatment decisions. Furthermore, the wide variation in CostRatio across models (ranging from 0.44 to 1.53) provides valuable insights into different models’ examination ordering strategies, enabling researchers to better understand the trade-offs between diagnostic thoroughness and resource efficiency, and identify models that can achieve clinical objectives while maintaining cost-conscious medical practice.

4.3 Model Scale Effects

Evaluation results reveal a clear correlation between model scale and clinical performance within the same model families. Examining the Llama-3.1 series, we observe substantial improvements as parameters increase from 70B to 405B: Acc^{diag} improves from 7.75% to 16.07%, while Acc^{treat} rises from 4.47% to 8.60%. Similarly, the Qwen2.5 series demonstrates consistent scaling benefits, with Acc^{diag} increasing from 17.47% (14B) to 21.76% (72B) and Acc^{treat} from 14.04% to 17.77%. These within-family improvements underscore that parameter scaling directly enhances models’ ability to process complex medical information and generate accurate clinical decisions.

However, the results also highlight that base model capabilities and training methodologies are equally crucial for medical tasks. Notably, Qwen2.5-14B achieves superior performance (Acc^{diag}: 17.47%, Acc^{treat}: 14.04%) compared to Llama-3.1-70B despite having significantly fewer parameters, suggesting that architectural design and pre-training quality on medical-relevant data can compensate for raw parameter count.

The scaling effects observed across different model families suggest that while increasing parameters generally improves medical reasoning capabilities, the efficiency of this scaling varies significantly based on the underlying model architecture and training approach. This indicates that future development of medical AI systems should balance parameter scaling with targeted improvements in




Figure 2: Acc^{diag} and Acc^{treat} distribution of Qwen and Llama family.

base model design and domain-specific pre-training strategies.

4.4 Comparison of Open-Weight and Closed-Weight Models

The performance gap between open and closed-weight models has narrowed significantly. Notably, open-weight Qwen3-235B-A22B achieves the highest diagnostic accuracy (Acc^{diag} 36.67%), surpassing the leading closed-weight GPT-5 (35.62%). However, GPT-5 retains the top spot in treatment recommendations (Acc^{treat} 38.33%), indicating closed models still excel in therapeutic strategies.

Among closed-weight models, GPT-4.1 ranks second (Acc^{diag} : 34.90%, Acc^{treat} : 37.00%) with remarkably low cost (CostRatio: 0.68). Doubao-1.5-Pro (Acc^{diag} : 33.25%, Acc^{treat} : 30.78%) also outperforms prominent models like Gemini-2.5-Flash and Claude-Sonnet-4, proving that top-tier performance does not necessitate high costs.

In the open-weight domain, Kimi-K2 achieves an Acc^{treat} of 31.10%, surpassing closed counterparts like Gemini-2.5-Flash and Claude-Sonnet-4. Meanwhile, DeepSeek-V3.1 offers a strong efficiency-accuracy balance (25.12%, 23.99%), outperforming GPT-4o. These results demonstrate open-weight models effectively rivaling top-tier proprietary systems.




Figure 3: Clinical Examination Cost-Efficiency Analysis. Comparison of (a) diagnostic and (b) treatment accuracy vs. CostRatio. The arrow points to the optimal balance. Legend: Blue circles = Open-weight; Red diamonds = Closed-weight.

4.5 Cost Sensitivity and Diagnostic-Therapeutic Balance

Results reveal distinct trade-offs between performance and cost across different models. We visualize accuracy versus CostRatio in Figure 3.

Overall, closed-weight models like Claude-Sonnet-4 (1.30), Gemini-2.5-Flash (1.41), and GPT-5 (0.93) exhibit cost ratios close to or slightly exceeding real-world medical scenarios. In contrast, certain open-weight models (e.g., DeepSeek-V3.1 at 0.69, Qwen2.5-14B at 0.66) demonstrate lower costs while maintaining competitive performance. However, high cost does not directly equate to high efficiency: at similar cost levels, significant gaps exist in diagnostic accuracy among different models. For instance, GPT-5 achieves a significantly higher diagnostic accuracy (35.62%) than Claude-Sonnet-4 (31.89%) at a lower cost (0.93 vs. 1.30), demonstrating superior performance-cost ratio.

