

Surviving ChatGPT in Healthcare

Zhengliang Liu¹, Lu Zhang², Zihao Wu¹, Xiaowei Yu², Chao Cao², Haixing Dai¹, Ninghao Liu¹, Jun Liu³, Wei Liu⁴, Quanzheng Li⁵, Dinggang Shen^{6, 7, 8}, Xiang Li⁵, Dajiang Zhu², Tianming Liu^{1*}

¹School of Computing, University of Georgia, United States, ²Department of Computer Science and Engineering, College of Engineering, University of Texas at Arlington, United States, ³Department of Radiology, Second Xiangya Hospital, Central South University, China, ⁴Department of Radiation Oncology, Mayo Clinic, United States, ⁵Department of Radiology, Massachusetts General Hospital, Harvard Medical School, United States, ⁶School of Biomedical Engineering, ShanghaiTech University, China, ⁷Shanghai United Imaging Intelligence Co., Ltd., China, ⁸Shanghai Clinical Research and Trial Center, China

Submitted to Journal:
Frontiers in Radiology

Specialty Section:
Artificial Intelligence in Radiology

Article type:
Review Article

Manuscript ID:
1224682

Received on:
18 May 2023

Revised on:
16 Jun 2023

Journal website link:
www.frontiersin.org

In review

Conflict of interest statement

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

Author contribution statement

Conceptualization: ZL, XL, WL, DS and TL. Methodology: ZL, XL and TL. Writing – original draft: ZL, LZ, XY, ZW, CC and HD. Writing – review and editing: NL, JL, XL, WL, DS and TL. Supervision: XL, WL and TL. All authors contributed to the article and approved the submitted version.

Keywords

ChatGPT, large language models (LLM), Artificial General Intelligence (AGI), GPT-4, Artificial Intelligence-AI

Abstract

Word count: 122

At the dawn of Artificial General Intelligence (AGI), the emergence of large language models such as ChatGPT show promise in revolutionizing healthcare by improving patient care, expanding medical access, and optimizing clinical processes. However, their integration into healthcare systems requires careful consideration of potential risks, such as inaccurate medical advice, patient privacy violations, the creation of falsified documents or images, overreliance on AGI in medical education, and the perpetuation of biases. It is crucial to implement proper oversight and regulation to address these risks, ensuring the safe and effective incorporation of AGI technologies into healthcare systems. By acknowledging and mitigating these challenges, AGI can be harnessed to enhance patient care, medical knowledge, and healthcare processes, ultimately benefiting society as a whole.

Contribution to the field

We underscore the importance of considering potential risks associated with the incorporation of AGI into healthcare. These risks include the provision of inaccurate medical advice, potential violation of patient privacy, exploitation of AGI to create falsified documents or images, over-reliance on AGI in medical education, and perpetuation of existing biases present in training data. This investigation holds particular importance for the field of Radiology, which has a profound connection with advancements in AI and machine learning. Radiology, with its reliance on imaging technology and pattern recognition, stands at the forefront of AI applications in healthcare. Therefore, radiology is deeply impacted by the developments in AGI and we believe it is critical to outline potential risks and challenges. We aim to spur broader discussions about the future of medicine in the AGI era.

Surviving ChatGPT in Healthcare

Zhengliang Liu¹, Lu Zhang², Zihao Wu¹, Xiaowei Yu², Chao Cao², Haixing Dai¹, Ninghao Liu¹, Jun Liu³, Wei Liu⁴, Quanzheng Li⁵, Dinggang Shen^{6,7,8}, Xiang Li⁵, Dajiang Zhu², and Tianming Liu¹

¹School of Computing, University of Georgia

²Department of Computer Science and Engineering, The University of Texas at Arlington

³Department of Radiology, Second Xiangya Hospital

⁴Department of Radiation Oncology, Mayo Clinic

⁵Department of Radiology, Massachusetts General Hospital and Harvard Medical School

⁶School of Biomedical Engineering, ShanghaiTech University

⁷Shanghai United Imaging Intelligence Co., Ltd.

⁸Shanghai Clinical Research and Trial Center

Abstract

At the dawn of Artificial General Intelligence (AGI), the emergence of large language models such as ChatGPT show promise in revolutionizing healthcare by improving patient care, expanding medical access, and optimizing clinical processes. However, their integration into healthcare systems requires careful consideration of potential risks, such as inaccurate medical advice, patient privacy violations, the creation of falsified documents or images, overreliance on AGI in medical education, and the perpetuation of biases. It is crucial to implement proper oversight and regulation to address these risks, ensuring the safe and effective incorporation of AGI technologies into healthcare systems. By acknowledging and mitigating these challenges, AGI can be harnessed to enhance patient care, medical knowledge, and healthcare processes, ultimately benefiting society as a whole.

1 Introduction

Large language models (LLM) such as ChatGPT and GPT-4 are making significant strides towards the development of Artificial General Intelligence (AGI) [1, 2]. Advanced generative AGI holds the potential to improve patient care [3, 4] and streamline healthcare processes [5, 3]. However, without proper oversight and regulation, the integration of LLMs into the healthcare system could introduce a range of unintended risks and consequences [6, 4]. It is essential to explore these potential risks and address them effectively to ensure that AGI serves as a beneficial aid in the medical field.

