



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

MedAgentSim: Self-Evolving Multi-Agent Simulations for Realistic Clinical Interactions

Mohammad Almansoori^{*,†}, Komal Kumar^{*}, and Hisham Cholakkal[Ⓜ]

Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, UAE
{mohammad.almansoori, komal.kumar, hisham.cholakkal}@mbzuai.ac.ae

^{*}Equal contribution [†]Corresponding author

Abstract. In this work, we introduce MedAgentSim, an open-source simulated clinical environment with doctor, patient, and measurement agents designed to evaluate and enhance LLM performance in dynamic diagnostic settings. Unlike prior approaches, our framework requires doctor agents to actively engage with patients through multi-turn conversations, requesting relevant medical examinations (e.g., temperature, blood pressure, ECG) and imaging results (e.g., MRI, X-ray) from a measurement agent to mimic the real-world diagnostic process. Additionally, we incorporate self improvement mechanisms that allow models to iteratively refine their diagnostic strategies. We enhance LLM performance in our simulated setting by integrating multi-agent discussions, chain-of-thought reasoning, and experience-based knowledge retrieval, facilitating progressive learning as doctor agents interact with more patients. We also introduce an evaluation benchmark for assessing the LLM’s ability to engage in dynamic, context-aware diagnostic interactions. While MedAgentSim is fully automated, it also supports a user-controlled mode, enabling human interaction with either the doctor or patient agent. Comprehensive evaluations in various simulated diagnostic scenarios demonstrate the effectiveness of our approach. Our codebase, simulation environment, and benchmark datasets are publicly available on the project page: <https://medagentsim.netlify.app/>.

Keywords: Multi Agents · Visual Agents · Self Improving Agents.

1 Introduction

Advancements in Large Language Models (LLMs) and Vision-Language Models (VLMs) have shown promising capabilities across various medical tasks, achieving human-level performance on several medical benchmarks [17]. These models have demonstrated the ability to encode clinical knowledge [25,29], retrieve relevant medical literature [33], and achieve high accuracy in single-turn medical question-answering tasks [2,12,17,32]. However, current medical LLM assessments often rely on static evaluation benchmarks, where models are provided with complete patient information and tasked with answering predefined questions, sometimes with multiple-choice options [8]. These assessments often fail to capture the complexity of real-world doctor-patient interactions, where diagnosis

is not a single-step process but a dynamic, multi-turn dialogue. Such multi-turn doctor-patient interactions are important in clinical scenarios, as patients often struggle to describe their symptoms accurately due to limited medical knowledge, ambiguous perceptions, or communication barriers [16]. Consequently, physicians play an active role in structuring these interactions, posing clarifying questions, and refining their assessments as new information emerges [35].

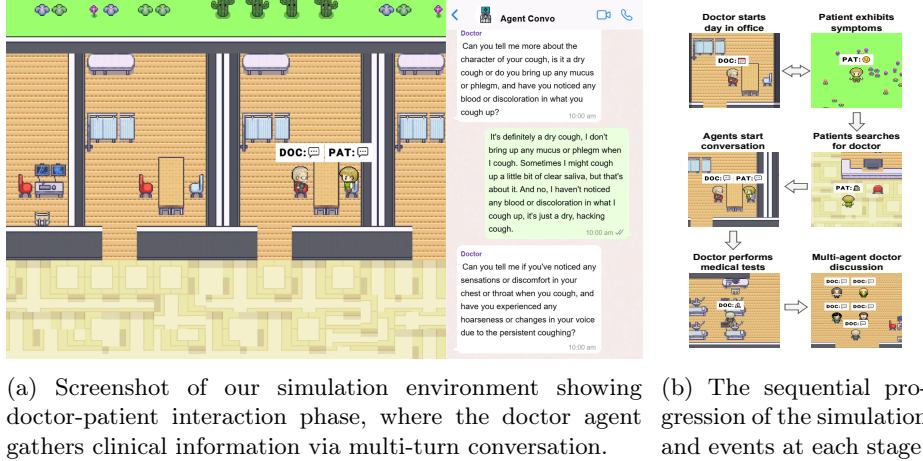


Fig. 1: Interactive clinical simulations in our MedAgentSim (best viewed when zoomed in).

