



OPEN ACCESS

EDITED BY

Carsten Grötzing, Charité University Medicine Berlin, Germany

REVIEWED BY

Flurin Eggmann, University of Basel, Switzerland
Simeng Zhu, Henry Ford Hospital, United States

*CORRESPONDENCE

Florian Putz
✉ Florian.Putz@uk-erlangen.de

RECEIVED 21 July 2023

ACCEPTED 23 August 2023

PUBLISHED 14 September 2023

CITATION

Huang Y, Gomaa A, Semrau S, Haderlein M, Lettmaier S, Weissmann T, Grigo J, Tkhatay HB, Frey B, Gaipf U, Distel L, Maier A, Fietkau R, Bert C and Putz F (2023) Benchmarking ChatGPT-4 on a radiation oncology in-training exam and Red Journal Gray Zone cases: potentials and challenges for ai-assisted medical education and decision making in radiation oncology. *Front. Oncol.* 13:1265024. doi: 10.3389/fonc.2023.1265024

COPYRIGHT

© 2023 Huang, Gomaa, Semrau, Haderlein, Lettmaier, Weissmann, Grigo, Tkhatay, Frey, Gaipf, Distel, Maier, Fietkau, Bert and Putz. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Benchmarking ChatGPT-4 on a radiation oncology in-training exam and Red Journal Gray Zone cases: potentials and challenges for ai-assisted medical education and decision making in radiation oncology

Yixing Huang^{1,2}, Ahmed Gomaa^{1,2}, Sabine Semrau^{1,2}, Marlen Haderlein^{1,2}, Sebastian Lettmaier^{1,2}, Thomas Weissmann^{1,2}, Johanna Grigo^{1,2}, Hassen Ben Tkhatay^{1,3}, Benjamin Frey^{1,2}, Udo Gaipf^{1,2}, Luitpold Distel^{1,2}, Andreas Maier³, Rainer Fietkau^{1,2}, Christoph Bert^{1,2} and Florian Putz^{1,2*}

¹Department of Radiation Oncology, University Hospital Erlangen, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany, ²Comprehensive Cancer Center Erlangen-EMN (CCC ER-EMN), Erlangen, Germany, ³Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany

Purpose: The potential of large language models in medicine for education and decision-making purposes has been demonstrated as they have achieved decent scores on medical exams such as the United States Medical Licensing Exam (USMLE) and the MedQA exam. This work aims to evaluate the performance of ChatGPT-4 in the specialized field of radiation oncology.

Methods: The 38th American College of Radiology (ACR) radiation oncology in-training (TXIT) exam and the 2022 Red Journal Gray Zone cases are used to benchmark the performance of ChatGPT-4. The TXIT exam contains 300 questions covering various topics of radiation oncology. The 2022 Gray Zone collection contains 15 complex clinical cases.

Results: For the TXIT exam, ChatGPT-3.5 and ChatGPT-4 have achieved the scores of 62.05% and 78.77%, respectively, highlighting the advantage of the latest ChatGPT-4 model. Based on the TXIT exam, ChatGPT-4's strong and weak areas in radiation oncology are identified to some extent. Specifically, ChatGPT-4 demonstrates better knowledge of statistics, CNS & eye, pediatrics, biology, and physics than knowledge of bone & soft tissue and gynecology, as per the ACR knowledge domain. Regarding clinical care paths, ChatGPT-4 performs better in diagnosis, prognosis, and toxicity than brachytherapy and dosimetry. It lacks proficiency in in-depth details of clinical trials. For the Gray Zone cases, ChatGPT-4 is able to suggest a personalized treatment approach to each case with high correctness and comprehensiveness. Importantly, it provides novel treatment aspects for many cases, which are not suggested by any human experts.

Conclusion: Both evaluations demonstrate the potential of ChatGPT-4 in medical education for the general public and cancer patients, as well as the potential to aid clinical decision-making, while acknowledging its limitations in certain domains. Owing to the risk of hallucinations, it is essential to verify the content generated by models such as ChatGPT for accuracy.

KEYWORDS

large language model, radiotherapy, natural language processing, artificial intelligence, Gray Zone, clinical decision support (CDS)

1 Introduction

With the recent advances in deep learning techniques such as transformer architectures (1), few-shot prompting (2), and reasoning (3), large language models (LLMs) have achieved breakthroughs in natural language processing. In these years, many LLMs have been developed and released to the public, including ChatGPT (2, 4), Bard (4), LLaMA (5), and PaLM (6). The capabilities of such LLMs range from simple text-related tasks like language translation and text refinement to complex ones like decision-making (7) and programming (8).

Like other fields, LLMs have also shown great potential in biomedical applications (9–11). Several domain-specific language models have been developed such as BioBERT (10), PubMedBERT (9), and ClinicalBERT (12). General-domain LLMs with fine-tuning on biomedical data have also achieved impressive results. For example, Med-PaLM (6) fine-tuned from PaLM has achieved 67.6% accuracy on the MedQA exam; ChatDoctor (11) fine-tuned from LLaMA (5) using doctor-patient conversation data for more than 700 diseases has achieved 91.25% accuracy on medication recommendations; HuaTuo (13) fine-tuned from LLaMA (5) is capable of providing advice on (traditional and modern) Chinese medicine with safety and usability. In addition to fine-tuning, hybrid models, which combine LLMs with models of other modalities, can extend the capabilities of general LLMs. For example, integrating ChatGPT with advanced imaging networks [e.g. the ChatCAD (14)] can overcome some of its limitation in image processing. Multimodal models may soon reach the goal of fully automatic diagnosis from medical images, along with automated medical report generation.

ChatGPT, being the most successful language model so far, has shown impressive performance in various domains without further fine-tuning, because of the large variety and amount of training data. It has been demonstrated successful in dozens of publicly-available official exams ranging from natural language processing like SAT EBRW reading and writing exams to subject-specific exams such as SAT Math, AP Chemistry and AP Biology exams, as reported in (15). ChatGPT is capable of improving its performance using reinforcement learning from human feedback (RLHF) (16). Because of its excellent performance on multidisciplinary subjects, ChatGPT becomes a very useful tool for diverse users. For domain knowledge in medicine, ChatGPT-3 achieved better than 50% accuracy across all the exams of the United States Medical Licensing Exam (USMLE) and exceeding 60% in certain analyses

(17); ChatGPT-3.5 was also reported beneficial for clinical decision support (18). Therefore, ChatGPT possesses the potential to enhance medical education for patients and decision support for clinicians.

In the field of radiation oncology, deep learning has achieved impressive results in various tasks (19), e.g., tumor segmentation (20, 21), lymph node level segmentation (22), synthetic CT generation (23), dose distribution estimation (24), and treatment prognosis (25, 26). With the wide spread of ChatGPT and its broad knowledge in medicine, ChatGPT has the potential to be a valuable tool for providing advice to cancer patients and radiation oncologists. Recently, ChatGPT's performance for the specialized domain of radiation oncology physics has been evaluated using a custom-designed exam with 100 questions (27), demonstrating the superiority of ChatGPT-4 to Bard, another LLM. Nevertheless, the field of radiation oncology covers diverse topics like statistics, biology, and anatomy specific oncology (e.g., gynecologic, gastrointestinal and genitourinary oncology) in addition to physics. To date, the performance of ChatGPT on radiation oncology using standard exams has not been benchmarked yet. Especially, its performance on real clinical cases has not been fully investigated. Consequently, the reliability of advice on radiation oncology provided by ChatGPT remains an open question (8).

In this work, the performance of ChatGPT on the American College of Radiation (ACR) radiation oncology in-training (TXIT) exam and the Red Journal Gray Zone cases is benchmarked. The performance difference between ChatGPT-3.5 and ChatGPT-4 is evaluated. Based on the two evaluations, the confidence zones and blind spots of ChatGPT in radiation oncology are revealed, highlighting its potential to medical education for patients and challenges for aiding clinicians in decision making.

2 Materials and methods

2.1 Benchmark on the ACR TXIT exam

To benchmark the performance of ChatGPT on radiation oncology, the 38th ACR TXIT exam (2021) is used. The exam sheet is publicly available on the ACR website¹. The TXIT exam

1 <https://www.acr.org/-/media/ACR/Files/DXIT-TXIT/ACR-2021-TXIT-Exam-Assembled.pdf>

covers the questions from six primary categories in radiation oncology: diagnosis, treatment decision, treatment planning, quality assurance, brachytherapy, and toxicity & management (28). The exam consists of 300 questions in total with 14 questions containing medical images. All questions are multiple choices with single answer. Among the 14 questions with medical images, 7 have a certain text description or list of answers from which the image content or the correct answer can be deduced (e.g., Question 1 as displayed in Figure 1). However, the other 7 questions (in particular, Questions 17, 86, 112, 116, 125, 143, and 164) are impossible to answer from the text information alone without access to the imaging information (e.g., Question 116 as displayed in Figure 1). Therefore, the later 7 questions were excluded in our evaluation.