One of the foremost concerns regarding AGI in healthcare is the risk of providing inaccurate medical advice [7, 6]. Since AI-generated content (AIGC) is based on vast amounts of internet data, there is a possibility that the information provided may be misleading or outright incorrect. For instance, ChatGPT might offer treatment suggestions that are either outdated or not suitable for a specific patient's condition. Such inaccuracies could result in patients receiving inappropriate treatments or even exacerbate their health issues.

In addition to the significant concern of inaccurate advice and conclusions, the potential violation of patient privacy is another area of concern [8]. Although AGI systems like ChatGPT are intended

to be securely designed, they may not yet fully comply with privacy regulations like HIPAA [9]. Consequently, patient data may be compromised, leading to the unauthorized access, harvesting and sharing of sensitive personal information. This risk highlights the importance of ensuring AGI technologies adhere to strict privacy standards before their widespread adoption in healthcare.

Powerful generative AGI could also be exploited to create fake documents or images, resulting in misinformation and negative consequences [10]. For example, unscrupulous individuals or companies might use AGI-generated medical images [11] to support false claims or promote unproven treatments, thereby misleading patients and healthcare professionals.

Moreover, the integration of large language models into medical training could have a detrimental effect on the education of future healthcare professionals [12, 13]. Students may rely too heavily on AGI-generated content, neglecting to develop the critical skills required for effective medical practice. By using AGI tools as shortcuts, they might not acquire the ability to differentiate between relevant and irrelevant information, which is crucial in the fast-paced and complex world of healthcare.

Furthermore, established AGI models may inadvertently perpetuate existing biases present in the data used for training [14]. Consequently, the outcomes generated by these models might be biased, leading to the reinforcement of stereotypes and potentially causing harm to certain demographics [15]. For instance, if a model’s training data contains biased information about a specific ethnic group, the model could produce advice that is detrimental to patients from that group.

In short, while AGI models like ChatGPT hold the potential to revolutionize the healthcare sector, it is essential to recognize and address the potential risks they pose. Proper oversight and regulation are necessary to ensure the integration of AGI technologies into the healthcare system is both safe and effective. By addressing these risks, AGI can be harnessed to significantly improve patient care, medical knowledge, and healthcare processes, ultimately benefiting society as a whole.

We divide the ensuing discussion into 6 sections. Each section addresses specific subsets of potential harms and risks posed by large language models such as ChatGPT. We aim to provide a comprehensive yet succinct summary of this subject to spur broader discussion and insights into the future of medicine in the era of AGI. In addition, in section 7, we discuss the potential risks of AGI models (including ChatGPT) in the context of radiology.

2 Truth or dare? Inaccurate medical advice from Dr. ChatGPT

The potential for AGI models, such as ChatGPT, to provide inaccurate medical advice is a critical concern within the medical community [16, 17]. As these systems generate content based on vast amounts of data [2, 5, 1], it is possible that the information provided could be misleading or incorrect [18]. This raises concerns about the trustworthiness of AGI-generated medical advice, as well as the potential risks to patients who rely on such guidance. This discussion has been ongoing since [19] the introduction of smaller pre-trained models (such as BERT or T5 [20]) to healthcare NLP, but it has become a more prominent concern with the rise of LLMs.

For instance, consider a scenario where a patient consults ChatGPT for recommendations on managing their diabetes. ChatGPT might provide advice based on outdated guidelines, which could lead to inappropriate dietary recommendations or incorrect medication dosages. If the patient were to follow this advice, they might experience adverse effects, such as uncontrolled blood sugar levels or complications arising from improper treatment. For example, in 2018, the American College of Physicians (ACP) issued a new recommendation to control type 2 diabetes patients’ A1C level to

between 7% to 8% instead of a previous target of below 7%, since research indicated that reducing the A1C too low through medication did not prevent macrovascular complications yet might lead to substantial harms [21]. If ChatGPT or similar models were trained on outdated standards, their responses might lack current medical validity.

Another example concerns the potential for ChatGPT to generate advice that does not consider individual patient factors. Clinical decision-making often requires a nuanced understanding of a patient’s medical history, concurrent health conditions, and potential contraindications. An AGI model like ChatGPT, without direct access to a patient’s medical records, may generate advice that is unsuitable or even harmful to the patient. For instance, a patient suffering from both hypertension and kidney disease may receive medication advice that is appropriate for managing hypertension but exacerbates their kidney condition [22], since ChatGPT is not aware of the full spectrum of the patient’s problems.

Additionally, the rapidly evolving nature of medical knowledge presents challenges for AGI models. With new research findings continuously emerging, it is vital that AGI models are updated regularly to ensure that their advice remains aligned with the latest evidence-based guidelines. However, the lag between the publication of new research and its integration into AGI models may result in patients receiving advice that is no longer considered best practice. Indeed, both ChatGPT and GPT-4 were only trained on data up to September 2021 [23] and consequently have no knowledge of more recent developments.