It is noteworthy that DeepSeek-V3.1 (CostRatio = 0.69) and GPT-4o (CostRatio = 0.74) present




Figure 4: Case illustration of diagnostic bias. By focusing on superficial features like routine blood work instead of MRI/biopsy results, the model incorrectly diagnosed prostate cancer as anemia and liver issues.

a compelling contrast: both models incur similar medical examination costs, yet DeepSeek-V3.1 outperforms GPT-4o in both diagnostic accuracy (25.12% vs. 22.13%) and treatment recommendation accuracy (23.99% vs. 18.68%). Similarly, comparing Qwen2.5-72B (CostRatio = 0.85) with GPT-5 (CostRatio = 0.93), we observe that despite similar cost levels, GPT-5 significantly outperforms Qwen2.5-72B in diagnostic accuracy (35.62% vs. 21.76%) and treatment recommendation accuracy (38.33% vs. 17.77%). These contrasts demonstrate that effective utilization of medical examinations is more critical than merely minimizing costs. The key challenge in medical AI is not just reducing examination costs, but effectively synthesizing the resulting information to reach accurate clinical conclusions—achieving the optimal balance between diagnostic accuracy and resource efficiency.

4.6 Case Study: Misleading Information Distorting Diagnostic Direction

When confronted with complex medical records, models are susceptible to misleading clues that can cause significant deviations in diagnostic direction.

In this case shown as Figure 4, the data input to the model included comprehensive clinical information: the patient’s complete blood count results, MRI imaging descriptions, and biopsy reports. The blood test results indicated mild anemia and mild liver function abnormalities, while the MRI and biopsy information revealed abnormal prostate signals and related histological features, with the actual diagnosis being prostate cancer. However, when generating its diagnosis, the model primarily focused on hematological indicators, outputting a diagnosis and treatment recommendation of “mild anemia and liver function abnormalities,” completely ignoring the malignant

tumor risk suggested by the imaging and record.

This outcome demonstrates that even with comprehensive clinical information, the model may exhibit systemic bias due to overemphasis on superficial features. Its reasoning process showed excessive sensitivity to prominent data (e.g., blood indicators) while undervaluing critical information (e.g., imaging and pathology reports). The diagnostic logic lacked hierarchical judgment capabilities, failing to perform causal reasoning and diagnostic hypothesis verification based on the overall clinical picture as clinical physicians do.

4.7 Evaluation Validity and Robustness

We validate our metrics through a hybrid strategy. First, human-defined ground truth ensures reliability; reference points are explicitly extracted by experts, limiting the semantic matcher to the annotated clinical key points. Second, judge sensitivity analysis assesses the stability of model rankings across different judge models (Table 3). Finally, human alignment checks compare our automated evaluation with medical-expert judgments (Table 4).

To assess robustness to evaluator choices, we use three judge/matcher configurations: Gemini-2.5-Flash with GPT-4o, GPT-5 with Claude-Sonnet-4, and GPT-4o with GPT-3.5. In each pair, the first model serves as the stage judge and the second as the semantic matcher. As shown in Table 3, absolute scores vary slightly, while the overall rankings remain broadly stable.

To ensure clinical validity, three board-certified physicians independently evaluated a random subset of 50 diagnostic conversations. Table 4 reports Fleiss’ kappa agreement between the medical experts and our automated evaluation pipeline. Agreement ranges from 0.88 to 0.94 for stage judgment (κ_J) and from 0.81 to 0.91 for semantic matching (κ_S).

5 Conclusion

We presented MMedAgentBench, a comprehensive benchmark designed to evaluate LLM-based medical agents in realistic multi-stage clinical workflows. Unlike existing medical benchmarks that focus on static, single-turn question answering, our framework captures the temporal complexity and resource constraints inherent in real clinical practice. Through a high-quality dataset comprising over 150 clinical cases and 500+ diagnostic stages, we enable systematic evaluation of medical agents

Model	Acc^{diag}	Acc^{treat}	Cost
Qwen2.5-14B	17.47	14.04	0.66
Qwen2.5-72B	21.76	17.77	0.85
Qwen3-14B	17.66	16.28	0.63
GPT-3.5	15.30	10.50	0.44
GPT-4o	22.13	18.68	0.74
GPT-5	35.62	38.33	0.93

(a) Gemini-2.5-Flash (Judge) + GPT-4o (Match).

