

---

Article

Not peer-reviewed version

---

# Towards Evaluating the Diagnostic Ability of LLMs

---

[Peter Sarvari](#) \* and [Zaid Al-fagih](#)

Posted Date: 7 March 2025

doi: [10.20944/preprints202409.0688.v4](https://doi.org/10.20944/preprints202409.0688.v4)

Keywords: generative AI; LLM; GPT-4; RAG; clinical medicine; diagnosis



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Towards Evaluating the Diagnostic Ability of LLMs

Peter Sarvari \* and Zaid Al-fagih

Rhazes AI, London, W1W 7LT

\* Correspondence: peter@rhazes.ai

**Abstract.** On average, one in ten patients die because of a diagnostic error and medical errors are the third largest cause of death in the US. While LLMs have been proposed to help doctors with diagnoses, no research results have been published on comparing the diagnostic ability of many popular LLMs on an openly accessible real-patient cohort. In this study, we compare LLMs from Google, OpenAI, Meta, Mistral, Cohere and Anthropic using a previously established evaluation methodology and explore improving their accuracy with RAG. We found that GPT-4o from OpenAI and Claude Sonnet 3.5 from Anthropic were the top performers with them only missing 0.5% of ground truth conditions that were clearly inferable from the available data; RAG further improved this error rate to 0.2%. While the results are promising, more diverse datasets, hospital pilots and close collaboration with physicians are needed to get a better understanding of the diagnostic ability of these models.

**Keywords:** generative AI; LLM; GPT-4; RAG; clinical medicine; diagnosis

## 1. Introduction

In the United States alone, medical errors are the third largest cause of death [1], and within that, diagnostic errors kill or permanently disable 800,000 people each year [2]. Research by The National Academy of Medicine as well as Newman-Toker et al. estimated that diagnostic errors are responsible for approximately 10% of patient deaths [3,4] and 6-17% of hospital complications [3]. 75% of diagnostic errors are cognitive errors [5] most commonly caused by premature closure, the failure to consider alternatives after an initial diagnosis has been established. Cognitive errors are also naturally linked to the overload and stress physicians have been experiencing with current burnout rates reaching the highest ever levels recorded [6]. Given the recent progress in Artificial Intelligence, large language models (LLMs) have been proposed to help with various aspects of clinical work, including diagnosis [7]. GPT-4, a LLM developed by OpenAI has shown promise in medical applications with its ability to pass medical board exams in multiple countries and languages [8–11].

Only a handful number of studies have attempted to compare the diagnostic ability of LLMs mostly on (1) clinical vignettes, (2) case records directly from clinics and (3) case reports, such as the New England Journal of Medicine (NEJM) Case Challenges. The latter are more complex than clinical vignettes and contain red herrings and other distractors to truly challenge a physician [12]. Khan et al. [12] used 10 case challenges, and compared diagnoses from GPT-3.5, GPT-4 (Bing) and Gemini 1.5 with the help of 10 physicians who filled out a grading rubric. Chiu et al. [13] used 102 case records from the Massachusetts General Hospital and showed that GPT-4 outperformed Bard and Claude 2 in its diagnostic accuracy based on the ICD-10 hierarchy. Shieh et al. [14] asked GPT-3.5 and GPT-4 to analyze 109 USMLE Step 2 clinical knowledge practice questions (vignettes) as well as 63 case reports from various journals. The researchers concluded that while GPT-4 was 87.2% accurate on the vignettes, it was only able to create a shortlist of differential diagnoses for 47 of the case reports (75%). Others scholars have assessed the capabilities of various LLMs within a given specialty, such as otolaryngology [15] and radiology [16].