In the realm of medical diagnostics, the potential for ChatGPT to misinterpret symptoms or overlook critical information could lead to diagnostic errors. For example, a patient may present with symptoms that align with multiple diagnoses. The model might generate advice based on the most common condition, while failing to consider a rare but more serious alternative diagnosis. Such oversight could have serious consequences for the patient, who may not receive the correct treatment in a timely manner.

Overall, there is a significant risk for large language models to produce ungrounded or unverified medical advice. It is necessary to raise awareness of this new challenge in the AGI landscape. We encourage efforts to instill correct knowledge into models like ChatGPT or establish guardrails that moderate generated medical content.

3 The Emperor’s New Clothes: Privacy concerns

Large language models (LLMs) like ChatGPT offer impressive capabilities [24, 25, 26, 27, 28], but they also come with significant privacy implications that need to be carefully addressed. One particular risk is the potential for models to inadvertently leak details from the data they were trained on. While this is a concern for all LLMs, there are additional challenges if a model trained on private data [8] were to become publicly accessible.

Datasets used to train language models can be substantial [1, 29], often reaching hundreds of gigabytes, and they draw from various sources and domains [30, 31, 32, 33, 34]. Consequently, even when trained on public data, these datasets can contain sensitive information, such as personally identifiable information (PII) including names, phone numbers, and addresses. This raises concerns that a model trained on such data could inadvertently expose private details in its output. It is crucial to identify and minimize the risks of such leaks and develop strategies to address these concerns with future models. This has long been a concern in applying language models to healthcare [35]. Prior to the advent of LLMs, BERT-based models were typically combined with differential

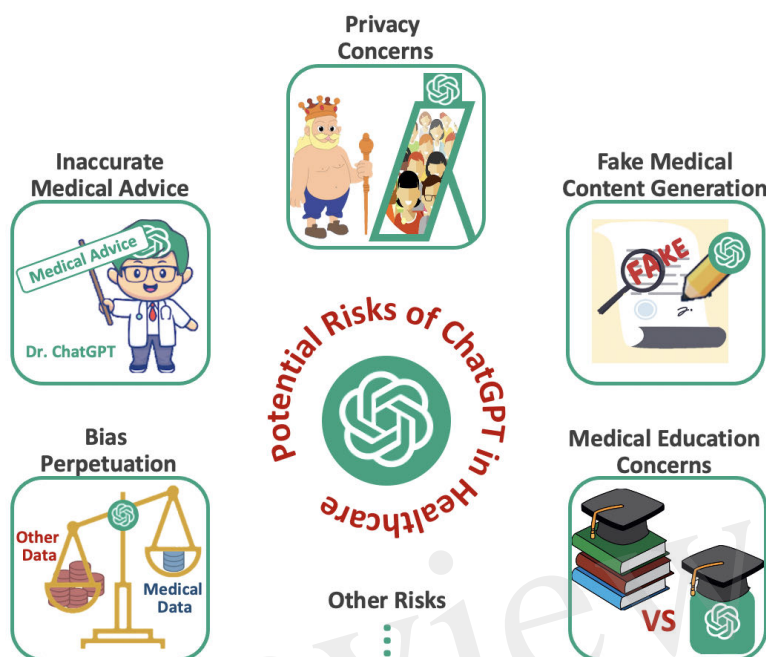


Figure 1: Potential risks of ChatGPT in healthcare.

privacy training and federated learning strategies to better protect privacy in healthcare applications [36].

Privacy and data protection regulation compliance is another significant concern associated with LLMs. These models have the capacity to “memorize” personal information, putting it at risk of being discovered by other users or potential attackers [37, 38]. This ability to retain and potentially reveal personal information calls for robust measures to ensure data privacy and prevent unauthorized access.

The use of LLMs in healthcare [25, 26] has also raised privacy and security concerns, particularly regarding sensitive medical information [39, 40]. Clinical notes, encompassing physician consultations, nursing assessments, lab results, and more, are often stored in free-text formats that may include identifiable or confidential patient information. Unauthorized access to this information poses significant risks to patient confidentiality and privacy. Regulations such as the U.S. Health Insurance Portability and Accountability Act (HIPAA) mandate the removal of re-identifying information from medical records before dissemination to preserve patient confidentiality [41, 42]. Researchers are actively exploring ways to mitigate these concerns by employing data masking techniques to conceal sensitive data and prevent unauthorized access [43, 8].

In conclusion, while LLMs offer impressive capabilities, the associated privacy concerns must be diligently addressed. Minimizing data leaks, protecting against unauthorized access, and implementing privacy-preserving techniques are crucial steps toward ensuring the responsible and ethical use of LLMs in safeguarding patient privacy and maintaining data confidentiality.

4 All that glitters is not gold: Fake medical content generation

Large language models have been increasingly used in fake medical content generation due to their ability to generate text that mimics [10] the writing style and language of medical professionals. These models are based on machine learning algorithms that have been trained on vast amounts of text data, which enables them to generate text that is both coherent and informative. Large language models can also be used to generate fake medical content on a wide range of topics, including diagnosis, treatment, and prognosis. This can be especially useful for medical writers and content creators who need to produce a high volume of content quickly and efficiently.