In this study, all questions (questions alone without additional texts) were entered into ChatGPT. Although no justification was requested in the input prompt, ChatGPT automatically provided certain explanations for its responses. To maintain consistency, no human feedback was given to ChatGPT. In regard to grading ChatGPT's answers, Question 71 has two correct answers A and B, while ChatGPT-3.5 and ChatGPT-4 both gave a single choice answer (A and B, respectively). Therefore, a score of 0.5 was assigned to this question for both ChatGPT-3.5 and ChatGPT-4. In addition, for Question 20 ChatGPT-3.5 suggested A or C depending on whether bone is considered part of the anatomically constrained area; For Question 135, ChatGPT-4 suggested that D may also be a viable option in addition to the correct answer C. For each case, a score of 0.5 was assigned.

In the initial evaluation of the TXIT exam (where the main findings are based on), the ChatGPT website interfaces for the default ChatGPT-3.5 and the advanced ChatGPT-4 were used. They were accessed in April 2023. While ChatGPT-3.5 was readily available to the public as the standard version, ChatGPT-4 at the time of writing came with a usage restriction of 25 messages per 3-hour window and was not offered free of charge. In the initial evaluation, a new chat session was recreated after every 5 questions to avoid memory problems. As the chat history influences ChatGPT's responses, for a more fair evaluation, ChatGPT-3.5's and ChatGPT-4's responses on the TXIT exam were assessed again with the ChatGPT API with a temperature parameter of 0.7 on August 10-15th 2023, where a new conversation was created after each question to avoid the influence of chat history.

2.2 Benchmark on the Red Journal Gray Zone cases

Within the field of radiation oncology as for the whole of medicine, treatment guidelines and available clinical evidence do not always provide a clear recommendation for every clinical case. These difficult clinical situations are referred to as Gray Zone cases (29), leaving room for differences of opinion and constructive debate in many patient scenarios. The vast collection of the Red Journal Gray Zone cases (29) provides ample data that can be used to benchmark the performance of ChatGPT on such real,

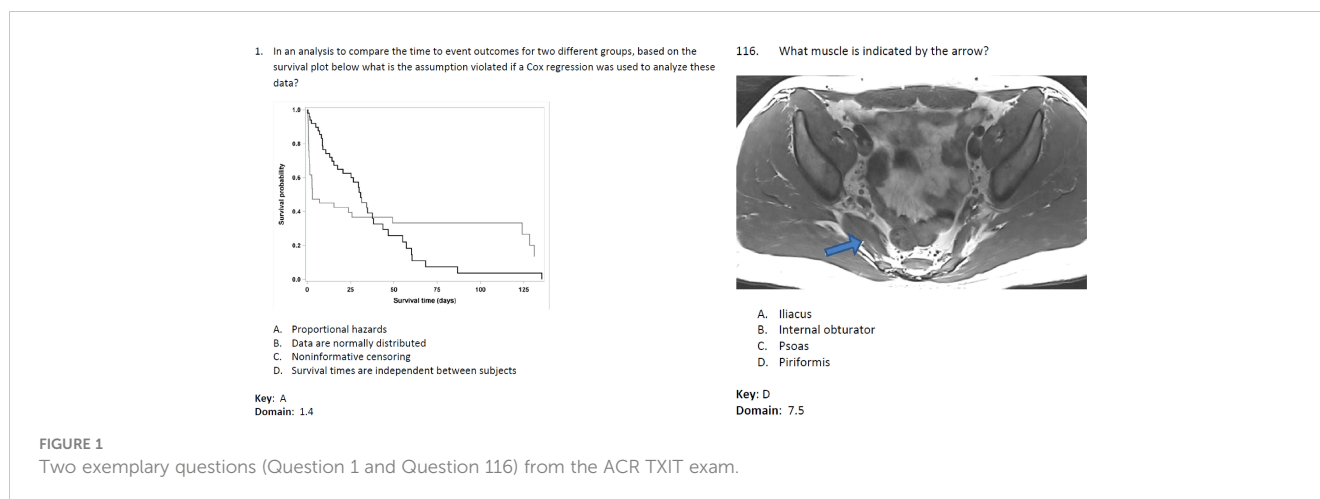
challenging clinical cases, which traditionally have required highly specialized domain experts and sophisticated clinical reasoning.

In this work, the 2022 collection of the Red Journal Gray Zone cases², in total 15 cases, were used for benchmark. Due to the superior performance of ChatGPT-4 to ChatGPT-3.5 based on the ACR TXIT exam, ChatGPT-4 was used for this evaluation. For each case, ChatGPT-4 was set to a role as an expert radiation oncologist by providing the prompt: "You are an expert radiation oncologist from an academic center", followed by the description of each patient's situation. For diagnostic medical images, only the text captions were provided. Based on the given information for the patient, ChatGPT-4's most favored therapeutic approach and its reasoning for the recommended approach were asked. Afterwards, other experts' recommendations to this case were provided to ChatGPT-4 and the following questions were asked:

- Summarize the recommendations of other experts in short sentences;
- Which expert's recommendation ChatGPT-4 thinks is the most proper for the patient;
- ChatGPT-4's initial recommendation is close to which expert's recommendation;
- Whether ChatGPT-4 will update its initial recommendation after seeing other experts' recommendations.

For all the Gray Zone cases a clinical expert (senior physician, board-certified radiation oncologist) evaluated the responses of ChatGPT-4 both in a qualitative and semiquantitative manner. The initial and updated recommendations of ChatGPT-4 were evaluated across 4 dimensions. First, the correctness of the responses was evaluated on a 4-point Likert scale with the following levels: 4 = "no mistakes", 3 = "mistake in detail aspect not relevant to the validity of the overall recommendation", 2 = "mistake in relevant aspect of the recommendation, but recommendation still clinically justifiable", 1 = "recommendation not clinically justifiable, because of incorrectness". Moreover, the comprehensiveness of the recommendation was also evaluated on a 4-point Likert scale with the following levels: 4 = "recommendation covers all relevant clinical aspects", 3 = "recommendation is missing some detail information, e.g., in regard to radiotherapy dose or target volume", 2 = "recommendation is missing relevant aspect, but overall recommendation is still clinically justifiable", 1 = "recommendation not clinically justifiable, because of incompleteness". Finally, novel valuable aspects in ChatGPT-4's response, not present in the real clinical experts' recommendations, as well as hallucinations were rated in a binary manner ("present" vs. "not present"). Hallucinations are responses generated by LLMs in a convincing appearance but actually are incorrect statements (30). Ratings for initial and revised recommendations were tested for difference using a paired Wilcoxon test. Aside from the initial and the final recommendation, all other responses by ChatGPT-4 were evaluated in a qualitative manner.

² <https://www.redjournal.org/content/grayzone>



3 Results

3.1 Results on the ACR EXIT exam

3.1.1 Overall performance difference between ChatGPT-3.5 and ChatGPT-4

Across the total of 293 questions, ChatGPT-3.5 and ChatGPT-4 attained accuracies of 63.14% and 74.06% respectively in our initial assessment via the website interface, both surpassing the standard pass rate of 60%. In the 5-time repeated assessment via the ChatGPT API, ChatGPT-3.5 and ChatGPT-4 achieved average accuracies of $62.05\% \pm 1.13\%$ and $78.77\% \pm 0.95\%$, respectively. The advanced ChatGPT-4 version exhibited a 16.72% increase in accuracy compared to the standard ChatGPT-3.5, illustrating its superior performance in the field of radiation oncology. Out of all the questions that were answered incorrectly, ChatGPT-3.5 and ChatGPT-4 were both incorrect for 51 of them. Additionally, ChatGPT-3.5 was incorrect for 57 questions that were correctly answered by ChatGPT-4, while ChatGPT-4 was incorrect for 25 questions that were correctly answered by ChatGPT-3.5. Please find the answers to all the questions in the [Supplementary Material](#). A copy of answers is available via [GitHub](#)³.

Figure 2 illustrates an exemplary instance of question-answering, where both ChatGPT-3.5 and ChatGPT4 successfully deliver accurate responses. In general, ChatGPT-3.5 provides longer answers with a quicker generation speed than ChatGPT-4. In this example, ChatGPT-3.5 refers to the NSABP B-51/RTOG 1304 trial (31) for its justification of the answer. Instead, ChatGPT-4 typically provides an answer with a shorter explanation. More frequently, ChatGPT-4 includes a cautionary message to prevent users from being inadvertently led towards potential health hazards, e.g., “*It is important to consult with a radiation oncologist and other members of the multidisciplinary cancer care team to determine the most appropriate treatment plan for the individual patient*” at the end of the answer to this exemplary question.