Model	Acc^{diag}	Acc^{treat}	Cost
Qwen2.5-14B	18.16	15.00	0.67
Qwen2.5-72B	20.93	17.61	0.85
Qwen3-14B	17.59	16.05	0.63
GPT-3.5	15.15	8.95	0.45
GPT-4o	23.02	18.59	0.74
GPT-5	37.04	38.15	0.95

(b) GPT-5 (Judge) + Claude-Sonnet-4 (Match).

Model	Acc^{diag}	Acc^{treat}	Cost
Qwen2.5-14B	16.27	14.26	0.66
Qwen2.5-72B	20.15	17.22	0.84
Qwen3-14B	18.05	18.21	0.59
GPT-3.5	15.90	9.87	0.45
GPT-4o	22.75	18.94	0.75
GPT-5	35.12	38.57	0.93

(c) GPT-4o (Judge) + GPT-3.5 (Match).

Table 3: Sensitivity analysis across three evaluator configurations. Panel captions list the stage judge followed by the semantic matcher. Acc^{diag} and Acc^{treat} denote diagnostic and treatment accuracy, respectively (in %); Cost denotes the clinical examination cost ratio.

across dynamic decision-making, cost control, and multi-stage reasoning dimensions.

MMedAgentBench establishes a crucial testbed for advancing medical AI toward clinical deployment. By requiring agents to navigate sequential clinical stages while balancing accuracy and resource efficiency, our benchmark pushes the development of more sophisticated medical reasoning capabilities. Future work can leverage this framework to develop specialized medical agents that better handle temporal dependencies, cost constraints, and the inherent uncertainty in clinical workflows, ultimately contributing to more practical and deployable medical AI systems.

Model	κ_J	κ_S
Qwen2.5-14B	0.88	0.81
Qwen2.5-72B	0.91	0.85
GPT-4o	0.90	0.88
GPT-5	0.94	0.91

Table 4: Human-LLM agreement (Fleiss’ kappa) for selected models. κ_J : stage judgment; κ_S : semantic matching.

6 Limitations

Simplification of Clinical Complexity. Real clinical practice involves nuances that the benchmark cannot fully capture, including non-verbal patient cues, urgency considerations, and real-time collaborative decision-making among healthcare teams. We assume perfect information retrieval once an examination is ordered, whereas real-world scenarios often involve incomplete, ambiguous, or delayed results.

Static Ground Truth Assumption. Our benchmark assumes fixed ground truth diagnoses and treatments annotated by experts. However, medical knowledge evolves, and legitimate clinical disagreements exist even among specialists. The benchmark may not adequately capture this inherent uncertainty and variability in medical practice.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant Nos. 62501337, 72293585, and 72293580), the National Science and Technology Major Project of New Generation Artificial Intelligence (Grant No. 2025ZD0124104), and the Shanghai Municipal Special Program for Basic Research on General AI Foundation Models (Grant No. 2025SHZDZX025D09). This work was conducted in collaboration with Shanghai Artificial Intelligence Laboratory.

References

- Anthropic. 2025. Claude 4. <https://www.anthropic.com/claude>. Available at <https://www.anthropic.com/claude>.
- ByteDance Seed. 2025. Doubao-1.5-pro. https://seed.bytedance.com/en/special/doubao_1_5_pro. Accessed: 2026-01-06.

- Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Naveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsson, Idan Szpektor, Nan-Jiang Jiang, and 3416 others. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#).
- DeepSeek-AI. 2024. [Deepseek-v3 technical report](#). Preprint, arXiv:2412.19437.
- Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadao Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, and 39 others. 2024. [Chatglm: A family of large language models from glm-130b to glm-4 all tools](#). Preprint, arXiv:2406.12793.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). Preprint, arXiv:2407.21783.
- Kaiyuan Ji, Yijin Guo, Zicheng Zhang, Xiangyang Zhu, Yuan Tian, and Ning Liu. 2026. Medomni-45°: a safety–performance benchmark for reasoning-oriented llms in medicine. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 35536–35544.
- Yixing Jiang, Kameron C Black, Gloria Geng, Danny Park, James Zou, Andrew Y Ng, and Jonathan H Chen. 2025. Medagentbench: a virtual ehr environment to benchmark medical llm agents. *NEJM AI*, 2(9):AIdbp2500144.
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14):6421.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William W Cohen, and Xinghua Lu. 2019. Pubmedqa: A dataset for biomedical research question answering. *arXiv preprint arXiv:1909.06146*.
- Junkai Li, Yunghwei Lai, Weitao Li, Jingyi Ren, Meng Zhang, Xinhui Kang, Siyu Wang, Peng Li, Ya-Qin Zhang, Weizhi Ma, and Yang Liu. 2025. [Agent hospital: A simulacrum of hospital with evolvable medical agents](#).
- Meta AI. 2024. [Introducing meta llama 3: The most capable openly available llm to date](#). Accessed: 2025-07-28.
- OpenAI. 2023a. Chatgpt. <https://chat.openai.com/chat>.
- OpenAI. 2023b. Gpt-4 technical report. Technical report, OpenAI. Available at <https://arxiv.org/abs/2303.08774>.
- Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2022. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Conference on health, inference, and learning*, pages 248–260. PMLR.
- Anusri Pampari, Preethi Raghavan, Jennifer Liang, and Jian Peng. 2018. emrqa: A large corpus for question answering on electronic medical records. *arXiv preprint arXiv:1809.00732*.
- Qwen Team. 2024. [Qwen2.5: A party of foundation models!](#)
- Samuel Schmidgall, Rojin Ziaei, Carl Harris, Ji Woong Kim, Eduardo Pontes Reis, Jeffrey Jopling, and Michael Moor. 2026. [AgentClinic: a multimodal benchmark for tool-using clinical AI agents](#). *npj Digital Medicine*, 9(1):499.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, Perry Payne, Martin Seneviratne, Paul Gamble, Chris Kelly, Abubakr Babiker, Nathanael Schärli, Aakanksha Chowdhery, Philip Mansfield, Dina Demner-Fushman, and 13 others. 2023. [Large language models encode clinical knowledge](#). *Nature*, 620(7972):172–180.
- Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, Zhuofu Chen, Jialei Cui, Hao Ding, Mengnan Dong, Angang Du, Chenzhuang Du, Dikang Du, Yulun Du, Yu Fan, and 150 others. 2025. [Kimi k2: Open agentic intelligence](#). Preprint, arXiv:2507.20534.
- Tencent Hunyuan Team. 2025. [Hunyuan-a13b technical report](#). Technical report, Tencent. Accessed: 2025-01-06.
- Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. Large language models in medicine. *Nature medicine*, 29(8):1930–1940.
- Yuan Tian, Kaiyuan Ji, Rongzhao Zhang, Yankai Jiang, Chunyi Li, Xiaosong Wang, and Guangtao Zhai. 2025a. Towards all-in-one medical image re-identification. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 30774–30786.
- Yuan Tian, Xiaoyue Ling, Cong Geng, Qiang Hu, Guo Lu, and Guangtao Zhai. 2025b. Smc++: Masked learning of unsupervised video semantic compression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