The use of LLMs for generating fake medical content has become more widespread due to their impressive ability to produce convincing and coherent text that mimics human writing [44]. The potential for LLMs, such as ChatGPT, to generate fake medical content raises several ethical concerns, particularly regarding patient safety and informed consent [45, 46]. The dissemination of fake medical content may lead to false diagnosis, inappropriate treatments, and further medical complications, causing harm to patients. Additionally, the spread of such fake medical content can have serious public health consequences, as it can fuel the promotion of unproven treatments or products that may be ineffective or even dangerous [16].

To address the risks associated with fake medical content generation using LLMs, several strategies have been proposed [10]. One approach is to develop advanced algorithms and tools that can detect fake medical content and prevent its spread. These tools could leverage machine learning and natural language processing techniques to analyze and validate the authenticity of medical content. Another approach is to increase awareness and education among healthcare professionals and the public about the risks associated with fake medical content. By promoting critical thinking and media literacy, individuals can be better equipped to identify fake medical content and make informed decisions about their health.

While LLMs have the potential to transform the content generation of medical data, their misuse for generating fake medical content raises significant ethical concerns. The development of advanced tools to detect and prevent the spread of fake medical content, along with increased awareness and education, can help mitigate the risks associated with the use of LLMs in medical content generation.

5 Veritas vos liberabit? AIGC knowledge compromises medical education

The potential misuse of ChatGPT and similar AGI models in medical education raises significant ethical concerns. It is plausible that medical students and trainees could employ these technologies to complete assignments unethically, misrepresenting their actual knowledge and skills through AGI-generated content (AIGC) [47, 48]. Such a scenario would not only result in a diminished educational experience for these students but could also jeopardize patient care when they transition into professional practice.

Moreover, relying on AGI tools to generate content might undermine students' capacity for learning, problem-solving, and generalization from known examples [49, 50]. The acquisition of these essential skills is paramount for medical professionals, who must be able to navigate complex situations and adapt to new challenges. If AGI-generated content becomes a crutch for students, they may fail to develop the critical thinking abilities necessary for successful medical practice.

By bypassing the rigorous process of learning and self-discovery, students risk hindering their cognitive

growth and reducing their aptitude for medical problem-solving. In the long term, this could lead to a workforce of medical professionals ill-equipped to handle the intricacies of their field, ultimately compromising the quality of healthcare and public trust in the healthcare system.

From another perspective, ChatGPT and similar models offer unique opportunities to democratize medical education that is previously not accessible to the public or medical students in disadvantaged regions. Open medical education offers noticeable benefits in avoiding unwanted treatment [51, 52], making informed decisions [51], promoting effective patient self-management [53] and achieving better clinical outcomes [54]. Indeed, ChatGPT enables any audience to quickly source medical information that is previously inaccessible [55], which has a significant social impact, especially for communities and individuals that benefit from the mass dissemination of medical knowledge.

It is crucial to establish guidelines and policies governing the use of AGI models in medical education to mitigate risks and ensure that their potential benefits are harnessed without sacrificing the integrity of the learning process.

6 One lie leads to another: Bias perpetuation

ChatGPT is a general-purpose LLM that is not specialized for medical problems, even though it can perform better than many models specifically fine-tuned on medical knowledge [56]. It is trained on large amounts of real Internet data, which might not be adequately representative of human diversity and could contain pre-existing bias. According to the latest AI Index Report issued by Stanford University, the probability and danger of bias will develop as the size and capabilities of large language models keep growing [57]. Previous models, such as BERT, also attracted concerns over bias in their training data [58]. Indeed, it is unavoidable for ChatGPT to contain inherent bias. The wide popularity of ChatGPT further exacerbates existing problems.

ChatGPT has shown bias against specific groups since its training data contains racial and sexist stereotypes [59]. The issues of underrepresentation, overrepresentation, and misrepresentation in web data are likely introducing various kinds of bias into ChatGPT. The outputs containing bias typically manifest in nuanced representations, which makes it difficult to recognize and correct bias and toxicity [60, 61, 62].

Also, fine-tuning ChatGPT on historical medical data has the potential to introduce or exacerbate bias that already exists within the data. For example, clinical practice biases, such as under-testing of marginalized communities, can impact the underlying clinical data and introduce bias during future training [63].

In addition, implicit bias from healthcare professionals can manifest in clinical notes, including segments of diagnoses and treatment decisions [64]. ChatGPT and GPT-4 might introduce these new biases into downstream applications if such notes are used for training.

OpenAI has released plugins for images and will certainly develop multimodal foundational models in the future (e.g., GPT-4 will have capabilities to process images in the near future). But fairness research indicates that the combination with additional information or modality may not necessarily improve performance and is likely to bring about new unfairness and bias [65]. For example, in CLIP, a language-vision model, historical race and gender bias are reinforced [66].