3.1.2 Domain dependent performance

All the questions of the ACR TXIT exam belong to 13 major knowledge domains, according to the TXIT table of specifications⁴. The 13 domains are statistics, bone & soft tissue, breast, central nervous system (CNS) & eye, gastrointestinal, genitourinary, gynecology, head & neck & skin, lung & mediastinum, lymphoma & leukemia, pediatrics, biology, and physics. The accuracies achieved by ChatGPT-3.5 and ChatGPT-4 for different domains are displayed in [Figure 3](#). Considering a 60% threshold, ChatGPT-3.5 only obtained 54.17%, 50%, 57.89%, 41.18%, 58.33%, and 40% for breast, gastrointestinal, genitourinary, gynecology, head & neck & skin and pediatrics, respectively. Notably, its accuracy for gynecology and pediatrics was only around 40%. In contrast, for statistics, CNS & eye, and biology, ChatGPT-3.5 achieved accuracies higher than 70%.

ChatGPT-4 attained a worse accuracy than ChatGPT-3.5 for bone & soft tissue and lymphoma & leukemia with accuracies of 50% and 45%, respectively. Note that there are only 4 valid questions for bone & soft tissue, and ChatGPT-4 answered 2 questions incorrectly (Question 18: *What postoperative RT dose is recommended for a high grade malignant peripheral nerve sheath tumor of the upper extremity following R1 resection?* Question 19: *What is the recommended preoperative GTV to CTV target volume expansion for an 8.5 cm high grade myxofibrosarcoma of the vastus lateralis muscle?*). ChatGPT-4 got the same bad performance (41.18%) on gynecology. ChatGPT-4 outperformed ChatGPT-3.5 in all other domains, with particularly impressive results of 100% accuracy for statistics, 90% accuracy for CNS & eye, 83.33% accuracy for gastrointestinal, and 86.67% accuracy for physics.

3.1.3 Clinical care related performance

Out of all the questions, the majority (totalling 190) are related to clinical care. Thus, beyond the standard ACR TXIT domain-dependent performance evaluation, these 190 questions—excluding those related to statistics, biology, and physics—are further

³ <https://github.com/YixingHuang/ChatGPT-Benchmark-on-Radiation-Oncology>

⁴ <https://www.acr.org/-/media/ACR/Files/DXIT-TXIT/ACR-TXIT-Table-of-Specifications.pdf>

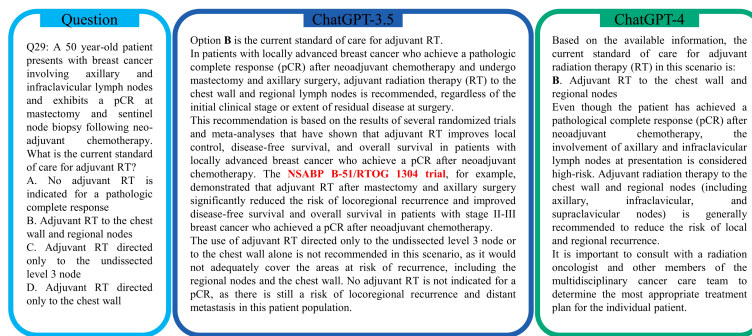


FIGURE 2
 An example of question-and-answer using ChatGPT for the ACR TXIT exam. ChatGPT-3.5 and ChatGPT-4 both provide the correct answer. However, ChatGPT-3.5 hallucinates the results of the NSABP B-51/RTOG 1304 trial (31), as the final findings are not publicly available yet.

classified into the following clinical care path-related categories: diagnosis, treatment decision, treatment planning, prognosis, toxicity, and brachytherapy. The prognosis category covers questions related to patient survival and tumor recurrence rate/risk, while the toxicity category covers questions related to side effects of treatment. Since dosimetry plays a crucial role in radiation therapy, all the dose-related questions are also grouped into one common category. Note that the categorization is not exclusive, which means that one question might belong to more than one category.

The accuracy distribution of ChatGPT-3.5 and ChatGPT-4 depending on the clinical care categories is displayed in Figure 4. Considering the 60% threshold, ChatGPT-3.5 passed this threshold only for diagnosis with a high accuracy of 71.88%. However, the accuracies for all the remaining categories were 56.72%, 50%, 36.84%, 46.15%, 16.67%, and 46.67%, respectively, all lower than 60%. In comparison, ChatGPT-4 passed the threshold for diagnosis, treatment decision, treatment planning and toxicity with accuracies of 81.25%, 62.5%, 71.05%, and 76.92%, respectively. But its accuracies for brachytherapy and dosimetry were lower than 60%, which were 16.67% and 56.67%, respectively. Out of all the categories, both ChatGPT-3.5 and ChatGPT-4 exhibited similarly unsatisfactory performance on brachytherapy, as they were only capable of correctly answering the same specific single question

from the total of six presented. Among other categories, ChatGPT-4 exhibited superior performance compared to ChatGPT-3.5, particularly in the areas of prognosis and toxicity, where ChatGPT-4 surpassed its predecessor by 30%.

3.1.4 Performance on clinical trials

Among all the questions, many questions are based on certain clinical trials [e.g., the Stockholm III (32), the CRITICS randomized trial (33), the PORTEC-3 trial (34), the German rectal study (35), and the ORIOLE phase 2 randomized clinical trial (36)] or guidelines [e.g., the 8th AJCC cancer staging manual (37)]. Such questions are also grouped into a category called trial/study, and the accuracies of ChatGPT-3.5 and ChatGPT-4 are displayed in Figure 4 as well.

ChatGPT-3.5 and ChatGPT-4 obtained the accuracies of 30.77% and 53.85% on the trial/study related questions, respectively, both of which were lower than 60%. ChatGPT-4 achieved 23% higher accuracy than ChatGPT-3. When we asked ChatGPT-3.5 and ChatGPT-4 whether they know a certain trial, e.g., the PORTEC-3 trial (34), both of them will provided a positive answer “Yes, I am familiar with the PORTEC-3 trial” and provide a short summary of the mentioned trial. This suggests that both ChatGPT3.5 and ChatGPT-4 have encountered such trials and studies in their training data. However, there is still a significant

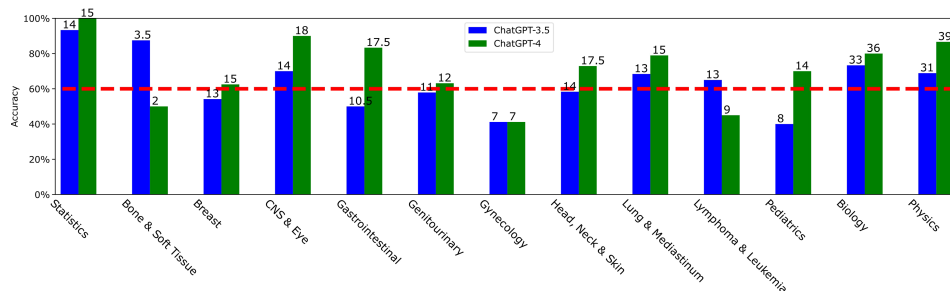


FIGURE 3
 The accuracy distribution for ChatGPT-3.5 and ChatGPT-4 depending on the question domain. The absolute number of correct answers for each domain is marked at the top of each bar. The domain number 1-13 correspond to statistics, bone & soft tissue, breast, CNS & eye, gastrointestinal, genitourinary, gynecology, head & neck & skin, lung & mediastinum, lymphoma & leukemia, pediatrics, biology, and physics, respectively. The X-axis labels are shifted to save space.

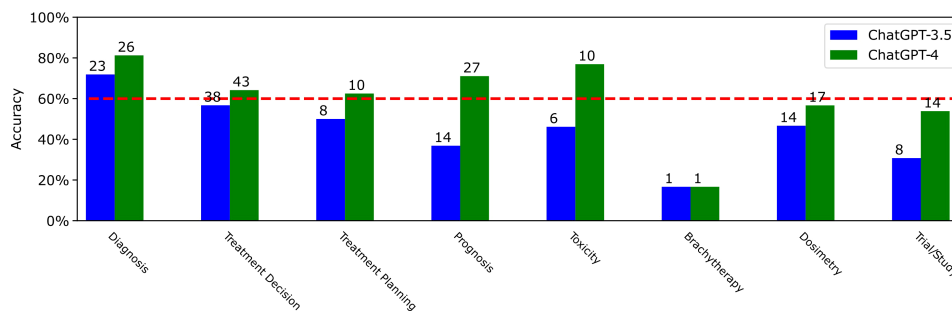


FIGURE 4

The accuracy distribution for ChatGPT-3.5 and ChatGPT-4 depending on the clinical care category. The absolute number of correct answers for each domain is marked at the top of each bar. The category number 1-8 correspond to diagnosis, treatment decision, treatment planning, prognosis, toxicity, brachytherapy, dosimetry, and trial/study, respectively. The X-axis labels are shifted to save space.

risk of them providing inaccurate answers. For example, for Question 107 (“*In the subset analysis of PORTEC-3 trial, patients with which histology MOST benefited from the addition of chemotherapy to RT? A. Endometrioid B. Carcinosarcoma C. Clear cell D. Serous*”), ChatGPT-3.5 answered A endometrioid, while ChatGPT-4 answered B carcinosarcoma, both of which are incorrect. When we copied and pasted the summary/abstract of the PORTEC-3 trial (34) into the conversation, i.e., leverage ChatGPT’s in-context learning capabilities (2), and asked ChatGPT the question again, both ChatGPT-3.5 and ChatGPT-4 provided the correct answer D based on the given summary.