- Yuan Tian, Guo Lu, Yichao Yan, Guangtao Zhai, Li Chen, and Zhiyong Gao. 2024. A coding framework and benchmark towards low-bitrate video understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(8):5852–5872.
- Yuan Tian, Shuo Wang, and Guangtao Zhai. 2025c. Medical manifestation-aware de-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 26363–26372.
- Yuan Tian, Shuo Wang, Rongzhao Zhang, Zijian Chen, Yankai Jiang, Chunyi Li, Xiangyang Zhu, Fang Yan, Qiang Hu, XiaoSong Wang, and Guangtao Zhai. 2025d. Semantics versus identity: a divide-and-conquer approach towards adjustable medical image de-identification. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20613–20625. IEEE.
- Yuan Tian, Min Zhou, Yitong Chen, Fang Li, Lingzi Qi, Shuo Wang, Xieyang Xu, Yu Yu, Shiqiong Xu, Chaoyu Lei, Yankai Jiang, Rongzhao Zhang, Jia Tan, Li Wu, Hong Chen, Xiaowei Liu, Wei Lu, Lin Li, Huifang Zhou, and 3 others. 2025e. Rofi: a deep learning-based ophthalmic sign-preserving and reversible patient face anonymizer. *NPI Digital Medicine*, 8(1):705.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. *Llama 2: Open foundation and fine-tuned chat models*.
- xAI. 2025. *Grok 4*. Online. Official announcement – Grok 4 model news.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. 2024. τ -bench: A benchmark for tool-agent-user interaction in real-world domains. *arXiv preprint arXiv:2406.12045*.
- Zicheng Zhang, Junying Wang, Yijin Guo, Farong Wen, Zijian Chen, Hanqing Wang, Wenzhe Li, Lu Sun, Yingjie Zhou, Jianbo Zhang, Bowen Yan, Ziheng Jia, Jiahao Xiao, Yuan Tian, Xiangyang Zhu, Kaiwei Zhang, Chunyi Li, Xiaohong Liu, Xiongkuo Min, and 2 others. 2025a. Aibench: Towards trustworthy evaluation under the 45°law. *Displays*, page 103255.
- Zicheng Zhang, Junying Wang, Farong Wen, Yijin Guo, Xiangyu Zhao, Xinyu Fang, Shengyuan Ding, Ziheng Jia, Jiahao Xiao, Ye Shen, Yushuo Zheng, Xiaorong Zhu, Yalun Wu, Ziheng Jiao, Wei Sun, Zijian Chen, Kaiwei Zhang, Kang Fu, Yuqin Cao, and 30 others. 2025b. Large multimodal models evaluation: A survey. *Science China Information Sciences*, 68(12):221301.
- Yuxuan Zhou, Xien Liu, Chen Ning, and Ji Wu. 2024. Multifaceteval: Multifaceted evaluation to probe llms in mastering medical knowledge. *arXiv preprint arXiv:2406.02919*.

A Appendix

A.1 Annotation Guidelines and Protocol

The creation of MMedAgentBench relied on a rigorous annotation protocol executed by board-certified physicians. These guidelines were provided to each annotator to ensure consistency, clinical accuracy, and relevance across all cases in the dataset. The objective was to establish a gold-standard ground truth for evaluating an AI agent’s multi-stage clinical reasoning.

A.1.1 General Principles

Annotators were instructed to approach each case as a real-world clinical scenario. The primary goal was to define what a competent, cost-conscious physician would conclude and do at each stage of the patient’s journey, based only on the information available at that point.

A.1.2 Per-Stage Annotation Form

Detailed instructions for the three per-stage annotation components summarized in Section 3.2 are provided below:

Part 1: Clinical Key Points.

This section captures the essential diagnostic and therapeutic conclusions for the stage. Annotators were asked to provide a list of short, distinct points.

Diagnostic Key Points (P_i^{diag}):

- **Objective:** List the most critical diagnoses, differential diagnoses, or significant findings that should be identified at this stage.
- **Instructions:**
 1. Focus on the most probable diagnoses and critical “can’t-miss” diagnoses (e.g., myocardial infarction, pulmonary embolism).
 2. Points should be justifiable based on the evidence presented up to stage s_i .
 3. Avoid vague statements. Instead of “abnormal lab results,” specify “leukocytosis with a left shift suggesting infection.”

Treatment Key Points (P_i^{treat}):

- **Objective:** List the essential treatment actions, management plans, or further consultations that should be initiated.
- **Instructions:**
 1. List specific actions (e.g., "Administer intravenous antibiotics," "Consult cardiology").
 2. Include key monitoring plans (e.g., "Monitor vital signs every 4 hours").
 3. The treatment plan should directly address the diagnostic findings.

Part 2: Reference Cost ($Cost_i^{ref}$).

This section establishes a cost baseline for efficient and effective clinical investigation.

- **Objective:** Identify a set of examinations that represents a reasonable and cost-effective diagnostic strategy for the stage.
- **Instructions:**
 1. From the list of all available examinations at stage s_i , select only those that are **essential** to confirm or rule out the primary differential diagnoses.
 2. Balance diagnostic yield against cost. Avoid ordering overly broad or low-yield tests unless clinically crucial.
 3. The final $Cost_i^{ref}$ is the sum of the prices of the selected essential examinations. Patient record items have a cost of zero.

Part 3: Stage Progression Criterion.

This defines the minimum requirements for an agent's performance to be considered successful for the current stage, allowing it to proceed to the next.