Fostering research and development to detect, mitigate, and prevent bias, toxicity, and other undesirable behaviors in large language models like ChatGPT and GPT-4 are crucial for a responsible AI future. By actively pursuing these objectives, we can ensure that these powerful tools serve as

inclusive, unbiased, and beneficial resources for users across diverse backgrounds. This pursuit not only safeguards the ethical foundations of AI but also greatly enhances its potential to positively impact society.

7 Potential Risks of Artificial General Intelligence (AGI) Models in Radiology

Inaccurate Medical Advice: A substantial risk associated with AGI models like ChatGPT involves the provision of inaccurate medical advice, particularly in the nuanced field of radiology. For instance, when interpreting a radiology report indicating the presence of a small pulmonary nodule, the model might suggest watchful waiting based on outdated guidelines, overlooking recent research that indicates a higher malignancy risk requiring more active intervention.

Privacy Concerns: Privacy is a paramount concern in the use of AGI models. For example, patients seeking guidance for understanding a radiology report could unwittingly disclose sensitive health data like a past diagnosis of breast cancer or a family history of genetic disorders. Potential security vulnerabilities in the wider AGI system could expose this sensitive data to misuse or unauthorized access.

Fake Medical Content Generation: The potential for AGI models to generate misleading or false medical content presents a significant risk. An individual with malevolent intent could misuse AGI tools to fabricate radiology reports, falsely indicating the presence or absence of a medical condition, such as fabricating a report showing a clean bill of health when the actual scan revealed lung nodules indicative of early-stage cancer.

Compromise of Medical Education: AGI models have the potential to unintentionally undermine the quality of medical education. For example, a medical trainee could become overly reliant on a multimodal AGI model for interpreting brain MRI scans for assignments. This over-reliance could deprive them of crucial learning experiences and inadvertently foster an environment of plagiarism and excessive dependence on automation. Consequently, this may lead to the undesired outcome of inadequately trained professionals tasked with handling complex or ambiguous clinical situations.

Bias Perpetuation: Lastly, AGI models can unintentionally propagate existing biases in medical data. For example, if the data used to train ChatGPT and GPT-4 over-represents Caucasian individuals, the models might be less adept at interpreting radiology reports concerning conditions more prevalent in other ethnic groups, such as the higher incidence of sarcoidosis in African American populations [67].

8 Conclusion

In conclusion, while the emergence of large language models such as ChatGPT offers promising prospects for revolutionizing healthcare, addressing the potential risks and challenges is of paramount importance. Future research should focus on developing robust methods to ensure the accuracy, reliability, and privacy compliance of medical contents generated by AGI, as well as monitoring and mitigating the biases that may be introduced during the training of these models. Additionally, guidelines and regulations should be established to govern the use of AGI models in medical education, promoting their responsible use and preserving the integrity of the learning process.

By acknowledging and addressing these challenges, Artificial General Intelligence can be harnessed to

revolutionize patient care and healthcare, ultimately benefiting society as a whole and significantly promoting national health. The interdisciplinary collaboration between AGI researchers, medical professionals, ethicists, and policy-makers will play a crucial role in shaping the future of medicine in the era of AGI, ensuring its safe and effective integration into healthcare systems worldwide.

References

- [1] Yiheng Liu, Tianle Han, Siyuan Ma, Jiayue Zhang, Yuanyuan Yang, Jiaming Tian, Hao He, Antong Li, Mengshen He, Zhengliang Liu, et al. Summary of chatgpt/gpt-4 research and perspective towards the future of large language models. *arXiv preprint arXiv:2304.01852*, 2023.
- [2] Lin Zhao, Lu Zhang, Zihao Wu, Yuzhong Chen, Haixing Dai, Xiaowei Yu, Zhengliang Liu, Tuo Zhang, Xintao Hu, Xi Jiang, et al. When brain-inspired ai meets agi. *arXiv preprint arXiv:2303.15935*, 2023.
- [3] Jianning Li, Amin Dada, Jens Kleesiek, and Jan Egger. Chatgpt in healthcare: A taxonomy and systematic review. *medRxiv*, pages 2023–03, 2023.
- [4] James H Thrall, Xiang Li, Quanzheng Li, Cinthia Cruz, Synho Do, Keith Dreyer, and James Brink. Artificial intelligence and machine learning in radiology: opportunities, challenges, pitfalls, and criteria for success. *Journal of the American College of Radiology*, 15(3):504–508, 2018.
- [5] Zhengliang Liu, Mengshen He, Zuwei Jiang, Zihao Wu, Haixing Dai, Lian Zhang, Siyi Luo, Tianle Han, Xiang Li, Xi Jiang, et al. Survey on natural language processing in medical image analysis. *Zhong nan da xue xue bao. Yi xue ban= Journal of Central South University. Medical Sciences*, 47(8):981–993, 2022.
- [6] Malik Sallam. Chatgpt utility in healthcare education, research, and practice: Systematic review on the promising perspectives and valid concerns. In *Healthcare*, volume 11, page 887. MDPI, 2023.
- [7] Claudia E Haupt and Mason Marks. Ai-generated medical advice—gpt and beyond. *JAMA*, 329(16):1349–1350, 2023.
- [8] Zhengliang Liu, Xiaowei Yu, Lu Zhang, Zihao Wu, Chao Cao, Haixing Dai, Lin Zhao, Wei Liu, Dinggang Shen, Quanzheng Li, et al. Deid-gpt: Zero-shot medical text de-identification by gpt-4. *arXiv preprint arXiv:2303.11032*, 2023.
- [9] Raju Vaishya, Anoop Misra, and Abhishek Vaish. Chatgpt: Is this version good for healthcare and research? *Diabetes & Metabolic Syndrome: Clinical Research & Reviews*, 17(4):102744, 2023.
- [10] Wenxiong Liao, Zhengliang Liu, Haixing Dai, Shaochen Xu, Zihao Wu, Yiyang Zhang, Xiaoke Huang, Dajiang Zhu, Hongmin Cai, Tianming Liu, et al. Differentiate chatgpt-generated and human-written medical texts. *arXiv preprint arXiv:2304.11567*, 2023.
- [11] Neal Mangaokar, Jiameng Pu, Parantapa Bhattacharya, Chandan K Reddy, and Bimal Viswanath. Jekyll: Attacking medical image diagnostics using deep generative models. In *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 139–157. IEEE, 2020.