Despite the aforementioned clinical significance, recent studies have highlighted that LLMs struggle in realistic clinical scenarios where they are not provided with all relevant information upfront [6,24]. Instead, they [6,24] shared only limited initial knowledge about the patient to the LLM and the LLMs are required to engage in a dynamic diagnostic process, systematically refining their understanding through patient dialogue. However, approaches such as AI Hospital [6] only introduced evaluation benchmarks, without enhancing LLMs for multi-turn interactions. Additionally, they relied on chat-based textual interaction simulations, where LLMs were not required to navigate complex environments or interact with medical tools.

Recently, LLM-driven game simulations were introduced in [11] for clinical settings, where closed-source AI agents based on OpenAI GPT-4o [20] were assigned roles such as doctors and patients [11]. These simulations were effective in capturing several aspects of real-world clinical complexity by requiring agents to navigate environments, interact with objects, and engage dynamically in decision-making. Additionally, these studies [4,11] incorporated memory-replay techniques to enhance agent performance. However, these approaches deviate from real-world clinical practice, as doctor agents are provided with a *pre-compiled, complete medical report of the patient*, rather than doctor agents

actively gathering patient information through interactive consultations. Furthermore, these simulations lack the ability to incorporate medical image-based diagnostic resources such as X-Rays and CT scans, which are critical in real medical decision-making. In addition to relying on closed-source LLMs like GPT-4o, many of these systems remain closed-source, limiting access to their data, code, and models, which hinders reproducibility and further research.

To address the limitations of existing methods, *we introduce MedAgentSim, an open-source, simulated hospital environment* designed to *evaluate and enhance* LLM performance in dynamic diagnostic settings. Unlike prior approaches, our framework, illustrated in Figure 1a, requires doctor agents to *actively engage with patients through multi-turn conversations, prompting medical examinations* to capture vital signs such as temperature, blood pressure, and electrocardiogram (ECG), *and requesting imaging results* (e.g., MRI, X-Ray) prior to making a diagnosis. Furthermore, we incorporate *self-improvement mechanisms*, allowing the models to iteratively refine their diagnostic strategies over time. We also *introduce an evaluation benchmark* designed to bridge the gap between static evaluations and real-world medical reasoning by assessing the LLM agent’s ability to engage in dynamic, context-aware diagnostic interactions, bringing it one step closer to practical clinical applications.

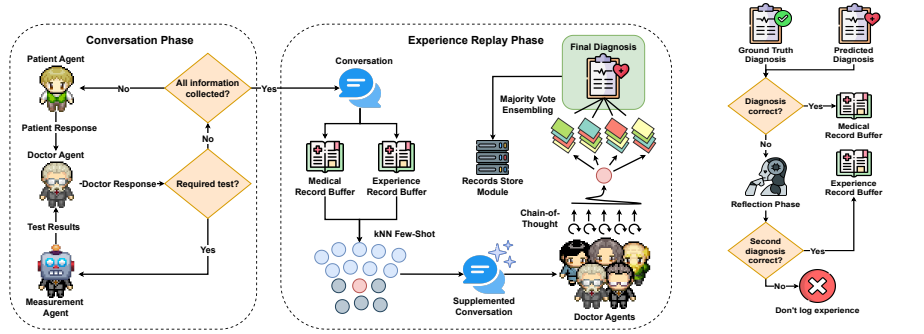
The **key contributions** of our method are summarized as below:

1. A game-based hospital simulation built with open-source LLMs [15,26], where LLM-powered doctor and patient agents interact in a realistic diagnostic setting. The system is fully automated and it also supports a user-controlled mode, allowing a human to take control of either the doctor or patient agent for real-time interaction with the AI counterpart.
2. A multi-agent LLM framework for realistic doctor-patient dialogue, where the doctor starts with no prior knowledge of the patient’s condition and need to ask questions for gathering relevant patient information. Test results are only provided if the doctor specifically requests the necessary tests, ensuring a process that closely mirrors real-world clinical consultations.
3. A multi-agent diagnostic pipeline that improves baseline LLM performance by incorporating self-improvement mechanisms, including multi-agent discussion, chain-of-thought (COT) [31] reasoning, and experience-based knowledge retrieval. The system enables progressive learning, where doctor agents refine their diagnostic capabilities as they interact with more patients.

2 Methodology: MedAgentSim

Figure 2 shows an overview of the proposed MedAgentSim comprising two key phases. At first, in the *Conversation Phase*, agents actively gather all relevant patient information necessary for diagnosis. Then, in the *Experience Replay Phase*, correctly diagnosed cases are stored as memory for future retrieval and learning. Next, we introduce our overall simulation architecture.