Many authors focused on evaluating the diagnostic ability of a single LLM: GPT-4 was the most popular choice as it was generally the most accurate LLM at the time. Eriksen et al. [17] asked GPT-4



to choose one from 6 diagnostic options for each of the 38 NEJM case challenges whereas Kanjee et al. [18] relied on NEJM clinopathologic case conferences and tasked GPT-4 to first, state the most likely diagnosis and second, give a list of differentials. Manual review by the authors concluded that in 45 out of the 70 cases the correct answer was included in the differentials (in 27 cases it was the most likely diagnosis). Shea et al. [19] used GPT-4 to diagnose 6 patients with extensive investigations but delayed definitive diagnoses and showed that GPT-4 has the potential to outperform clinicians and alternative diagnostic tools such as Isabel DDx companion. Fabre et al. [20] took 10 NEJM cases and while concluded that the final diagnosis was correctly identified by the AI in 8 cases (it was included in the list of differentials), the investigators also assessed treatment suggestions and found that GPT-4 failed to suggest adequate treatment for 7 cases. Notably, some researchers focused on assessing agreement between doctors and GPT-4, rather than evaluating the accuracy directly: Hirosawa et al. [21,22], measured the Cohen's Kappa coefficient in two different studies, relying on cases from the American Journal of Case Reports as well as 52 complex case reports published by the authors. In both cases, the researchers found fair to good agreement (0.63 [21] and 0.86 [22], respectively) between doctors and GPT-4.

These evaluation strategies work for case challenges but would not suffice for a large cohort of highly comorbid real patients, such as MIMIV-IV [23], where patients might suffer from multiple conditions concurrently. To solve this, Sarvari et al. [24] outlined a methodology to use AI-Assisted evaluation ("LLM-as-a-judge [25]") to quickly estimate the diagnostic accuracy of different models on a set of highly comorbid real hospital patients. This automated evaluation not only allows for evaluating on larger datasets (we increased the sample size 10-fold from the <100 typically seen in evaluations based on clinical cases to 1000), but also facilitates quick benchmarking of multiple models, which is our goal in this study. Automated evaluation gave reliable estimates as judged by three medical doctors in the aforementioned study [24], and as AI models improve, we only expect this to become better. Moreno et al. [26] also hints at non-human evaluation as a method to allow for a larger-scale beta test and Zack et al. [27] actively employs this method to match generated diagnoses to ground truth ones and shares the prompt in the supplementary material.

Despite the recent successes, there are subdomains where GPT-4o is proven to be inferior to alternative AI methods, or human diagnosis, particularly when it comes to medical image analysis. GPT-4o was found to perform poorly in detecting pneumonia from pediatric chest X-rays compared to traditional CNN-based methods [28]. Zhang et al. [29] compared GPT-4o to 3 medical doctors in their abilities to diagnose 26 glaucoma cases and found using Likert scales that GPT-4o performed worse than the lowest scoring doctor in the completeness category. Cai et al. [30] assessed the clinical utility of GPT-4o in recognizing abnormal blood cell morphology, an important component of hematologic diagnostics in 70 images. The LLM achieved an accuracy of only 70% (compared to 95.42% accuracy of hematologists) as reviewed by two experts in the field.

## 2. Methods

### 2.1. Models

We compared the following models for diagnosis in our analysis: Gemini 1 (*gemini-pro-vision* via Google Vertex AI API used on 2024/03/26 with *temperature* set to zero), Gemini 1.5 (*Gemini-1.5-pro-preview-0409* via Google Vertex AI API used 2024/05/08, *temperature* set to zero), MedLM (*medlm-medium* via Google Vertex AI API used 2024/05/08, *temperature*: 0.2, *top\_p*: 0.8, *top\_k*: 40), LLaMA 3.1 (*Meta-Llama-3.1-405B-Instruct* deployed on Microsoft Azure used via API on 2024/08/22), Mistral 2 (*Mistral-large-2407* deployed on Microsoft Azure used via API on 2024/08/22), Command R Plus (*command-r-plus* via Cohere API used 2024/06/24), GPT-4-Turbo (*gpt-4-1106-preview* via OpenAI API used 2024/02/14), GPT-4o (*gpt-4o* via OpenAI API used 2024/08/29, *temperature* set to zero), Claude 3.5 Sonnet (*claude-3-5-sonnet-20240620* via Anthropic API used 2024/07/23, *temperature* set to zero). The automated evaluation was done by GPT-4-Turbo (*gpt-4-1106-preview* via OpenAI API, *temperature* set to zero) on the same day when the diagnostic models were run. Note: when parameters are not

mentioned, they were not explicitly set, and their default values have been used. The reported hit rate is the average across all the ground truth diagnoses of the 1000 sample patients.