-
- [12] Andrej Thurzo, Martin Strunga, Renáta Urban, Jana Surovková, and Kelvin I Afrashtehfar. Impact of artificial intelligence on dental education: A review and guide for curriculum update. *Education Sciences*, 13(2):150, 2023.
- [13] Mohammad Hosseini, Catherine A Gao, David M Liebovitz, Alexandre M Carvalho, Faraz S Ahmad, Yuan Luo, Ngan MacDonald, Kristi L Holmes, and Abel Kho. An exploratory survey about using chatgpt in education, healthcare, and research. *medRxiv*, pages 2023–03, 2023.
- [14] Brady D Lund, Ting Wang, Nishith Reddy Mannuru, Bing Nie, Somipam Shimray, and Ziang Wang. Chatgpt and a new academic reality: Artificial intelligence-written research papers and the ethics of the large language models in scholarly publishing. *Journal of the Association for Information Science and Technology*, 2023.
- [15] Emilio Ferrara. Should chatgpt be biased? challenges and risks of bias in large language models. *arXiv preprint arXiv:2304.03738*, 2023.
- [16] Yiqiu Shen, Laura Heacock, Jonathan Elias, Keith D Hentel, Beatriu Reig, George Shih, and Linda Moy. Chatgpt and other large language models are double-edged swords, 2023.
- [17] Florin Eggmann, Roland Weiger, Nicola U Zitzmann, and Markus B Blatz. Implications of large language models such as chatgpt for dental medicine. *Journal of Esthetic and Restorative Dentistry*, 2023.
- [18] Hyunsu Lee. The rise of chatgpt: Exploring its potential in medical education. *Anatomical Sciences Education*, 2023.
- [19] Regina Merine and Saptarshi Purkayastha. Risks and benefits of ai-generated text summarization for expert level content in graduate health informatics. In *2022 IEEE 10th International Conference on Healthcare Informatics (ICHI)*, pages 567–574. IEEE, 2022.
- [20] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020.
- [21] Gerald W Smetana, David M Nathan, David C Dugdale, and Risa B Burns. To what target hemoglobin a1c level would you treat this patient with type 2 diabetes? grand rounds discussion from beth israel deaconess medical center. *Annals of Internal Medicine*, 171(7):505–513, 2019.
- [22] Dan Pugh, Peter J Gallacher, and Neeraj Dhaun. Management of hypertension in chronic kidney disease. *Drugs*, 79:365–379, 2019.
- [23] Alejandro Lopez-Lira and Yuehua Tang. Can chatgpt forecast stock price movements? return predictability and large language models. *arXiv preprint arXiv:2304.07619*, 2023.
- [24] Jason Holmes, Zhengliang Liu, Lian Zhang, Yuzhen Ding, Terence T Sio, Lisa A McGee, Jonathan B Ashman, Xiang Li, Tianming Liu, Jiajian Shen, et al. Evaluating large language models on a highly-specialized topic, radiation oncology physics. *arXiv preprint arXiv:2304.01938*, 2023.
- [25] Chong Ma, Zihao Wu, Jiaqi Wang, Shaochen Xu, Yaonai Wei, Zhengliang Liu, Lei Guo, Xiaoyan Cai, Shu Zhang, Tuo Zhang, et al. Impressiongpt: An iterative optimizing framework for radiology report summarization with chatgpt. *arXiv preprint arXiv:2304.08448*, 2023.
-