Simulation Environment. The proposed hospital simulation environment builds upon Generative Agents [21], transforming it into an interactive healthcare set-



(a) In *Conversation Phase*, the doctor and patient agents engage in an interactive dialogue, allowing the doctor to gather vital information and request necessary diagnostic tests, such as blood tests and X-Rays, from the measurement agent. As results are provided, the conversation continues until the doctor has sufficient information. Once ready to diagnose, the process transitions to *Experience Replay Phase*. Here, past doctor-patient interactions are analyzed through memory buffers, retrieving relevant cases as few-shot examples to enrich the current dialogue. A team of doctor agents then evaluates this enhanced conversation using *COT reasoning* and *majority-vote ensembling* to reach a consensus, producing a well-informed diagnosis.

(b) The record storing module progressively maintains a *medical records buffer* for storing correct diagnoses and an *experience records buffer* for tracking cases where initial misdiagnoses were later corrected upon reflection.

Fig. 2: (a) Overview of the proposed MedAgentSim comprising *Conversation* and *Experience Replay* phases. (b) Our records store module for progressive learning.

ting where autonomous virtual characters, commonly referred to as non-playable characters (NPCs), simulate real-world hospital dynamics. These NPCs, powered by an LLM, can move freely, initiate conversations, and interact with medical equipment, making real-time decisions based on the unfolding scenario.

(a) Agent Roles. The simulation consists of three core agent types: the *patient agent*, the *doctor agent*, and the *measurement agent*. The patient experiences symptoms and seeks medical attention from the doctor, who is responsible for diagnosing and treating conditions. The measurement agent provides diagnostic test results but only when explicitly requested, requiring the doctor to actively gather information rather than receiving all patient data upfront. Figure 1b showcases a sample scenario, demonstrating how agents navigate the environment and engage in clinical workflows. This baseline framework is referred to as Multi-Agent Clinic [24].

(b) Agent Interaction Modes. Both the doctor and patient agents can function in one of three distinct modes, determining how they generate and process information during interactions. In *Generation Mode*, the patient agent autonomously creates a case, generating illnesses, symptoms, and test results, which are internally stored. The doctor agent must actively extract relevant de-

tails through questioning. In *Dataset Mode*, patient responses are derived from a predefined dataset, ensuring consistency with structured medical knowledge, while the doctor agent follows the same interactive probing process. Finally, in *Control Mode*, a human user can assume control of either the doctor or patient, enabling real-time interactions with the AI-driven counterpart. This mode facilitates testing and supports potential real-world deployment, where real patients could engage with an AI-powered doctor or vice-versa.

Memory and Self-Improvement. Doctor-patient consultations take place through natural language interactions, where the doctor questions the patient, infers possible conditions, and orders tests. If a medical test is not requested, its results remain unavailable, mirroring real-world diagnostic constraints. Once the doctor is ready to make a diagnosis, the conversation undergoes a experience replay phase, refining the model’s decision-making over time.

(a) Records Buffer. To enable progressive learning, the system maintains a record storage and retrieval mechanism that captures both successful and corrected diagnoses. It consists of two dynamically expanding libraries: the *Medical Records Buffer*, which stores correctly diagnosed cases, and the *Experience Records Buffer*, which retains misdiagnosed cases that were later corrected through reflection. During a new consultation, the system uses k-nearest neighbors (KNN) to retrieve relevant past cases. The Medical Records Buffer provides full conversations and diagnoses, while the Experience Records Buffer extracts key insights from the reflection process. This approach builds on prior experiences, as studies show that LLMs benefit from failure-driven learning [34].

(b) COT and Ensembling. The retrieved information is then incorporated into the consultation, enriching the doctor’s contextual understanding. A multi-agent system processes the updated input, where multiple doctor agents independently assess the case and propose diagnoses. These assessments are aggregated and refined using COT [31] reasoning and majority-vote ensembling [17], producing a final diagnosis.

(c) Records Storage. Once finalized, the system converts all case data, including conversation history, diagnosis, medical images, and lab results, into CLIP [23] embeddings. Correct diagnosis embeddings are added to the Medical Records Buffer, while incorrect cases undergo a reflection phase where the doctor analyzes the mistake before making a second attempt. If the revised diagnosis is correct, only the CLIP-embedded reflection insights are stored in the Experience Records Buffer; otherwise, the case is discarded to ensure learning is based on meaningful examples. Figure 2b illustrates the full reflection and storage process.