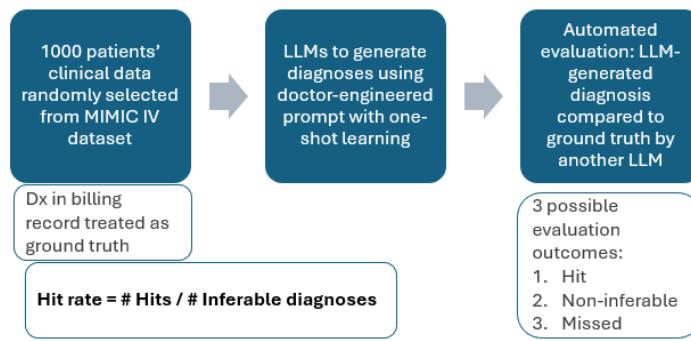
GPT-4o with retrieval augmented generation (RAG) was implemented via Azure Cognitive Search. A document containing Laboratory Test Reference Ranges from The American Board of Internal Medicine updated January 2024 [31] was vectorized (embedded by the *Ada-002* model from OpenAI) and indexed to be used for RAG with default overlap and chunk size (1024). The 5 closest matches were retrieved using the cosine similarity metric and the output was generated on 2024/06/26 with *temperature* of zero, *strictness* parameter of 3 and the *inScope* flag set to False.

## 2.2. Diagnosis and Automated Evaluation

The MIMIC-IV data sample containing 1000 hospital admissions and the diagnostic and evaluation prompts were taken from. The evaluation methodology is summarized in *Figure 1*.

Our initial idea was to simply compare the predicted ICD codes to the ICD codes extracted from the patients' billing reports (ground truth) and examine what proportion was guessed correctly. However, the MIMIC-IV data did not contain patient history (previous diagnoses, medications), patient physical examinations and other useful measurements such as ECG. Of course, without medication records, we would not know if the patient is suffering from a coagulation disorder or is taking anticoagulants and without ECG we cannot diagnose atrial fibrillation. Hence, such diagnoses are not inferable from the data, and we exclude them. Further, given the lack of patient diagnostic history and the very specific ICD code names, it may not be possible to distinguish between diseases with different onsets (acute vs chronic) or between diseases with differing degrees of severity. Hence, we deem the prediction correct if the predicted and the ground truth diagnoses are two related diseases (e.g., caused by the same pathogen, affecting the same organ) which are indistinguishable given the patient data. In this case, the further tests the LLM is instructed to suggest in the prompt from Sarvari et al. [24] are of crucial importance to understand to exact disease pathology. There are also ICD codes that do not correspond to diagnoses (e.g., Do Not Resuscitate, homelessness, unemployment) and we exclude such codes from this study. We define a correct prediction as a 'hit', and the failure to predict a ground truth diagnosis as a 'miss'.

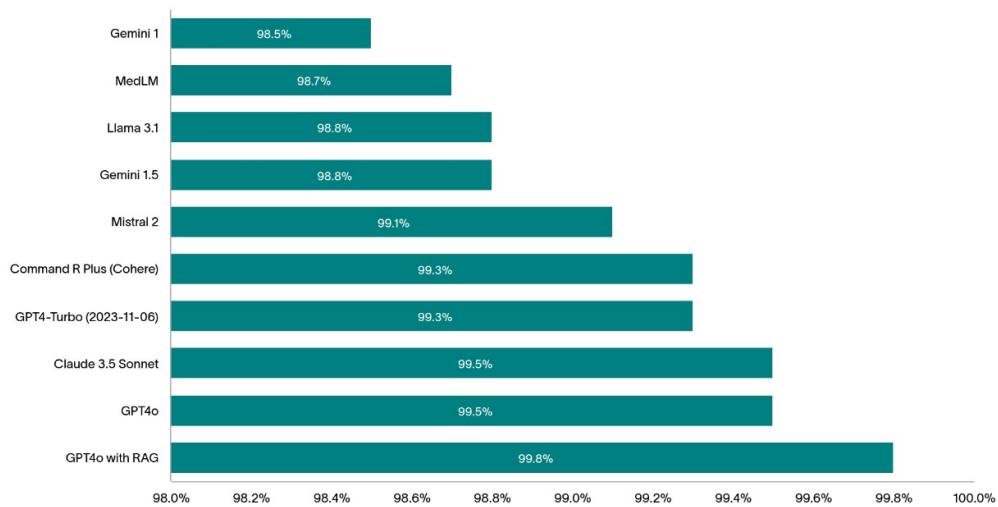
In terms of the evaluation metrics, we solely focus on hit rate (also called recall, true positive rate, sensitivity) in this study. The rationale is as follows: for every single disease in the world, the patient may have it or not have it. As such, when making predictions, the LLM is effectively executing binary classifications for every single disease. Of course, even a highly comorbid patient will not have 99.99%+ of the possible diseases and hence the metrics related to negative selected elements, such as specificity are very close to 1 by default and are not meaningful to report. As a result, the meaningful metrics here are precision and hit rate. However, a good quantification of precision is challenging in this case because false positives are difficult to establish as not every single medical condition ends up on the billing report of the patient. Hence, it's unclear and subjective whether certain well-reasoned diagnostic predictions should be marked false positives just because they did not show up on the patient's billing report. As a solution, we report the hit rate while (1) indirectly constraining the number of predictions by limiting the LLM output tokens to 4096 and (2) ensuring explainability by asking the LLM to reason why it predicted certain conditions.