-
- [26] Zihao Wu, Lu Zhang, Chao Cao, Xiaowei Yu, Haixing Dai, Chong Ma, Zhengliang Liu, Lin Zhao, Gang Li, Wei Liu, et al. Exploring the trade-offs: Unified large language models vs local fine-tuned models for highly-specific radiology nli task. *arXiv preprint arXiv:2304.09138*, 2023.
- [27] Tianyang Zhong, Yaonai Wei, Li Yang, Zihao Wu, Zhengliang Liu, Xiaozheng Wei, Wenjun Li, Junjie Yao, Chong Ma, Xiang Li, et al. Chatabl: Abductive learning via natural language interaction with chatgpt. *arXiv preprint arXiv:2304.11107*, 2023.
- [28] Haixing Dai, Zhengliang Liu, Wenxiong Liao, Xiaoke Huang, Zihao Wu, Lin Zhao, Wei Liu, Ninghao Liu, Sheng Li, Dajiang Zhu, et al. Chataug: Leveraging chatgpt for text data augmentation. *arXiv preprint arXiv:2302.13007*, 2023.
- [29] Wenxiong Liao, Zhengliang Liu, Haixing Dai, Zihao Wu, Yiyang Zhang, Xiaoke Huang, Yuzhong Chen, Xi Jiang, Dajiang Zhu, Tianming Liu, et al. Mask-guided bert for few shot text classification. *arXiv preprint arXiv:2302.10447*, 2023.
- [30] Saed Rezayi, Haixing Dai, Zhengliang Liu, Zihao Wu, Akarsh Hebbar, Andrew H Burns, Lin Zhao, Dajiang Zhu, Quanzheng Li, Wei Liu, et al. Clinicalradiobert: Knowledge-infused few shot learning for clinical notes named entity recognition. In *Machine Learning in Medical Imaging: 13th International Workshop, MLMI 2022, Held in Conjunction with MICCAI 2022, Singapore, September 18, 2022, Proceedings*, pages 269–278. Springer, 2022.
- [31] Saed Rezayi, Zhengliang Liu, Zihao Wu, Chandra Dhakal, Bao Ge, Chen Zhen, Tianming Liu, and Sheng Li. Agribert: Knowledge-infused agricultural language models for matching food and nutrition. In Lud De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 5150–5156. International Joint Conferences on Artificial Intelligence Organization, 7 2022. AI for Good.
- [32] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- [33] Homgmin Cai, Wenxiong Liao, Zhengliang Liu, Xiaoke Huang, Yiyang Zhang, Siqi Ding, Sheng Li, Quanzheng Li, Tianming Liu, and Xiang Li. Coarse-to-fine knowledge graph domain adaptation based on distantly-supervised iterative training. *arXiv preprint arXiv:2211.02849*, 2022.
- [34] Zhengliang Liu, Xinyu He, Lei Liu, Tianming Liu, and Xiaoming Zhai. Context matters: A strategy to pre-train language model for science education. *arXiv preprint arXiv:2301.12031*, 2023.
- [35] Thomas Vakili and Hercules Dalianis. Are clinical bert models privacy preserving? the difficulty of extracting patient-condition associations. In *HUMAN@ AAAI Fall Symposium*, 2021.
- [36] Priyam Basu, Tiasa Singha Roy, Rakshit Naidu, Zumrut Muftuoglu, Sahib Singh, and Fate-mehsadat Miresghallah. Benchmarking differential privacy and federated learning for bert models. *arXiv preprint arXiv:2106.13973*, 2021.
- [37] Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom B Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *USENIX Security Symposium*, volume 6, 2021.
- [38] Jiayuan Ye, Aadyaa Maddi, Sasi Kumar Murakonda, Vincent Bindschaedler, and Reza Shokri. Enhanced membership inference attacks against machine learning models. In *Proceedings of the*
-