3 Experiments

Experimental Details. We conducted extensive experiments to evaluate the effectiveness of MedAgentSim in a real-world doctor-agent setting. Our study leveraged a diverse set of both open-source models available on Hugging Face [3] and proprietary models, tested across three primary benchmarks: NEJM [24], MedQA [8], and MIMIC-IV [9]. For VLM tasks, we utilized the NEJM dataset,

which includes 15 complex real-world cases along with an extended set, NEJM Extended, of 120 additional cases. MedQA comprises 106 simulated diagnostic scenarios, while its extended variant, MedQA Extended, contains 214 distinct cases. Additionally, MIMIC-IV features 288 clinical cases, providing a diverse set of real-world medical interactions. As these datasets are primarily formatted for QA tasks, they are not directly compatible with our simulation pipeline. To address this, we preprocess the data using GPT-4o, converting it into a structured JSON format, where the doctor, patient information, and test results are assigned to the doctor agent, patient agent, and measurement agent, respectively. Model accuracy is evaluated using a *binary true/false metric* for the final diagnosis, with an LLM serving as the evaluator to account for variability in generated responses. Both the dataset conversion process and accuracy logs were manually reviewed to ensure reliability.

All models were deployed using vLLM [10] on a 4×48 GB NVIDIA RTX A6000 setup. For vision-language tasks, we integrated QwenVL [27], for the Qwen family of models, and LLaVA 1.5 [14] for the remaining models, with LLaVA demonstrating strong performance in medical image interpretation, particularly in generating descriptive reports for X-Rays, MRIs, and other imaging modalities. The visual game simulation was developed using Phaser, a web-based game engine [22], with the map designed in Tiled, a 2D level editor [13]. Game assets were sourced from Generative Agents [21].

Results and Analysis. Table 1 compares the performance of the baseline Multi-Agent Clinic and our proposed MedAgentSim across key medical benchmarks, covering both language-based and vision-based tasks. MedAgentSim integrates LLaVA 1.5-Mistral, a multi-modal model combining visual encoding with large language models.

The results show that MedAgentSim significantly outperforms the baseline across all benchmarks, particularly in multi-modal tasks. In the NEJM benchmark, MedAgentSim achieves 26.7% with LLaMA 3.3, a substantial improvement over the baseline Multi-Agent Clinic, where models struggle to exceed 20.0%. This gap widens in NEJM Extended, where MedAgentSim reaches 28.3% with LLaMA 3.3, surpassing the best baseline performance of 24.2%. These findings indicate that MedAgentSim is better equipped to interpret medical images and generate accurate clinical insights.

For language-based reasoning, MedAgentSim consistently demonstrates superior performance. In MedQA, it achieves 70.8% with LLaMA 3.3, while the best-performing baseline model records 62.3%. Similarly, in MedQA Extended, MedAgentSim attains 72.0%, a notable increase over the 63.6% baseline. The most significant performance boost is observed in MIMIC-IV, where MedAgentSim reaches 79.5%, far exceeding the highest baseline score of 42.7%.

In addition to automated benchmarks, a preliminary human study was performed in which one agent (doctor, patient, or measurement) was replaced with a real human. Doctors viewing the clips failed to identify the human in 62.5% of cases, indicating high behavioral realism. Full results are on the project page¹.

¹ <https://medagentsim.netlify.app>

Table 1: Performance of Multi-Agent Clinic (Basic) and MedAgentSim (Our) models across medical benchmarks. We used diverse LLMs including closed source. For visual language tasks, we use LLava 1.5 [14] for visual encoding.