**Figure 1.** Summary of the evaluation methodology.

### 3. Results

The 1000 randomly selected patients are highly comorbid with an average of 14.4 distinct diagnostic codes per patient (min:1, max: 39, IQR: 10). The bar chart on *Figure 2* shows the diagnostic hit rate of the models we tested in this study.



**Figure 2.** The diagnostic hit rates of the tested LLM.

The top-performing models (without RAG) assessed in this study were GPT-4o and Claude 3.5 Sonnet with both achieving 99.5% hit rate. *Table 1* summarizes the most common hits and misses by these two top-performing LLMs.

To boost the diagnostic hit rate of one of the best performing models, GPT-4o, we reduced the errors in the predictions by plugging the knowledge gaps in the model using RAG. The 1000 patient EHR contained a total of 14403 ground truth diagnoses. Among the 7604 inferable diagnoses, GPT-4o found the exact condition, or one deemed directly related to it (i.e., equally reasonable to infer given the patient data) in 7586 cases giving it a diagnostic hit rate (sensitivity) of 99.8%. The 7586 hits the model made correspond to 1733 unique diagnoses. In *Table 2*, we display the most common correctly identified diagnoses as well as the ones that were missed more than once.

**Table 1.** Most common diagnostic hits and misses made by the best performing models, GPT-4o and Claude 3.5 Sonnet.

Type	Name	Occurrence (GPT4o)	Occurrence (Claude)
Hit	Kidney failure, unspecified	217	216
Hit	Diabetes mellitus (no mention of complication, type II or unspecified)	132	137
Hit	Acidosis	129	126
Hit	Congestive heart failure, unspecified	125	128
Miss	Dehydration	8	5
Miss	Diabetes	8	3
Miss	Hypertension	4	2
Miss	Hypotension	2	5

**Table 2.** Most common diagnostic hits and misses made by GPT-4o with RAG.

Type	Name	Occurrence
Hit	Kidney failure, unspecified	218
Hit	Acidosis	129
Hit	Congestive heart failure, unspecified	127
Miss	Diabetes mellitus without complications	8
Miss	Hypoxemia	2

#### 4. Discussion

In this paper we compared the diagnostic ability of multiple large language models using a previously established method on a subset of the MIMIC-IV dataset. The method uses the ICD codes from the patient record as the ground truth and (1) removes not inferable diagnoses and (2) accepts similar ICD diagnoses as correct predictions when there's not enough data to infer the exact code.