It is worth noting that in the example of Figure 2, the NSABP B-51/RTOG 1304 has not been fully published yet and only a meeting update on the course of the study without any definitive results was included in the short abstract of (31). Therefore, ChatGPT-3.5 is hallucinating the results here in Figure 2.

3.2 Results for clinical decision making in the Gray Zone cases

3.2.1 An exemplary Gray Zone case

ChatGPT-4’s recommendations for one exemplary Gray Zone case (38) are displayed in Figure 5. In this example, a 55-year-old woman was treated in the past for mixed invasive lobular and ductal carcinoma of the left breast developed contralateral nodal recurrence after an interval of 10 years. More details about this patient can be found in (38) and its detailed responses to other cases are available in our GitHub repository. ChatGPT-4 proposed a combination of endocrine therapy (39) with an aromatase inhibitor and a CDK4/6 inhibitor, regional nodal irradiation, close monitoring and appropriate supportive care. ChatGPT-4 made such a recommendation based on the patient’s genotype and historic treatment. In addition, ChatGPT-4 provided a concise summary of the five experts’ recommendations, as displayed in Figure 5. ChatGPT-4 stated that its initial recommendation aligned most closely with Expert 3’s recommendation because both suggested focusing on treating the current contralateral axillary nodal metastases and involved the use of endocrine therapy and a CDK2/6 inhibitor as part of the systemic therapy. Nevertheless, ChatGPT-4

favors Expert 2’s recommendation instead of Expert 3’s since Expert 2’s recommendation considers the patient’s suboptimal initial treatment and the potential benefits of aggressive locoregional therapy in controlling the disease. Therefore, after seeing all five experts’ opinions, ChatGPT-4 tended to update its recommendation, “drawing inspiration from Expert 2’s recommendation”. This exemplary case demonstrates the potential of ChatGPT-4 in assisting decision making for intricate Gray Zone cases.

3.2.2 Overall performance based on ChatGPT-4’s self-assessment

For the Gray Zone cases, the recommendations provided by clinical experts for each case were voted by other expert raters and published on the Red Journal website. The distribution of votes for the 15 cases in the 2022 collection is displayed in Table 1. For instance, for the first case (40), Expert 1, Expert 2, and Expert 3 received 61.54%, 15.38%, and 23.08% of votes, respectively, out of a total of 13 votes, while Expert 3 and Expert 4 received no votes. A total of 59 expert recommendations were evaluated for the 15 Gray Zone cases, resulting in an average vote of 25.42% (15/59) for each expert. ChatGPT-4’s initial recommendation for each case typically shares common points with a specific expert, which has two implications: (a) ChatGPT-4’s recommendation is comparable to that of a human expert, and (b) ChatGPT-4’s recommendation provides complementary information to that of other individual experts. For certain cases [Case #10 (41), Case #13 (42), and Case #15 (43)], ChatGPT-4’s recommendation covered points from two experts, indicating that ChatGPT-4’s recommendation was more comprehensive. If we consider the closest expert vote to ChatGPT-4’s recommendation as an approximate evaluation metric, ChatGPT-4’s recommendation would receive an average vote of 28.76%. Similarly, ChatGPT-4’s preferred recommendation from other experts received an average vote of 24.99%. Both values (28.76% and 24.99%) are close to the experts’ average vote of 25.42%.note

In the evaluation of ChatGPT-4’s recommendations for the Gray Zone cases, ChatGPT-4 consistently highlights the significance of a multidisciplinary team in seeking an individualized treatment plan that is balanced or comprehensive. Its preferred recommendations typically take into account the

Question
You are an expert radiation oncologist from an academic center. Please give your recommendation on the following case: "... (text description)"
Question for the expert:
Please give a specific recommendation of your most favored therapeutic approach in this patient case based on the available information and your reasoning for it.

ChatGPT-4
Based on the available information, the patient has **locally recurrent breast cancer** with contralateral axillary lymph node involvement. Considering the patient's history and pathology findings, my recommendation for the therapeutic approach would include the following:
Systemic therapy: Given that the patient has hormone receptor-positive (ER/PR+) and HER2-negative disease, initiating endocrine therapy with an aromatase inhibitor (AI) would be my first choice. In addition, considering the prior discontinuation of tamoxifen due to intolerance, AI would be a better option for this patient. Additionally, since the patient has locally advanced and recurrent disease, I would also recommend adding a CDK4/6 inhibitor, such as palbociclib, ribociclib, or abemaciclib, in combination with the AI to improve the efficacy of treatment.
Radiation therapy: As the patient has not received radiation therapy in the past and given the involvement of contralateral axillary lymph nodes with extranodal extension, I would recommend considering regional nodal irradiation, which includes the chest wall (in the case of a mastectomy) and regional lymph nodes (axillary, supraclavicular, and internal mammary). This will help in reducing the risk of further local-regional recurrence and may improve overall survival.
Close monitoring: The patient should undergo regular follow-ups with physical examinations, and imaging as needed, to closely monitor her response to the treatment and to detect any signs of disease progression or new metastases.
Supportive care: It is essential to address any potential side effects or complications of the treatments mentioned above, such as fatigue, bone health issues, and lymphedema. This would involve coordinating with a multidisciplinary team of healthcare professionals including medical oncologists, radiation oncologists, surgeons, physical therapists, and other supportive care services.
In summary, my favored therapeutic approach for this patient would be the combination of endocrine therapy with an aromatase inhibitor and a CDK4/6 inhibitor, regional nodal irradiation, close monitoring, and appropriate supportive care. This recommendation is based on the patient's hormone receptor-positive, HER2-negative status, prior treatment history, and the need to address local-regional disease control while minimizing the risk of further recurrence.

Question
Which expert's recommendation is your initial suggestion closest to?

ChatGPT-4
My initial recommendation aligns most closely with Expert 3's recommendation. Expert 3 suggests radiating the right-sided axillary levels 1 to 3 and supraclavicular nodes only and emphasizes the critical role of systemic therapy in reducing the chance of metastases. This approach focuses on treating the current contralateral axillary nodal metastases and involves the use of endocrine therapy and a CDK4/6 inhibitor as part of the systemic therapy.
While Expert 3's recommendation omits irradiation of the bilateral chest walls and the left axilla/supraclavicular nodes to reduce morbidity, it acknowledges the importance of long-term follow-up and potential salvage treatment in case of isolated recurrence. This recommendation balances the need for effective treatment of the current disease presentation while minimizing potential side effects and considering the patient's overall prognosis.

Question
After seeing the other experts' recommendations, are you going to update your own recommendation of your most favored therapeutic approach in this patient case? If so, what is your updated recommendation?

ChatGPT-4
After reviewing the other five experts' recommendations, I would like to update my initial recommendation to provide a more comprehensive approach for the patient, **drawing inspiration from Expert 2's recommendation.**
Updated recommendation:
Perform thorough imaging, including bilateral magnetic resonance imaging, to assess for potential primary or recurrent tumors in the residual ipsilateral or contralateral breast tissue.
Administer bilateral comprehensive radiation therapy to the chest wall and regional nodal irradiation (RNI), targeting undissected axillary nodal basins, supraclavicular nodes, and internal mammary nodes.
Initiate systemic therapy, including endocrine therapy with an aromatase inhibitor, a CDK4/6 inhibitor, and discuss the possibility of chemotherapy to mitigate the risk of distant recurrence and maximize progression-free survival and overall survival.
Conduct close monitoring and long-term follow-up to detect any signs of recurrence or progression early and provide appropriate supportive care.
This updated recommendation is primarily based on Expert 2's suggestion. It takes into account the patient's suboptimal initial treatment and the potential benefits of aggressive locoregional therapy in controlling the disease. By targeting both sides, it addresses the risk of residual disease on the left side and the possibility of aberrant lymphatic drainage after the axillary lymph node dissection. Furthermore, the recommendation includes a thorough assessment for potential primary or recurrent tumors through bilateral magnetic resonance imaging.
Adding systemic therapy to the treatment plan is essential to minimize the risk of distant metastases and maximize progression-free survival and overall survival. The multidisciplinary approach mentioned in this recommendation ensures a comprehensive treatment plan tailored to the patient's needs and overall prognosis.

Question
Please summarize the five experts' recommendations in short sentences.