2022 ACM SIGSAC Conference on Computer and Communications Security, pages 3093–3106, 2022.

- [39] Maryam Tayefi, Phuong Ngo, Taridzo Chomutare, Hercules Dalianis, Elisa Salvi, Andrius Budrionis, and Fred Godtlielsen. Challenges and opportunities beyond structured data in analysis of electronic health records. *Wiley Interdisciplinary Reviews: Computational Statistics*, 13(6):e1549, 2021.
- [40] Jay Urbain, George Kowalski, Kristen Osinski, Robert Spaniol, Mei Liu, Bradley Taylor, Lemuel R Waitman, et al. Natural language processing for enterprise-scale de-identification of protected health information in clinical notes. In *AMIA Annual Symposium Proceedings*, volume 2022, page 92. American Medical Informatics Association, 2022.
- [41] HHS Office for Civil Rights. Standards for privacy of individually identifiable health information. final rule. *Federal register*, 67(157):53181–53273, 2002.
- [42] Tanbir Ahmed, Md Momin Al Aziz, and Noman Mohammed. De-identification of electronic health record using neural network. *Scientific reports*, 10(1):1–11, 2020.
- [43] N Sharma, Ashima Anand, and Amit Kumar Singh. Bio-signal data sharing security through watermarking: a technical survey. *Computing*, pages 1–35, 2021.
- [44] Jocelyn Gravel, Madeleine D’Amours-Gravel, and Esli Osmanliu. Learning to fake it: limited responses and fabricated references provided by chatgpt for medical questions, 2023.
- [45] Luigi De Angelis, Francesco Baglivo, Guglielmo Arzilli, Gaetano Pierpaolo Privitera, Paolo Ferragina, Alberto Eugenio Tozzi, and Caterina Rizzo. Chatgpt and the rise of large language models: The new ai-driven infodemic threat in public health. *Frontiers in Public Health*, 2023.
- [46] Sajan B. Patel, Kyle Lam, and Michael Liebreuz. Chatgpt: Friend or foe. *Lancet Digit. Health*, 5, 2023.
- [47] Michael R King and ChatGPT. A conversation on artificial intelligence, chatbots, and plagiarism in higher education. *Cellular and Molecular Bioengineering*, 16(1):1–2, 2023.
- [48] Rehan Ahmed Khan, Masood Jawaid, Aymen Rehan Khan, and Madiha Sajjad. Chatgpt-reshaping medical education and clinical management. *Pakistan Journal of Medical Sciences*, 39(2):605, 2023.
- [49] Ahmed Tlili, Boulus Shehata, Michael Agyemang Adarkwah, Aras Bozkurt, Daniel T Hickey, Ronghuai Huang, and Brighter Agyemang. What if the devil is my guardian angel: Chatgpt as a case study of using chatbots in education. *Smart Learning Environments*, 10(1):15, 2023.
- [50] Mohammadreza Farrokhnia, Seyyed Kazem Banihashem, Omid Noroozi, and Arjen Wals. A swot analysis of chatgpt: Implications for educational practice and research. *Innovations in Education and Teaching International*, pages 1–15, 2023.
- [51] Yvonne Chan, Jonathan C Irish, Stephen J Wood, Lorne E Rotstein, Dale H Brown, Patrick J Gullane, and Gina A Lockwood. Patient education and informed consent in head and neck surgery. *Archives of Otolaryngology–Head & Neck Surgery*, 128(11):1269–1274, 2002.
- [52] Cara Tannenbaum, Philippe Martin, Robyn Tamblyn, Andrea Benedetti, and Sara Ahmed. Reduction of inappropriate benzodiazepine prescriptions among older adults through direct patient education: the empower cluster randomized trial. *JAMA internal medicine*, 174(6):890–898, 2014.

-
- [53] Richard H Osborne, Jean M Spinks, and Ian P Wicks. Patient education and self-management programs in arthritis. *Medical journal of Australia*, 180(5):S23, 2004.
- [54] Deanna J Attai, Michael S Cowher, Mohammed Al-Hamadani, Jody M Schoger, Alicia C Staley, and Jeffrey Landercasper. Twitter social media is an effective tool for breast cancer patient education and support: patient-reported outcomes by survey. *Journal of medical Internet research*, 17(7):e188, 2015.
- [55] Som Biswas. Chatgpt and the future of medical writing, 2023.
- [56] Harsha Nori, Nicholas King, Scott Mayer McKinney, Dean Carignan, and Eric Horvitz. Capabilities of gpt-4 on medical challenge problems. *arXiv preprint arXiv:2303.13375*, 2023.
- [57] D Zhang, Nestor Maslej, Erik Brynjolfsson, John Etchemendy, Terah Lyons, James Manyika, Helen Ngo, Juan Carlos Niebles, Michael Sellitto, Ellie Sakhaee, et al. The ai index 2022 annual report. ai index steering committee. *Stanford Institute for Human-Centered AI*, 2022.
- [58] Hammaad Adam, Aparna Balagopalan, Emily Alsentzer, Fotini Christia, and Marzyeh Ghassemi. Mitigating the impact of biased artificial intelligence in emergency decision-making. *Communications Medicine*, 2(1):149, 2022.
- [59] Ameet Deshpande, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, and Karthik Narasimhan. Toxicity in chatgpt: Analyzing persona-assigned language models, 2023.
- [60] Malik Sallam. The utility of chatgpt as an example of large language models in healthcare education, research and practice: Systematic review on the future perspectives and potential limitations. *medRxiv*, pages 2023–02, 2023.
- [61] Calvin D Lawrence. *Hidden in White Sight: How AI Empowers and Deepens Systemic Racism*. CRC Press, 2023.
- [62] David A Asch. An interview with chatgpt about health care. *NEJM Catalyst Innovations in Care Delivery*, 4(2), 2023.
- [63] Jenna Wiens, W Nicholson Price, and Michael W Sjoding. Diagnosing bias in data-driven algorithms for healthcare. *Nature Medicine*, 26(1):25–26, 2020.
- [64] Chloë FitzGerald and Samia Hurst. Implicit bias in healthcare professionals: a systematic review. *BMC medical ethics*, 18(1):1–18, 2017.
- [65] Can Cui, Haichun Yang, Yaohong Wang, Shilin Zhao, Zuhayr Asad, Lori A. Coburn, Keith T. Wilson, Bennett A. Landman, and Yuankai Huo. Deep multi-modal fusion of image and non-image data in disease diagnosis and prognosis: A review, 2023.
- [66] Abeba Birhane, Vinay Uday Prabhu, and Emmanuel Kahembwe. Multimodal datasets: misogyny, pornography, and malignant stereotypes, 2021.
- [67] Mehdi Mirsaeidi, Roberto F Machado, Dean Schraufnagel, Nadera J Sweiss, and Robert P Baughman. Racial difference in sarcoidosis mortality in the united states. *Chest*, 147(2):438–449, 2015.