Others have used ICD chapters [13] and 515 CCSR categories and 22 CCSR bodies [32] to compare the diagnostic predictions to the ground truth and reported accuracies at these different levels. While this method is very much helpful for creating a fast and objective evaluation framework, it does not consider if the data available is enough to arrive to the ground truth diagnosis (or to a similar one within the same CCSR category) resulting in a more conservative reported diagnostic accuracy. In other words, by using this method, one assumes that the information in the data used (MIMIC-III in the case of Mohammadi et al. [32]) is sufficient to make the reported ICD diagnoses. In addition, one major drawback of attempting to predict ICD chapters and CCSR categories is that two physiologically very different diseases may end up in the same category. For example, 'Type 1 diabetes mellitus without complications' (ICD-10 code: *E109*) and 'Type 2 diabetes mellitus without complications' (ICD-10 code *E119*) belong to the same CCSR category 1 of *END002*. This means that if the LLM predicted type 1 diabetes, but the patient was suffering from type 2 diabetes, the prediction would be deemed correct, even though in practice this would be a serious misdiagnosis. Ironically, closely related conditions may end up in different CCSR categories: 'chronic kidney disease, stage 1' (ICD-10 code *N181*) is in the *GEN003* CCSR category whereas 'Hypertensive chronic kidney disease with stage 1 through stage 4 chronic kidney disease, or unspecified chronic kidney disease' (ICD-10 code *I129*) is in the *CIR008* CCSR category. This means that we would penalize the LLM if it does not know that the chronic kidney disease was of hypertensive origin even if it does not have access to the

patient history proving so (note that patient blood pressure may appear normal in hypertensive kidney disease due to medication).

Our method uses a more subjective assessment, where we let the LLM agent conducting the evaluation decide whether the prediction is acceptable based on its similarity to the ground truth and given the available data. For example, mixing up type 1 and type 2 diabetes would be considered a miss if there is relevant antibody and C peptide data. At the very least, the model would suggest a further C peptide test (as instructed via the prompt in Sarvari et al. [24]) if not already in the data, to confirm the diagnosis. Another advantage of our approach is that it makes the reported hit rate less data dependent by removing the non-inferable diagnoses. However, in an ideal case, complete and detailed patient EHR data is available from multiple hospitals, locations and demographics to test the diagnostic ability of LLMs. While the hit rate of these LLMs on such dataset might be different, we would expect the relative rankings of these models to stay the same.

Throughout our analysis, we took care to report the exact dates when the experiments were conducted to account for potential silent model changes that have happened since. Note that due to the stochastic nature of the LLMs, the same model ran twice may give different results. However, repeating experiments (and the evaluation) have resulted in very similar results without a change to the first number after the decimal point. Hence, we chose the report hit rates with one decimal precision.

RAG is helpful for boosting hit rate as it allows the model to refer to up-to-date clinical reference ranges and diagnostic guidelines: whenever there's a blood result or symptom mentioned in the text, the LLM receives further information on the diagnostic guidelines resulting in fewer hallucinations [33] and more consistent output. For example, using RAG helped the GPT-4o model accurately diagnose dehydration which was the top missed diagnosis as reported in *Table 1*. We presume this is because of the serum and plasma osmolarity reference ranges in the document used for RAG. We expect fine-tuning these models with a medic-curated dataset will further increase their diagnostic abilities and this is something we're currently experimenting with.

Finally, we would like to draw attention to the shortcomings of this study: first, we only considered one single dataset, coming from a single hospital. This dataset didn't contain all information that doctors normally use for diagnosing patients, resulting in excluding some important diagnoses from the analysis as they were deemed non-inferable. In fact, in practice, decision making goes beyond text-based data from the electronic patient record and without an AI system taking multimodal inputs sitting alongside a doctor as part of a proper hospital pilot, it'll be very difficult to truly compare diagnostic ability of LLMs to that of doctors. In this study we allowed LLMs to make many predictions, however, in practice doctors may need to rely on one single diagnosis and treatment plan, which is their current best estimate. In addition, evaluation was done by an LLM and has not been reviewed manually by a human, let alone a clinician. Moreover, this paper did not assess model biases in the predictions made by the different models, which would be an essential first step towards hospital deployment of LLMs. Readers looking to learn more about this topic are directed to Zack et al. [27]. This paper also doesn't consider images and only takes natural language as an input; this is a crucial limitation especially as aforementioned studies indicated shortcomings of GPT-4o in medical image analysis [28–30]. It is also worth highlighting that while here we only tested the performance of LLMs in English language, recent research suggests consistent diagnostic performance of GPT-4o across 9 different languages [34].