Baseline	Size/Type	NEJM	NEJM Ext.	MedQA	MedQA Ext.	MIMIC-IV
Multi-Agent Clinic						
Claude [1]	3.5	—	—	62.3	63.6	42.7
ChatGPT [20]	4o	26.7	25.8	52.8	52.3	34.4
ChatGPT [19]	4	13.3	19.2	35.8	33.2	24.7
ChatGPT [18]	3.5	—	—	36.8	34.6	27.8
LLaMA 3.3 [15]	70B	20.0	24.2	54.7	53.3	36.8
LLaMA 3 [5]	70B	6.7	5.0	19.8	17.3	13.9
LLaMA 2 [28]	70B	—	—	4.7	2.8	8.3
Mixtral [7]	8×7B	6.7	2.5	37.7	39.3	30.2
Mistral [26]	24B	6.7	3.3	45.3	41.1	21.9
Qwen2 [30]	VL-7B	0.0	1.7	20.8	16.8	25.7
Qwen2.5 [27]	72B	0.0	2.5	38.7	41.6	21.2
MedAgentSim (Ours)						
ChatGPT [20]	4o	26.7	27.5	66.0	67.8	75.3
LLaMA 3.3 [15]	70B	26.7	28.3	70.8	72.0	79.5
Mistral [26]	24B	13.3	9.2	53.8	49.5	56.6
Qwen2 [30]	VL-7B	6.7	4.2	31.3	29.2	38.2
Qwen2.5 [27]	72B	6.7	4.2	55.7	57.5	66.0

3.1 Ablation Study

Impact of MedAgentSim Strategies. Table 2 summarizes the impact of adding incremental reasoning strategies on model accuracy. The integration of measurement, memory augmentation, COT [31] reasoning, and ensembling progressively improves diagnostic performance. Notably, the LLaMa 3.3 70B model benefits significantly from memory and COT strategies, achieving a final accuracy boost of 16.1%.

Table 2: Incremental improvements in model accuracy as measurement, memory, COT [31], and ensembling techniques are added.

Mistral 24B	Accuracy	LLaMa 3.3 70B	Accuracy
Baseline	45.3%	Baseline	54.7%
+ Measurement	47.2%	+ Measurement	59.4%
+ Memory	51.9%	+ Memory	65.1%
+ COT	52.8%	+ COT	68.9%
+ Ensembling	53.8%	+ Ensembling	70.8%

Model Sensitivity and Bias Reduction. The effectiveness of these strategies in mitigating bias is visualized in Figure 3. The left subfigure quantifies the baseline model’s susceptibility to biases, measured as accuracy fluctuations

across different diagnostic categories. The right subfigure highlights the stabilization effect of enhanced reasoning strategies, which reduce variance and improve robustness across bias types.

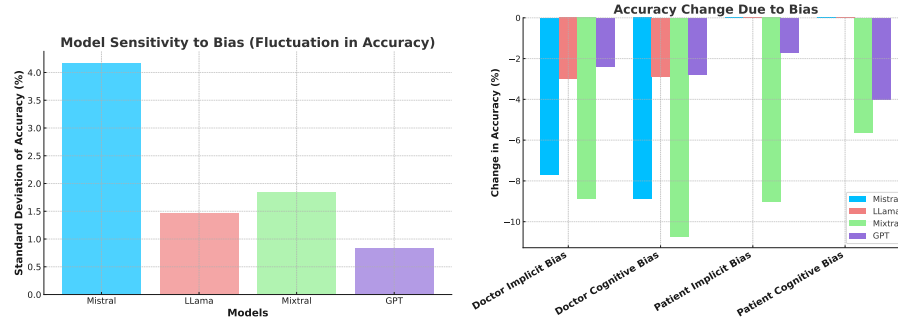


Fig. 3: The left figure shows the initial bias distribution, while the right figure illustrates bias reduction after incorporating additional features.

Impact of biases in the diagnosis. To further examine the role of biases in diagnostic accuracy, we present a radar plot, illustrated in Figure 4, comparing model performance under different cognitive and implicit bias conditions. The results indicate that Mixtral [7] and Mistral [26] exhibit greater susceptibility to patient cognitive biases, whereas LLaMa and GPT demonstrate higher stability.

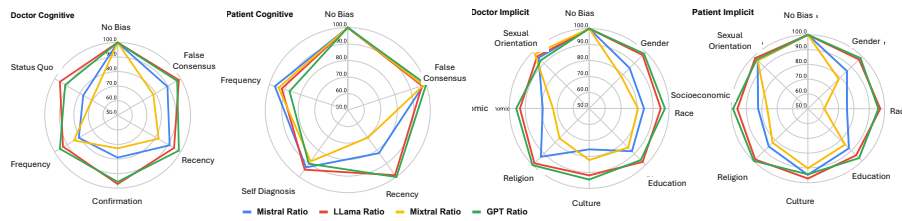


Fig. 4: Impact of Cognitive and Implicit Biases on Model Accuracy. This radar plot visualizes the accuracy variations of different models under various bias conditions. Larger deviations from the center indicate greater robustness to biases, while more compact shapes suggest higher sensitivity.