ChatGPT-4
Expert 1: Recommends right postmastectomy radiation therapy (PMRT) to the chest wall and regional nodes, regardless of systemic therapy or axillary surgery. Omission of left-side radiation is suggested.
Expert 2: Suggests bilateral comprehensive radiation therapy to the chest wall and regional nodal irradiation (RNI), with discussion of additional systemic therapy.
Expert 3: Advises radiation to the right-sided axillary levels 1 to 3 and supraclavicular nodes only, with systemic therapy as a critical component for reducing the chance of metastases.
Expert 4: Recommends postoperative comprehensive radiation therapy with curative intent, targeting bilateral chest walls and lymph node areas, alongside subsequent aromatase inhibitor and CDK 4/6 inhibitor.
Expert 5: Favors right-sided locoregional radiation therapy with a hypofractionated dose to the chest wall and regional lymph nodes. The provision of chemotherapy would not alter this recommendation.

Question
As an expert radiation oncologist, which of the above experts' recommendation do you think the most proper for the patient described above?

ChatGPT-4
As an expert radiation oncologist, **I believe the recommendation provided**

FIGURE 5
An example of ChatGPT-4's recommendation for the Gray Zone case #8 (38): A viewpoint on isolated contralateral axillary lymph node involvement by breast cancer: regional recurrence or distant metastasis? Note that the local recurrence statement in ChatGPT-4's summary is incorrect.

patient's historic treatment responses, survival prognosis, and potential risks and toxicity. Additionally, it respects the personal preferences and priorities of the patient.

3.2.3 Clinical evaluation of ChatGPT-4's recommendations

The initial recommendations of ChatGPT-4 generally showed a similar structure: After a summary of the clinical case, ChatGPT-4 provided its recommendations for the patient vignette; At the end of the recommendations, more generic, universally agreed statements e.g., on the value of interdisciplinary discussion, post-treatment follow-up and supportive treatment, were typically observed (e.g., Figure 5). In the initial case summary, ChatGPT-4 typically was able to capture the relevant aspect of the patient case in a correct

manner. Errors were rarely observed in the case summary, but one example can be seen in the exemplary case (Case #8 (38) in Figure 5). In Figure 5, the patient case was described as "locally recurrent breast cancer with contralateral axillary lymph node involvement." However, this is not fully correct, as the patient had no local recurrence but only contralateral axillary lymph node metastasis. As the information on local recurrence is not present in the case vignette, this is also rated as hallucination (Table 1).

Generally, the initial recommendations of ChatGPT-4 showed a surprising amount of correctness and comprehensiveness. On a scale of 1-4, mean correctness of the 15 initial ChatGPT-4 recommendations was 3.5 and mean comprehensiveness was 3.1. Importantly, no recommendation was rated with the lowest score of 1, meaning that all ChatGPT-4 recommendations were seen as

TABLE 1 The performance of ChatGPT-4’s initial recommendations and revised recommendations on the Gray Zone cases.

Case ID	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Distribution of votes for the Gray Zone clinical expert recommendations:															
Expert 1	61.54	20	5.56	7.14	71.43	60	40	8	29.41	60	37.5	62.5	0	25	16.67
Expert 2	15.38	26.67	0	57.14	0	40	20	0	52.94	30	25	12.5	100	50	50
Expert 3	0	33.33	55.56	35.71	14.29	0	40	32	5.88	10	25	12.5	0	25	33.33
Expert 4	0	20	38.89	0	–	–	20	11.76	–	12.5	12.5	–	–	–	–
Expert 5	23.08	–	–	–	14.29	–	–	40	–	–	–	–	–	–	–
GPT-4’s self-assessment:															
Closest	E3	E2	E1	E1	E4	E2	E1	E3	E3	E1+E2 E3	E1	E2+E3	E2	E2+E3	
Favourite	E3	E3	E4	E1	E2	E2	E2	E2	E2	E1+E2 E3	E2	E1	E2	E2	
Senior physician’s assessment: Initial recommendation															
Correctness	4	4	3	4	4	4	3	2	4	3	4	3	4	3	4
Comprehensi.	3	4	3	2	3	2	4	2	4	4	3	3	2	4	4
Novel aspects	Yes	Yes	No	Yes	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes
Hallucination	No	No	No	No	No	No	No	Yes	No	Yes	No	No	No	No	No
Revised recommendation															
Correctness	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4
Comprehensi.	3	4	3	4	4	4	3	4	4	4	3	4	3	4	4
Novel aspects	Yes	No	No	Yes	No	Yes	No	Yes	No	No	No	No	No	No	Yes
Hallucination	No	No	No	No	No	No	No	No	No	No	No	No	No	No	No

Closest: ChatGPT-4’s initial recommendation is closest to which expert’s recommendation.
 Favourite: Which expert’s recommendation is the most proper for the patient.
 Comprehensi., Comprehensiveness.

clinically justifiable. Hallucinations were common with 13.3% (2/15) of recommendations.

Interestingly, ChatGPT-4 showed some signs of clinical reasoning, in which specific features of the case vignette were integrated into the decision making beyond standard clinical algorithms or in which recommendations were obtained by multi-step reasoning. For example in Figure 5, ChatGPT-4 recommended an aromatase inhibitor (AI) by referring to the patient’s prior intolerance to Tamoxifen (“[...] considering the prior discontinuation of tamoxifen due to intolerance, AI would be a better option for this patient. [...]”). A quite impressive example for multi-step clinical reasoning can be found in the recommendation for Case #9 (44) (Supplementary Material) on a recurrent prostate cancer case, in which an isolated supraclavicular lymph node showed PSMA-uptake: “The 18F-PSMA PET/CT shows no evidence of locoregional recurrence but identifies a hypermetabolic lymph node in the supraclavicular region on the left side. This finding could represent metastatic involvement, although it is unusual for prostate cancer to metastasize to this location without involving pelvic lymph nodes first. [...] it is important to consider a biopsy of the suspicious supraclavicular lymph node to confirm its nature and guide further management, which might include directed radiation therapy to the supraclavicular region if it is confirmed to be metastatic”.

Finally, it is interesting to note that we have also observed novel, valuable aspects in ChatGPT-4’s recommendations that are not

present in the real clinical experts’ responses with a frequency of 80.0% (12/15). These include the recommendation for considering tumor-treating field in a patient with grade 2 glioma [Case #10 (45)], the recommendation to incorporate potential drug interaction when choosing concurrent chemotherapy in a liver transplant patient [Case #13 (42)], the recommendation for genomic profiling to find targetable molecular alterations for systemic therapy in a patient case with oligometastatic breast cancer [Case #2 (46)] as well as the consideration for additional immunotherapy in a patient with anal cancer [Case #1 (40)].

In the subsequent conversation with ChatGPT-4 following its initial recommendation, the LLM generally was able to summarize the relevant aspects of the clinical experts’ recommendation well. In the updated ChatGPT-4 recommendation following in-context learning with the clinical expert recommendations, we observed a significant increase in correctness (mean, 4.0 vs. 3.5, paired p = 0.020), comprehensiveness (mean, 3.7 vs. 3.1, p = 0.046), no hallucinations (0.0% vs. 13.3%) but also reduced novel aspects (33.3% vs. 80.0%). Interestingly, in the revised ChatGPT-4 recommendation particularly valuable aspects of the clinical experts’ recommendation were incorporated. This e.g., includes the recommendation for neoadjuvant chemotherapy in a patient with metastatic anal cancer that received prior prostate brachytherapy [Case #1 (40)] to select, if the patient will benefit from locoregional treatment as well as the

recommendation to perform bilateral chest wall MRI in a patient case with contralateral supraclavicular lymph node recurrence from breast cancer (Case #8 (38) in Figure 5).

4 Discussion

4.1 Potential for medical education in radiation oncology

Both ChatGPT-3.5 and ChatGPT-4 have demonstrated a certain level of proficiency in grasping fundamental concepts of radiation oncology. For instance, both versions of ChatGPT are proficient in identifying common types of cancer and have a certain awareness of clinical trials and studies. In the evaluation of Gray Zone cases, ChatGPT-4 provided reasonable explanations to the recommended treatment approach and received high ratings for correctness and comprehensiveness. As such, ChatGPT possesses the potential to offer medical education on radiation oncology to the general public and cancer patients, promoting radiation oncology education into a new stage (47), when the risk of content hallucination is accounted for by proper verification.

4.2 Potential to assist in clinical decision making

ChatGPT-4's decent performance on the topics of diagnosis, treatment decision, treatment planning, prognosis and toxicity (Figure 4) and its reasonable responses on the Gray Zone cases (Figure 5; Table 1) indicate its potential to assist in clinical decision making. As a single human expert may fail to consider all aspects of an intricate Gray Zone case, ChatGPT-4's recommendation can provide valuable complementary information in certain cases, potentially leading to a more comprehensive treatment approach. Especially, ChatGPT is capable of suggesting novel treatment ideas (novel aspects in Table 1), for example, using tumor-treating field (45) for the patient case with gliomas (41), which was not suggested by any of the human experts. Therefore, such general artificial intelligence like ChatGPT can in turn improve human decision-making by increasing novelty (48).