## 5. Conclusions

In this study we compared the diagnostic ability of 9 different LLMs from 6 different companies on 1000 electronic patient records. We found that GPT-4o from OpenAI and Claude Sonnet 3.5 from Anthropic were the top performers with them only missing 0.5% of ground truth conditions that were clearly inferable from the available data. Open-source models, such as Mistral 2 and LLaMA 3.1 performed reasonably well, better than the closed-source models from Google, but worse than alternatives from Cohere, Anthropic and OpenAI. We showed how retrieval augmented generation

further improved the hit rate of GPT-4o and even though the numbers look very promising, we cautioned against drawing conclusions about the diagnostic abilities of these models in a real hospital setting.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

**Data Availability Statement.** The MIMIC-IV data is available to approved researchers on PhysioNet, and the SQL code used to transform this dataset is available at <https://github.com/sarvarip/MIMIC-SQL>

**Funding:** The authors have no funding to declare.

**Author contributions.** Conceptualization, Peter Sarvari; Methodology, Peter Sarvari and Zaid Al-fagih; Software, Peter Sarvari; Validation, Peter Sarvari and Zaid Al-fagih; Formal analysis, Peter Sarvari; Investigation, Peter Sarvari and Zaid Al-fagih; Resources, Peter Sarvari and Zaid Al-fagih; Data curation, Peter Sarvari and Zaid Al-fagih; Writing – original draft, Peter Sarvari; Writing – review & editing, Peter Sarvari and Zaid Al-fagih; Visualization, Zaid Al-fagih; Project administration, Peter Sarvari.

## References

1. Sameera, V., Bindra, A. & Rath, G. P. Human errors and their prevention in healthcare. *J Anaesthesiol Clin Pharmacol* **37**, 328–335 (2021).
2. Newman-Toker, D. E. et al. Burden of serious harms from diagnostic error in the USA. *BMJ Qual Saf* **33**, 109–120 (2024).
3. National Academies of Sciences. *Improving Diagnosis in Health Care*. (National Academies Press, Washington, D.C., 2015). doi:10.17226/21794.
4. Newman-Toker, D. E. et al. Rate of diagnostic errors and serious misdiagnosis-related harms for major vascular events, infections, and cancers: toward a national incidence estimate using the “Big Three”. *Diagnosis* **8**, 67–84 (2021).
5. Thammasitboon, S. & Cutrer, W. B. Diagnostic Decision-Making and Strategies to Improve Diagnosis. *Curr Probl Pediatr Adolesc Health Care* **43**, 232–241 (2013).
6. Wise, J. Burnout among trainees is at all time high, GMC survey shows. *BMJ* o1796 (2022) doi:10.1136/bmj.o1796.
7. Topol, E. J. Toward the eradication of medical diagnostic errors. *Science (1979)* **383**, (2024).
8. Madrid-García, A. et al. Harnessing ChatGPT and GPT-4 for evaluating the rheumatology questions of the Spanish access exam to specialized medical training. *Sci Rep* **13**, 22129 (2023).
9. Rosoł, M., Gąsior, J. S., Łaba, J., Korzeniewski, K. & Młyńczak, M. Evaluation of the performance of GPT-3.5 and GPT-4 on the Polish Medical Final Examination. *Sci Rep* **13**, 20512 (2023).
10. Brin, D. et al. Comparing ChatGPT and GPT-4 performance in USMLE soft skill assessments. *Sci Rep* **13**, 16492 (2023).
11. Kung, T. H. et al. Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models. *PLOS Digital Health* **2**, e0000198 (2023).
12. Khan, M. P. & O’Sullivan, E. D. A comparison of the diagnostic ability of large language models in challenging clinical cases. *Front Artif Intell* **7**, (2024).
13. Chiu, W. H. K. et al. Evaluating the Diagnostic Performance of Large Language Models on Complex Multimodal Medical Cases. *J Med Internet Res* **26**, e53724 (2024).
14. Shieh, A. et al. Assessing ChatGPT 4.0’s test performance and clinical diagnostic accuracy on USMLE STEP 2 CK and clinical case reports. *Sci Rep* **14**, 9330 (2024).
15. Warrier, A., Singh, R., Haleem, A., Zaki, H. & Eloy, J. A. The Comparative Diagnostic Capability of Large Language Models in Otolaryngology. *Laryngoscope* **134**, 3997–4002 (2024).
16. Sonoda, Y. et al. Diagnostic performances of GPT-4o, Claude 3 Opus, and Gemini 1.5 Pro in “Diagnosis Please” cases. *Jpn J Radiol* **42**, 1231–1235 (2024).
17. Eriksen, A. V., Möller, S. & Ryg, J. Use of GPT-4 to Diagnose Complex Clinical Cases. *NEJM AI* **1**, (2024).