4 Conclusion

We introduced MedAgentSim, a multi-agent framework for interactive doctor-patient simulations that enhances diagnostic accuracy through structured reasoning, measurement-based decision-making, and self-improvement mechanisms. Our results demonstrate that memory, COT prompting, and ensembling significantly improve performance in realistic clinical scenarios. Additionally, our bias analysis highlights disparities in model robustness, emphasizing the need for fairness-aware AI in clinical applications. By bridging the gap between static benchmarks and real-world diagnostic reasoning, MedAgentSim provides a more adaptive approach to AI-driven healthcare.

Acknowledgments. We gratefully acknowledge support for this work from the Meta Llama Impact Innovation Award, the Meta Regional Research grant (Project OMER), the Google research award, the NVIDIA Academic grant, and the MBZUAI-WIS research grant (P008).

Disclosure of Interests. The authors declare no competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

1. Anthropic: Introducing claude 3.5 sonnet (2024), available at <https://www.anthropic.com/news/claude-3-5-sonnet>
2. Chen, Z., Cano, A.H., Romanou, A., Bonnet, A., Matoba, K., Salvi, F., Pagliardini, M., Fan, S., Köpf, A., Mohtashami, A., et al.: Meditron-70b: Scaling medical pretraining for large language models. arXiv preprint arXiv:2311.16079 (2023)
3. Community, H.F.: Huggingface (2024), <https://huggingface.co/models>
4. Du, Z., Zheng, L., Hu, R., Xu, Y., Li, X., Sun, Y., Chen, W., Wu, J., Cai, H., Ying, H.: Lms can simulate standardized patients via agent coevolution. arXiv preprint arXiv:2412.11716 (2024)
5. Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al.: The llama 3 herd of models. arXiv preprint arXiv:2407.21783 (2024), available at <https://arxiv.org/abs/2407.21783>
6. Fan, Z., Wei, L., Tang, J., Chen, W., Siyuan, W., Wei, Z., Huang, F.: Ai hospital: Benchmarking large language models in a multi-agent medical interaction simulator. In: Proceedings of the 31st International Conference on Computational Linguistics. pp. 10183–10213 (2025)
7. Jiang, A.Q., Sablayrolles, A., Roux, A., Mensch, A., Savary, B., Bamford, C., Chaplot, D.S., de las Casas, D., Hanna, E.B., Bressand, F., Lengyel, G., Bour, G., Lample, G., Lavaud, L.R., Saulnier, L., Lachaux, M.A., Stock, P., Subramanian, S., Yang, S., Antoniak, S., Scao, T.L., Gervet, T., Lavril, T., Wang, T., Lacroix, T., Sayed, W.E.: Mixtral of experts. arXiv preprint arXiv:2401.04088 (2024), available at <https://arxiv.org/abs/2401.04088>
8. Jin, D., Pan, E., Oufattole, N., Weng, W.H., Fang, H., Szolovits, P.: What disease does this patient have? a large-scale open domain question answering dataset from medical exams. Applied Sciences **11**(14), 6421 (2021)