4.3 Challenges in clinical decision making for certain topics

While both versions of ChatGPT exhibited a grasp of essential concepts in radiation oncology, their knowledge was limited or superficial when it comes to certain topics such as gynecology, brachytherapy, dosimetry, and clinical trials, based on the TXIT exam. Consequently, these areas may elicit a relatively high rate of false responses when queried in-depth. In the evaluation of the Gray Zone cases, the recommendations closest to ChatGPT-4's response or its favored expert recommendations have received low votes from the raters for some cases, like Case 1, Case 4, Case 5, and Case 13 in Table 1. Therefore, despite of its potential, ChatGPT still has certain limitations in clinical decision making.

4.4 Addressing the risk of hallucination

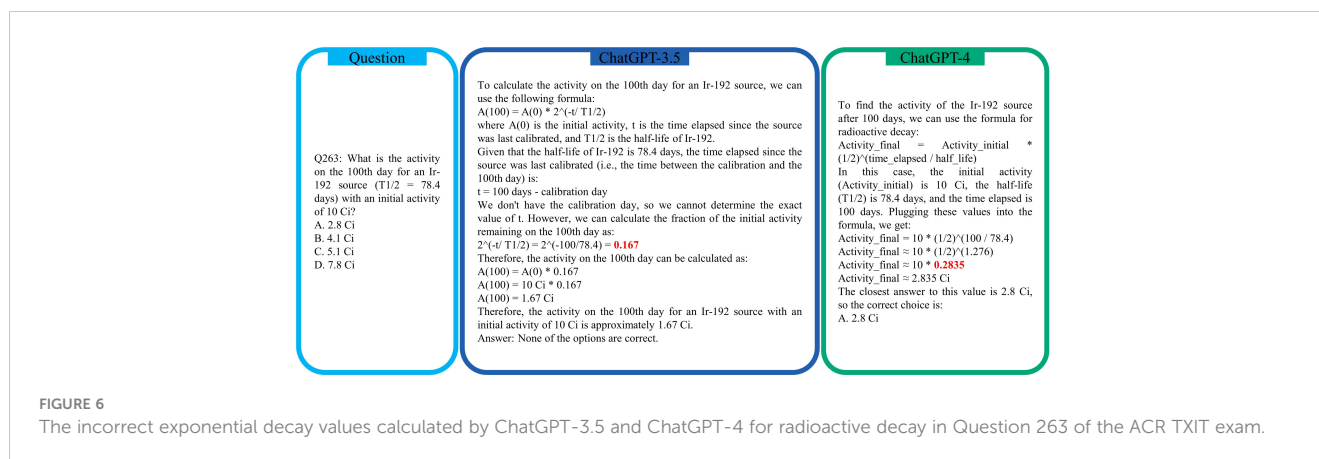
LLMs can hallucinate information (49) when generating responses, which becomes a widely-known limitation. Especially, for the field of medicine, it is critical that the provided information is correct. We mostly observed hallucinations in the context of clinical trials and citations. For example in this evaluation, ChatGPT-3.5 cited results of the NSABP B-51/RTOG 1204 study (Figure 2), which have not been published yet; ChatGPT-4 falsely added the local recurrence in the summary of Case #8 in Figure 5. The tendency of LLMs for hallucinations may be particularly problematic as they may be missed by the less proficient reader, since the hallucinations frequently appear very plausible in the context of the text and are phrased in a convincing manner. Because of the risk of content hallucination, answers and recommendations by ChatGPT always need to be verified. Potential solutions to reduce the risk of hallucination include in-context learning as well as model fine-tuning on medical studies and guidelines.

4.5 Responses on math need cross-check

Regarding simple math questions like calculating the mode (Question 6), mean (Question 7), and median (Question 8) of a sequence, both ChatGPT-3.5 and ChatGPT-4.0 were able to solve them correctly. However, for exponential/radioactive decay, although ChatGPT-3.5 and ChatGPT-4 both know the right mathematical expression, they nevertheless are likely to provide an incorrect value. For example, for the radioactive decay calculation in Question 263, both ChatGPT-3.5 and ChatGPT-4 could not obtain the correct value for $2^{(-100/78.4)} = 2^{(-1.2755)} = 0.413$, as displayed in Figure 6. We observed that if this question is asked multiple times in new conversations, each time a different value will be generated by ChatGPT-3.5 and ChatGPT-4. When ChatGPT-4 was asked to calculate the intermediate step $(1/2)^{(100/78.4)}$ specifically, it was able to provide a more accurate answer of 0.4129 and hence updated its previous answer to the correct answer B 4.1 Ci, while ChatGPT-3.5 still failed with an inaccurate value 0.2456. It is worth noting that at the time of writing (May 2023), ChatGPT-4 has the new feature of using external plug-ins. With the enabled plug-in of Wolfram Alpha, ChatGPT-4 can do mathematical calculation accurately and deliver the correct answer B directly for Question 263.

4.6 Analysing medical images using visual input is impossible yet

Despite being equipped with a new feature of visual input, ChatGPT-4 falls short in its ability to describe the content of medical images. In the evaluation of the ACR TXIT exam, when presented with an image such as the one in Question 116 as depicted in Figure 1 via a URL link, ChatGPT-4 failed to provide any meaningful context regarding the image. Instead, it provided a generic response, stating that it could not view images as an AI language model. While ChatGPT is capable of generating



descriptions of images based on accompanying text descriptions, such descriptions may not always align with the actual content of the image. The same observation was also reported in (50), where ChatGPT failed to identify retinal fundus images. As such, the current version of ChatGPT is not capable of analysing medical images in the same manner as a radiation oncologist and providing relevant diagnoses and treatment recommendations. Combining ChatGPT with medical imaging processing networks like ChatCAD (14) is promising in enhancing its capacity to analyze medical images in radiation oncology.

4.7 Potential for summarizing guidelines with in-context learning

In the TXIT evaluation, ChatGPT demonstrated its proficiency in accurately answering Question 107 when presented with a summary of the PORTEC-3 trial (34). In the evaluation of Gray Zone cases, ChatGPT was also able to provide concise summaries of other experts' opinions and update its initial recommendation based on other experts' opinions. The risk of hallucination was reduced after seeing other experts' opinions. As clinical guidelines are frequently updated, many clinicians may not be familiar with the latest details. Fortunately, ChatGPT is adept at summarizing text documents and, with in-context learning techniques (2), it is capable of rapidly acquiring new knowledge. When information is presented in context, ChatGPT can provide suggestions with a high degree of confidence. Consequently, ChatGPT has the potential to considerably assist clinicians in understanding updated guidelines and providing up-to-date treatment recommendations to patients based on the latest guidelines. Due to the extensive nature of many guidelines, like the NCCN cancer treatment guidelines, which often exceed the 4k token limit of the current ChatGPT interface, access to higher token limits, such as 32k tokens, is necessary for effective utilization in such applications.

4.8 Improvement with further domain-specific fine-tuning

At the time of this manuscript preparation, ChatGPT is not qualified as a specialist in radiation oncology yet, since ChatGPT

still lacks in-depth knowledge in many areas, as revealed by the TXIT exam and the Gray Zone cases. Similar to other fine-tuned LLMs like Med-PaLM (6), ChatDoctor (11), and HuaTuo (13), a radiation oncology domain-specific, fine-tuned LLM can be trained to better assist radiation oncologists in decision making for real clinical cases. The vast collection of Red Journal Gray Zone cases (29) and the latest guidelines provides ample data that can be automatically extracted for use in training such a domain-specific model, which is promising as our future work.

4.9 Capacity extension with external plug-ins

As a LLM, the fundamental function of ChatGPT is text generation. Hence, it has many limitations in many specialized tasks, for example, mathematical calculation in Figure 6. Some of such limitations can be addressed by the new feature of external plug-ins. Especially, as the training data for ChatGPT-4 dates back to 2021 September, it has very limited knowledge on the latest updates of guidelines for radiation oncology. Enabling the internet browsing function can improve ChatGPT-4's responses to such queries.

4.10 Regulatory approvals necessary for healthcare applications

The potential of LLM chatbots including ChatGPT-4 in assisting with medical diagnoses and patient care is vast, but their integration into the medical domain comes with a responsibility to ensure the highest standards of safety and efficacy. International consensus, echoing the principles laid out in proposed regulations from both the EU and the US, underscores the importance of ensuring bias control, transparency, oversight, and validation for AI in healthcare. However, current LLM chatbots fall short of adhering to these stringent principles. This disparity underscores the pressing need for LLM chatbots to undergo a rigorous regulatory approval process as medical devices. Obtaining this validation not only

establishes the accuracy and clinical efficacy of these models but also instills trust in medical professionals and patients. Gilbert et al. (51) further emphasize this imperative, advocating for a standardized regulatory framework that classifies LLM chatbots within the purview of medical devices, ensuring their safe and effective use in healthcare applications including radiation oncology.