18. Kanjee, Z., Crowe, B. & Rodman, A. Accuracy of a Generative Artificial Intelligence Model in a Complex Diagnostic Challenge. *JAMA* **330**, 78 (2023).
19. Shea, Y.-F., Lee, C. M. Y., Ip, W. C. T., Luk, D. W. A. & Wong, S. S. W. Use of GPT-4 to Analyze Medical Records of Patients With Extensive Investigations and Delayed Diagnosis. *JAMA Netw Open* **6**, e2325000 (2023).
20. Fabre, B. L. et al. Evaluating GPT-4 as an academic support tool for clinicians: a comparative analysis of case records from the literature. *ESMO Real World Data and Digital Oncology* **4**, 100042 (2024).
21. Hirosawa, T. et al. Evaluating ChatGPT-4's Accuracy in Identifying Final Diagnoses Within Differential Diagnoses Compared With Those of Physicians: Experimental Study for Diagnostic Cases. *JMIR Form Res* **8**, e59267 (2024).
22. Mizuta, K., Hirosawa, T., Harada, Y. & Shimizu, T. Can ChatGPT-4 evaluate whether a differential diagnosis list contains the correct diagnosis as accurately as a physician? *Diagnosis* **11**, 321–324 (2024).
23. Johnson, A. E. W. et al. MIMIC-IV, a freely accessible electronic health record dataset. *Sci Data* **10**, 1 (2023).
24. Sarvari, P., Al-fagih, Z., Ghuwel, A. & Al-fagih, O. A systematic evaluation of the performance of GPT-4 and PaLM2 to diagnose comorbidities in MIMIC-IV patients. *Health Care Science* **3**, 3–18 (2024).
25. Kahng, M. et al. LLM Comparator: Interactive Analysis of Side-by-Side Evaluation of Large Language Models. *IEEE Trans Vis Comput Graph* **31**, 503–513 (2025).
26. Moreno, A. C. & Bitterman, D. S. Toward Clinical-Grade Evaluation of Large Language Models. *International Journal of Radiation Oncology\*Biology\*Physics* **118**, 916–920 (2024).
27. Zack, T. et al. Assessing the potential of GPT-4 to perpetuate racial and gender biases in health care: a model evaluation study. *Lancet Digit Health* **6**, e12–e22 (2024).
28. Chetla, N. et al. Evaluating ChatGPT's Efficacy in Pediatric Pneumonia Detection From Chest X-Rays: Comparative Analysis of Specialized AI Models. *JMIR AI* **4**, e67621 (2025).
29. Zhang, J. et al. A comparative study of GPT-4o and human ophthalmologists in glaucoma diagnosis. *Sci Rep* **14**, 30385 (2024).
30. Cai, X., Zhan, L. & Lin, Y. Assessing the accuracy and clinical utility of GPT-4O in abnormal blood cell morphology recognition. *Digit Health* **10**, (2024).
31. American Board of Internal Medicine. ABIM Laboratory Test Reference Ranges - January 2024. <https://www.abim.org/Media/bfijryql/laboratory-reference-ranges.pdf> (2024).
32. Shah-Mohammadi, F. & Finkelstein, J. Accuracy Evaluation of GPT-Assisted Differential Diagnosis in Emergency Department. *Diagnostics* **14**, 1779 (2024).
33. Wang, C. et al. Potential for GPT Technology to Optimize Future Clinical Decision-Making Using Retrieval-Augmented Generation. *Ann Biomed Eng* **52**, 1115–1118 (2024).
34. Chimirri, L. et al. Consistent Performance of GPT-4o in Rare Disease Diagnosis Across Nine Languages and 4967 Cases. Preprint at <https://doi.org/10.1101/2025.02.26.25322769> (2025).

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.