- **Objective:** Define the primary clinical goal of the current stage.
- **Instructions:**
 1. Answer the question: "What is the minimum level of understanding or action required at this stage to ensure patient safety and effective progression of care?"
 2. This is a holistic judgment. For example, the criterion might be "The agent must recognize the signs of sepsis and initiate fluid resuscitation and antibiotics" or "The agent must correctly interpret the

MRI as highly suspicious for malignancy and recommend a biopsy."

3. An agent failing to meet this criterion would not proceed, simulating a critical clinical error.

A.2 Recruitment and Payment

All data annotation for MMedAgentBench was conducted manually by a team of trained medical annotators with relevant clinical backgrounds. Each annotation underwent a multi-stage verification pipeline, including internal cross-checking and random sampling review, to minimize potential bias and human error. A subset of the annotated data was further proofread and refined by the authors to ensure accuracy and adherence to medical logic. Reasonable compensation was provided to all annotators according to their workload and contribution, ensuring fairness and quality throughout the entire annotation process.

A.3 Additional Data Governance Details

Medical data processing requires attention to protecting patient identity while retaining clinically relevant information. Prior work has examined re-identification and information preservation under de-identification (Tian et al., 2025c,a,e,d). Before annotation, direct identifiers and sensitive metadata were removed from the JSON records, and patient IDs were replaced with one-way cryptographic hashes used to retain record linkage across clinical stages. Only de-identified records entered the workflow. Medical experts reviewed the field-handling procedure and safeguards.

A.4 Dataset Statistics

Table 5 summarizes the distribution of patients across clinical departments in our dataset. The 167 cases span six major departments, covering a diverse range of conditions with varying diagnostic complexity and multi-stage reasoning requirements.

A.5 Prompt Design

Figure 5 presents the system prompts for our diagnostic reasoning framework. The Controller prompt guides the agent in iteratively selecting relevant medical information from a structured catalog, enforcing strict title matching and cost-effective retrieval. The Finalize prompt instructs the model to synthesize collected evidence into structured diagnostic conclusions.

Table 5: Distribution of patients across clinical departments.

Department	Patients	Percentage
Internal Medicine	42	25.1%
Surgery	35	21.0%
Cardiology	28	16.8%
Neurology	24	14.4%
Oncology	21	12.6%
Emergency Medicine	17	10.2%

Figure 6 shows the evaluation prompts. The Stage-Level Judge assesses whether candidate answers align with reference answers derived from clinical guidelines, allowing semantic equivalence rather than exact matching. The Semantic Point

Matching prompt enables fine-grained evaluation of whether specific clinical points are adequately covered.

B LLM Usage

LLM was used for drafting and language polishing. Except these uses, no AI tools were involved in study design, core experiments, result analysis, or interpretation. The authors are fully responsible for the accuracy of the work.

Controller System Prompt

You are a medical reasoning assistant. The system provides:

- Current stage available catalog;
- Historical known titles list (items selected in previous stages, used for retrieving historical information).

Task: Based on the patient's condition and cost-effectiveness, select one precise item title from the 'Current Stage Catalog', or finalize the current stage.

Available Actions (output strictly in JSON format):

```
{"action": "pick_report"|"pick_patient"|"get_history"|"finalize",
  "title": "exact title from catalog",
  "reason": "brief rationale"}
```

Rules:

1. You may ONLY select from titles in the 'Current Stage Catalog' (must be exact title match); do NOT invent items outside the catalog.
2. Historical known titles are ONLY for retrieving historical information.
3. Avoid repeatedly selecting the same item; when additional information provides limited value, select finalize.

Finalize System Prompt

You are a clinical physician. Based on the provided evidence, generate the conclusion for the current stage in strict JSON format:

```
{"diagnosis": "...",
  "treatment": "...",
  "used_evidence_titles": [...],
  "notes": "..."}

```

Requirements:

1. used_evidence_titles must ALL come from the provided evidence titles.
2. Diagnosis and treatment should be concise and definitive.
3. If evidence is insufficient or multiple retrievals return empty results, state in notes: 'No additional relevant records/reports found within the current scope.'

Figure 5: System prompts for the Controller and Finalize modules in our diagnostic reasoning framework.




Figure 6: Evaluation prompts for stage-level judgment and semantic point matching used in our automated assessment pipeline. Both prompts are grounded in reference clinical guidelines to ensure evaluation validity.