9. Johnson, A.E., Bulgarelli, L., Shen, L., Gayles, A., Shammout, A., Horng, S., Pollard, T.J., Hao, S., Moody, B., Gow, B., et al.: Mimic-iv, a freely accessible electronic health record dataset. *Scientific data* **10**(1), 1 (2023)
10. Kwon, W., Li, Z., Zhuang, S., Sheng, Y., Zheng, L., Yu, C.H., Gonzalez, J.E., Zhang, H., Stoica, I.: Efficient memory management for large language model serving with pagedattention (2023), <https://arxiv.org/abs/2309.06180>
11. Li, J., Lai, Y., Li, W., Ren, J., Zhang, M., Kang, X., Wang, S., Li, P., Zhang, Y.Q., Ma, W., et al.: Agent hospital: A simulacrum of hospital with evolvable medical agents. *arXiv preprint arXiv:2405.02957* (2024)
12. Liévin, V., Hother, C.E., Motzfeldt, A.G., Winther, O.: Can large language models reason about medical questions? *Patterns* **5**(3) (2024)
13. Lindeijer, T.: A free and open source, easy to use, and flexible full-featured level editor. (2019), <https://www.mapeditor.org/>
14. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. *Advances in neural information processing systems* **36**, 34892–34916 (2023)
15. Meta: Llama 3.3-70b instruct (2024), available at <https://huggingface.co/meta-llama/Llama-3.3-70B-Instruct>
16. Meyer, A.N., Giardina, T.D., Khawaja, L., Singh, H.: Patient and clinician experiences of uncertainty in the diagnostic process: current understanding and future directions. *Patient Education and Counseling* **104**(11), 2606–2615 (2021)
17. Nori, H., Lee, Y.T., Zhang, S., Carignan, D., Edgar, R., Fusi, N., King, N., Larson, J., Li, Y., Liu, W., et al.: Can generalist foundation models outcompete special-purpose tuning? case study in medicine. *arXiv preprint arXiv:2311.16452* (2023)
18. OpenAI: Chatgpt-3.5 (2022), available at <https://openai.com/blog/chatgpt-3-5>
19. OpenAI: Chatgpt-4 (2023), available at <https://openai.com/blog/chatgpt-4>
20. OpenAI: Chatgpt-4o (2024), available at <https://openai.com/blog/chatgpt-4o>
21. Park, J.S., O’Brien, J., Cai, C.J., Morris, M.R., Liang, P., Bernstein, M.S.: Generative agents: Interactive simulacra of human behavior. In: *Proceedings of the 36th annual acm symposium on user interface software and technology*. pp. 1–22 (2023)
22. Phaser Studio, I.: A fast, fun and free open source html5 game framework (2018), <https://phaser.io/>
23. Radford, A., Kim, J.W., Hallacy, A., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. *Proceedings of the 38th International Conference on Machine Learning* (2021), available at <https://openai.com/research/clip>
24. Schmidgall, S., Ziaei, R., Harris, C., Reis, E., Jopling, J., Moor, M.: Agentclinic: a multimodal agent benchmark to evaluate ai in simulated clinical environments. *arXiv preprint arXiv:2405.07960* (2024)
25. Singhal, K., Azizi, S., Tu, T., Mahdavi, S.S., Wei, J., Chung, H.W., Scales, N., Tanwani, A., Cole-Lewis, H., Pfohl, S., et al.: Large language models encode clinical knowledge. *Nature* **620**(7972), 172–180 (2023)
26. Team, M.A.: Mistral small 3: A latency-optimized 24b-parameter model (2025), available at <https://mistral.ai/news/mistral-small-3>
27. Team, Q.: Qwen2.5-72b: A 72 billion parameter language model (2024), available at <https://huggingface.co/Qwen/Qwen2.5-72B>
28. Touvron, H., Martin, L., Stone, K.R., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023), available at <https://arxiv.org/abs/2307.09288>

29. Vaid, A., Landi, I., Nadkarni, G., Nabeel, I.: Using fine-tuned large language models to parse clinical notes in musculoskeletal pain disorders. *The Lancet Digital Health* **5**(12), e855–e858 (2023)
30. Wang, P., Bai, S., Tan, S., Wang, S., Fan, Z., Bai, J., Chen, K., Liu, X., Wang, J., Ge, W., Fan, Y., Dang, K., Du, M., Ren, X., Men, R., Liu, D., Zhou, C., Zhou, J., Lin, J.: Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191* (2024), available at <https://arxiv.org/abs/2409.12191>
31. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* **35**, 24824–24837 (2022)
32. Wu, C., Lin, W., Zhang, X., Zhang, Y., Xie, W., Wang, Y.: Pmc-llama: toward building open-source language models for medicine. *Journal of the American Medical Informatics Association* p. ocae045 (2024)
33. Xiong, G., Jin, Q., Lu, Z., Zhang, A.: Benchmarking retrieval-augmented generation for medicine. *arXiv preprint arXiv:2402.13178* (2024)
34. Yang, Z., Li, P., Liu, Y.: Failures pave the way: Enhancing large language models through tuning-free rule accumulation. *arXiv preprint arXiv:2310.15746* (2023)
35. Zhong, C., Liao, K., Chen, W., Liu, Q., Peng, B., Huang, X., Peng, J., Wei, Z.: Hierarchical reinforcement learning for automatic disease diagnosis. *Bioinformatics* **38**(16), 3995–4001 (2022)