4.11 Growing capabilities of ChatGPT

In this ACR TXIT exam evaluation, ChatGPT-4 has demonstrated its superiority to ChatGPT-3.5 in both the general radiation oncology field and various knowledge sub-domains, despite its slower generation speed and limited access. Notably, the questions on clinical trials suggest that ChatGPT-3.5 and ChatGPT-4 were trained on similar data sets. Therefore, the enhanced performance of ChatGPT-4 may be attributed more to its superior interpretability and generation capabilities than to the potentially increased amount of training data. In the evaluation of Gray Zone cases, ChatGPT-4 is able to update its own recommendation based on other experts' recommendations. With ongoing technical advancements, continuously expanding training data, more feedback via RLHF (16), and more external plug-ins, future iterations of ChatGPT are expected to deliver even more impressive performance in all medical fields, including radiation oncology.

4.12 Limitations of this study

The study in this work is not without its limitations. The ACR TXIT exam and the Gray Zone cases in this work represent only a narrow spectrum of knowledge in radiation oncology, as only 293 questions from the TXIT exam and 15 cases from the Gray Zone collection were evaluated. While the gaps in knowledge such as gynecology and brachytherapy were detected in this work, other benchmark tests are likely to find different deficiencies across medical domains.

Another limitation is that our performance benchmark applies to ChatGPT-3.5 and ChatGPT-4 in a specific time window (from April to August 2023). LLMs including ChatGPT are frequently updated. Therefore, the performance is highly likely to vary when new updates are applied to ChatGPT in the future.

In the work of (27), the superiority of ChatGPT over Bard in radiation oncology physics has been demonstrated. Because of its superior performance, this work focuses on the benchmark of ChatGPT's performance in a broader field of clinical radiation oncology. However, ChatGPT's performance has not been compared comprehensively with other LLMs on the TXIT exam and the Gray Zone cases in this work. In this work, the performance of LLaMA-2 (52) on the TXIT exam has been preliminarily evaluated. It achieved an average accuracy of 34.81% (more details in the [Supplementary Material](#) in our GitHub), which

was lower than ChatGPT-4. LLaMA-2 with more parameters such as the 13b and 70b models are likely to have better

performance. However, due to our hardware limitation, we are not able to evaluate them at the current stage. Some LLMs, in particular Med-PaLM (6) and its newer version Med-PaLM 2 (53), are fine-tuned specifically for medical applications, which have the potential to outperform ChatGPT-4 in radiation oncology. Nevertheless, such comparison will be our future work once we have granted the access to Med-PaLM (version 1 or 2).

The TXIT exam has a standard answer to each question and hence it can be evaluated objectively and accurately. However, for the complex Gray Zone cases, no gold standards exist to assess the accuracy of ChatGPT-4's responses. As a consequence, intra-rater and inter-rater variability is one major limitation in our evaluation on the Gray Zone cases. To draw more definitive conclusions, a robust research design, such as providing concordance training for evaluators before the assessment, is recommended. Nonetheless, the present analysis provides valuable insights into expert perceptions of ChatGPT-4's proficiency in clinical decision-making within radiation oncology.

5 Conclusion

This study benchmarks the performance of ChatGPT-3.5 and ChatGPT-4 on the 38th ACR TXIT exam in radiation oncology and the 2022 Red Journal Gray Zone cases. For the TXIT exam, ChatGPT-3.5 and ChatGPT-4 have achieved accuracies of 63.65% and 74.57%, respectively, indicating the advantage of the latest ChatGPT-4 model. Based on the TXIT exam, ChatGPT-4's strong and weak areas in radiation oncology are identified to some extent. Specifically, ChatGPT-4 demonstrates better knowledge in statistics, CNS & eye, pediatrics, biology, and physics than in bone & soft tissue and gynecology, as per the ACR knowledge domain. Regarding clinical care paths, ChatGPT-4 performs better in diagnosis, prognosis, and toxicity than brachytherapy and dosimetry. And it lacks proficiency in in-depth details for clinical trials. For the Gray Zone cases, ChatGPT-4 is able to suggest a personalized treatment approach to each case with high correctness and comprehensiveness considering the patient's historic treatment response, personal priority, and quality of life. Most importantly, it provides novel treatment aspects for many cases, which are not suggested by any human experts. Both evaluations have demonstrated the potential of ChatGPT in medical education for the general public and cancer patients, as well as the potential to aid clinical decision-making, while acknowledging its limitations in certain domains. Despite these promising results, ChatGPT-4 is not competent for clinical use yet. ChatGPT's answers currently always have to be verified, because of the risk of hallucination, which is one of main remaining issues that will need to be addressed by future developments.

Data availability statement

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found in the article/[Supplementary Material](#).

Author contributions

YH: Conceptualization, Data curation, Investigation, Methodology, Software, Visualization, Writing – original draft, Writing – review & editing. AG: Investigation, Software, Writing – review & editing. SS: Investigation, Methodology, Validation, Writing – review & editing. MH: Investigation, Methodology, Validation, Writing – review & editing. SL: Investigation, Methodology, Validation, Writing – review & editing. TW: Investigation, Methodology, Validation, Writing – review & editing. JG: Investigation, Methodology, Software, Writing – review & editing. HT: Investigation, Methodology, Software, Writing – review & editing. BF: Investigation, Methodology, Software, Validation, Writing – review & editing. UG: Investigation, Methodology, Software, Validation, Writing – review & editing. LD: Investigation, Methodology, Software, Validation, Writing – review & editing. AM: Investigation, Resources, Supervision, Validation, Writing – review & editing. RF: Funding acquisition, Investigation, Resources, Supervision, Validation, Writing – review & editing. CB: Investigation, Methodology, Software, Validation, Writing – review & editing. FP: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Writing – original draft, Writing – review & editing.

References

- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. *Adv Neural Inf Process Syst* (2017) 30:1–13. Available at: https://papers.nips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.
- Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, et al. Language models are few-shot learners. *Adv Neural Inf Process Syst* (2020) 33:1877–901.
- Wei J, Wang X, Schuurmans D, Bosma M, Chi E, Le Q, et al. Chain of thought prompting elicits reasoning in large language models. *NeurIPS* (2022) 1–14.
- Thapa S, Adhikari S. ChatGPT, bard, and large language models for biomedical research: opportunities and pitfalls. *Ann Biomed Eng* (2023), 1–5. doi: 10.1007/s10439-023-03284-0#citeas
- Touvron H, Lavril T, Izacard G, Martinet X, Lachaux MA, Lacroix T, et al. LLaMA: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023) 1–17.
- Singhal K, Azizi S, Tu T, Mahdavi SS, Wei J, Chung HW, et al. Large language models encode clinical knowledge. *Nat* (2023) 620:172–80.
- Wang FY, Yang J, Wang X, Li J, Han QL. Chat with ChatGPT on industry 5.0: Learning and decision-making for intelligent industries. *IEEE/CAA J Autom Sin* (2023) 10:831–4.
- Ebrahimi B, Howard A, Carlson DJ, Al-Hallaq H. ChatGPT: Can a natural language processing tool be trusted for radiation oncology use? *Int J Radiat Oncol Biol Phys* (2023) 116(5):977–83.
- Gu Y, Tinn R, Cheng H, Lucas M, Usuyama N, Liu X, et al. Domain-specific language model pretraining for biomedical natural language processing. *ACM Trans Comput Healthc (Health)* (2021) 3:1–23.
- Lee J, Yoon W, Kim S, Kim D, Kim S, So CH, et al. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* (2020) 36:1234–40.
- Li Y, Li Z, Zhang K, Dan R, Jiang S, Zhang Y. Chatdoctor: A medical chat model fine-tuned on a large language model Meta-AI (LLaMA) using medical domain knowledge. *Cureus* (2023) 15(6):e40895. doi: 10.7759/cureus.40895
- Alsentzer E, Murphy J, Boag W, Weng WH, Jindi D, Naumann T, et al. Publicly available clinical BERT embeddings. *Proc Clin NLP Workshop* (2019) 72–8.
- Wang H, Liu C, Xi N, Qiang Z, Zhao S, Qin B, et al. Huatuo: Tuning LLaMA model with chinese medical knowledge. *arXiv preprint arXiv:2304.06975* (2023) 1–6.
- Wang S, Zhao Z, Ouyang X, Wang Q, Shen D. ChatCAD: Interactive computer-aided diagnosis on medical image using large language models. *arXiv preprint arXiv:2302.07257* (2023) 1–11.
- Open AI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774* (2023) 1–100.
- Christiano PF, Leike J, Brown T, Martic M, Legg S, Amodei D. Deep reinforcement learning from human preferences. *Adv Neural Inf Process Syst* (2017) 30:1–9. Available at: https://papers.nips.cc/paper_files/paper/2017/hash/d5e2c0adad503c91f91df240d0cd4e49-Abstract.html.
- Kung TH, Cheatham M, Medenilla A, Sillos C, De Leon L, Elepano C, et al. Performance of ChatGPT on USMLE: Potential for AI-assisted medical education using large language models. *PloS Digit Health* (2023) 2:e0000198.
- Liu S, Wright AP, Patterson BL, Wanderer JP, Turer RW, Nelson SD, et al. Using AI-generated suggestions from ChatGPT to optimize clinical decision support. *J Am Med Inform Assoc* (2023) 30(7):1237–45.
- Sahiner B, Pezeshk A, Hadjiiski LM, Wang X, Drukker K, Cha KH, et al. Deep learning in medical imaging and radiation therapy. *Med Phys* (2019) 46:e1–e36.
- Kamnitsas K, Ledig C, Newcombe VF, Simpson JP, Kane AD, Menon DK, et al. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Med Image Anal* (2017) 36:61–78.
- Huang Y, Bert C, Sommer P, Frey B, Gaip U, Distel LV, et al. Deep learning for brain metastasis detection and segmentation in longitudinal mri data. *Med Phys* (2022) 49:5773–86.
- Weissmann T, Huang Y, Fischer S, Roesch J, Mansoorian S, Ayala Gaona H, et al. Deep learning for automatic head and neck lymph node level delineation provides expert-level accuracy. *Front Oncol* (2023) 13.
- Wang H, Liu X, Kong L, Huang Y, Chen H, Ma X, et al. Improving cbct image quality to the ct level using reggan in esophageal cancer adaptive radiotherapy. *Strahlenther Onkol* (2023) 1–13.
- Xing Y, Nguyen D, Lu W, Yang M, Jiang S. A feasibility study on deep learning-based radiotherapy dose calculation. *Med Phys* (2020) 47:753–8.

Funding

The authors declare no financial support was received for the research, authorship, and/or publication of this article.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fonc.2023.1265024/full#supplementary-material>

25. Yang Y, Wang M, Qiu K, Wang Y, Ma X. Computed tomography-based deep-learning prediction of induction chemotherapy treatment response in locally advanced nasopharyngeal carcinoma. *Strahlenther Onkol* (2022) 198:183–93.
26. Hagag A, Gomaa A, Kornek D, Maier A, Fietkau R, Bert C, et al. Deep learning for cancer prognosis prediction using portrait photos by StyleGAN embedding. *arXiv preprint arXiv:2306.14596* (2023) 1–12.
27. Holmes J, Liu Z, Zhang L, Ding Y, Sio TT, McGee LA, et al. Evaluating large language models on a highly-specialized topic, radiation oncology physics. *Front Oncol* (2023) 13.
28. Rogacki K, Gutiontov S, Goodman C, Jeans E, Hasan Y, Golden DW. Analysis of the radiation oncology in-training examination content using a clinical care path conceptual framework. *Appl Radiat Oncol* (2021). 10(3):41–51.
29. Palma DA, Yom SS, Bennett KE, Corry J, Zietman AL. Introducing: The red journal gray zone. *Int J Radiat Oncol Biol Phys* (2017) 97:1.
30. Open AI. GPT-4 system card. *OpenAI.com* (2023), 6–6.
31. Mamounas EP, Bandos H, White JR, Julian TB, Khan AJ, Shaitelman SF, et al. NRG Oncology/NSABP B-51/RTOG 1304: Phase III trial to determine if chest wall and regional nodal radiotherapy (CWRNRT) post mastectomy (Mx) or the addition of RNRT to whole breast rt post breast-conserving surgery (BCS) reduces invasive breast cancer recurrence-free interval (IBCR-FI) in patients (pts) with pathologically positive axillary (PPAx) nodes who are ypN0 after neoadjuvant chemotherapy (NC). *J Clin Oncol* (2019) 37:TPS600–0.
32. Erlandsson J, Holm T, Pettersson D, Berglund A, Cedermark B, Radu C, et al. Optimal fractionation of preoperative radiotherapy and timing to surgery for rectal cancer (stockholm iii): a multicentre, randomised, non-blinded, phase 3, non-inferiority trial. *Lancet Oncol* (2017) 18:336–46.
33. Cats A, Jansen EP, van Grieken NC, Sikorska K, Lind P, Nordmark M, et al. Chemotherapy versus chemoradiotherapy after surgery and preoperative chemotherapy for resectable gastric cancer (critics): an international, open-label, randomised phase 3 trial. *Lancet Oncol* (2018) 19:616–28.
34. de Boer SM, Powell ME, Mileskin L, Katsaros D, Bessette P, Haie-Meder C, et al. Adjuvant chemoradiotherapy versus radiotherapy alone in women with high-risk endometrial cancer (portec-3): patterns of recurrence and *post-hoc* survival analysis of a randomised phase 3 trial. *Lancet Oncol* (2019) 20:1273–85.
35. Kirchheiner K, Nout RA, Lindegaard JC, Haie-Meder C, Mahantshetty U, Segedin B, et al. Dose-effect relationship and risk factors for vaginal stenosis after definitive radio (chemo) therapy with imageguided brachytherapy for locally advanced cervical cancer in the embrace study. *Radiother Oncol* (2016) 118:160–6.
36. Phillips R, Shi WY, Deek M, Radwan N, Lim SJ, Antonarakis ES, et al. Outcomes of observation vs stereotactic ablative radiation for oligometastatic prostate cancer: the oriole phase 2 randomized clinical trial. *JAMA Oncol* (2020) 6:650–9.
37. Amin MB, Greene FL, Edge SB, Compton CC, Gershenwald JE, Brookland RK, et al. The eighth edition AJCC cancer staging manual: continuing to build a bridge from a population-based to a more “personalized” approach to cancer staging. *CA: Cancer J Clin* (2017) 67:93–9.
38. Al-Rashdan A, Cao J. A viewpoint on isolated contralateral axillary lymph node involvement by breast cancer: regional recurrence or distant metastasis? *Int J Radiat Oncol Biol Phys* (2022) 113:489.
39. Tremont A, Lu J, Cole JT. Endocrine therapy for early breast cancer: updated review. *Ochsner J* (2017) 17:405–11.
40. Tchelebi LT. Sowing the seeds: A case of oligometastatic anal cancer 12 years after prostate brachytherapy. *Int J Radiat Oncol Biol Phys* (2022) 114:827–8.
41. Scarpelli D, Jaboin JJ. Exploring the role of resection post-radiation therapy in gliomas. *Int J Radiat Oncol Biol Phys* (2022) 113:11.
42. Prcic M, Murgic J, Frobe A. Radiation therapy for cure or palliation: Case of the immunosuppressed patient with multiple primary cancers and liver transplant. *Int J Radiat Oncol Biol Phys* (2022) 112:581–2.
43. Johnson J, Kudrimoti M, Kaushal A. Synopsis of supraclavicular sarcoma: synthesis of stratagem and solutions. *Int J Radiat Oncol Biol Phys* (2021) 112:4–5. Available at: [https://www.redjournal.org/article/S0360-3016\(21\)00466-1/fulltext](https://www.redjournal.org/article/S0360-3016(21)00466-1/fulltext).
44. Berghen C, De Meeler G. Postoperative radiation therapy in prostate cancer: Timing, duration of hormonal treatment and the use of PSMA PET-CT. *Int J Radiat Oncol Biol Phys* (2022) 113:252–3.
45. Hottinger AF, Pacheco P, Stupp R. Tumor treating fields: a novel treatment modality and its use in brain tumors. *Neuro-Oncol* (2016) 18:1338–49.
46. Goodman CR, Strauss JB. One way or another: An oligorecurrence after an oligometastasis of an estrogen receptor-positive breast cancer. *Int J Radiat Oncol Biol Phys* (2022) 114:571–2.
47. Oertel M, Pepper NB, Schmitz M, Becker JC, Eich HT. Digital transfer in radiation oncology education for medical students—single-center data and systemic review of the literature. *Strahlenther Onkol* (2022) 198:765–72.
48. Shin M, Kim J, van Opheusden B, Griffiths TL. Superhuman artificial intelligence can improve human decision-making by increasing novelty. *PNAS* (2023) 120:e2214840120.
49. Azamfirei R, KudChadkar SR, Fackler J. Large language models and the perils of their hallucinations. *Crit Care* (2023) 27:1–2.
50. Waisberg E, Ong J, Masalkhi M, Kamran SA, Zaman N, Sarker P, et al. GPT-4: a new era of artificial intelligence in medicine. *Ir J Med Sci* (2023), 1–4.
51. Gilbert S, Harvey H, Melvin T, Vollebregt E, Wicks P. Large language model ai chatbots require approval as medical devices. *Nat Med* (2023), 1–3.
52. Touvron H, Martin L, Stone K, Albert P, Almahairi A, Babaei Y, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023) 1–77.
53. Singhal K, Tu T, Gottweis J, Sayres R, Wulczyn E, Hou L, et al. Towards expert-level medical question answering with large language models. *arXiv preprint arXiv:2305.09617* (2023) 1–30.