

ChatGPT vs. Scholars: A comparative examination of literature reviews conducting by humans and AI

Mehrnaz Mostafapour, Jacqueline H Fortier, Karen Pacheco, Heather Murray,
Gary Garber

Submitted to: JMIR AI
on: January 18, 2024

Disclaimer: © The authors. All rights reserved. This is a privileged document currently under peer-review/community review. Authors have provided JMIR Publications with an exclusive license to publish this preprint on its website for review purposes only. While the final peer-reviewed paper may be licensed under a CC BY license on publication, at this stage authors and publisher expressly prohibit redistribution of this draft paper other than for review purposes.

Table of Contents

Original Manuscript..... 5

Supplementary Files..... 19

0..... 19

Multimedia Appendixes 20

Multimedia Appendix 4..... 20

ChatGPT vs. Scholars: A comparative examination of literature reviews conducting by humans and AI

Mehrnaz Mostafapour¹ PhD; Jacqueline H Fortier¹ MSc; Karen Pacheco¹ MSc; Heather Murray^{1,2} MD, MSc; Gary Garber^{1,3,4} MD, FRCPC

¹Canadian Medical Protective Association Ottawa CA

²Department of Emergency Medicine, Queen's University Kingston CA

³Ottawa Hospital Research Institute Ottawa CA

⁴Faculty of Medicine, Department of Medicine and the School of Public Health and Epidemiology, University of Ottawa Ottawa CA

Corresponding Author:

Gary Garber MD, FRCPC

Canadian Medical Protective Association

875 Carling Ave.

Ottawa

CA

Abstract

Background: With the rapid evolution of artificial intelligence (AI), particularly large language models (LLMs) like ChatGPT-4, there is increasing interest in their potential to assist in scholarly tasks, including conducting literature reviews. However, the efficacy of AI-generated reviews compared to traditional human-led approaches remains underexplored.

Objective: This study aims to compare the quality of literature reviews conducted by OpenAI's ChatGPT-4 model with those conducted by human researchers, focusing on the relational dynamics between physicians and patients.

Methods: The study included two literature reviews on the same topic namely, exploring factors affecting relational dynamics between physicians and patients in medico-legal contexts. One review used OpenAI's GPT-4, last updated in September 2021, and the other was conducted by human researchers. The human review involved a comprehensive literature search using medical subject headings and keywords in Ovid MEDLINE followed with a thematic analysis of the literature to synthesize information from selected articles. The AI-generated review employed a new prompt engineering approach, using iterative and sequential prompts to generate results. Comparative analysis was based on qualitative measures such as accuracy, response time, consistency, breadth and depth of knowledge, contextual understanding, and transparency.

Results: GPT-4 produced an extensive list of relational factors rapidly. The AI model demonstrated an impressive breadth of knowledge but exhibited limitations in depth and contextual understanding, occasionally producing irrelevant or incorrect information. In comparison, human researchers provided a more nuanced and contextually relevant review. The comparative analysis assessed the reviews based on criteria including accuracy, response time, consistency, breadth and depth of knowledge, contextual understanding, and transparency. While GPT-4 showed advantages in response time and breadth of knowledge, human-led reviews excelled in accuracy, depth of knowledge, and contextual understanding.

Conclusions: The study suggests that GPT-4, with structured prompt engineering, can be a valuable tool for conducting preliminary literature reviews by providing a broad overview of topics quickly. However, its limitations necessitate careful expert evaluation and refinement, making it an assistant rather than a substitute for human expertise in comprehensive literature reviews. Moreover, this research highlights the potential and limitations of using AI tools like GPT-4 in academic research, particularly in the fields of health services and medical research. It underscores the necessity of combining AI's rapid information retrieval capabilities with human expertise for more accurate and contextually rich scholarly outputs.

(JMIR Preprints 18/01/2024:56537)

DOI: <https://doi.org/10.2196/preprints.56537>

Preprint Settings

1) Would you like to publish your submitted manuscript as preprint?

Please make my preprint PDF available to anyone at any time (recommended).

Please make my preprint PDF available only to logged-in users; I understand that my title and abstract will remain visible to all users.

Only make the preprint title and abstract visible.

✓ **No, I do not wish to publish my submitted manuscript as a preprint.**

2) If accepted for publication in a JMIR journal, would you like the PDF to be visible to the public?

✓ **Yes, please make my accepted manuscript PDF available to anyone at any time (Recommended).**

Yes, but please make my accepted manuscript PDF available only to logged-in users; I understand that the title and abstract will remain visible to all users.

Yes, but only make the title and abstract visible (see Important note, above). I understand that if I later pay to participate in <http://www.jmir.org/preprint/56537>



Original Manuscript

ChatGPT vs. Scholars: A comparative examination of literature reviews conducting by humans and AI

Abstract

Background: With the rapid evolution of artificial intelligence (AI), particularly large language models (LLMs) like ChatGPT-4, there is increasing interest in their potential to assist in scholarly tasks, including conducting literature reviews. However, the efficacy of AI-generated reviews compared to traditional human-led approaches remains underexplored.

Objectives: This study aims to compare the quality of literature reviews conducted by OpenAI's ChatGPT-4 model with those conducted by human researchers, focusing on the relational dynamics between physicians and patients.

Methods: The study included two literature reviews on the same topic namely, exploring factors affecting relational dynamics between physicians and patients in medico-legal contexts. One review used OpenAI's GPT-4, last updated in September 2021, and the other was conducted by human researchers. The human review involved a comprehensive literature search using medical subject headings and keywords in Ovid MEDLINE followed with a thematic analysis of the literature to synthesize information from selected articles. The AI-generated review employed a new prompt engineering approach, using iterative and sequential prompts to generate results. Comparative analysis was based on qualitative measures such as accuracy, response time, consistency, breadth and depth of knowledge, contextual understanding, and transparency.

Results: GPT-4 produced an extensive list of relational factors rapidly. The AI model demonstrated an impressive breadth of knowledge but exhibited limitations in depth and contextual understanding, occasionally producing irrelevant or incorrect information. In comparison, human researchers provided a more nuanced and contextually relevant review. The comparative analysis assessed the reviews based on criteria including accuracy, response time, consistency, breadth and depth of knowledge, contextual understanding, and transparency. While GPT-4 showed advantages in response time and breadth of knowledge, human-led reviews excelled in accuracy, depth of knowledge, and contextual understanding.

Conclusions: The study suggests that GPT-4, with structured prompt engineering, can be a valuable tool for conducting preliminary literature reviews by providing a broad overview of topics quickly. However, its limitations necessitate careful expert evaluation and refinement, making it an assistant rather than a substitute for human expertise in comprehensive literature reviews. Moreover, this research highlights the potential and limitations of using AI tools like GPT-4 in academic research, particularly in the fields of health services and medical research. It underscores the necessity of combining AI's rapid information retrieval capabilities with human expertise for more accurate and contextually rich scholarly outputs.

Keywords: OpenAIs; Chat GPT; AI vs. Human; Literature search; Chat GPT performance evaluation

Introduction

Artificial intelligence (AI) is a rapidly evolving technology that combines computer programming with large datasets to enable software to perform tasks. Generative AI uses this technology to synthesize content; the system is trained on large volumes of data to identify patterns until it can recognize those patterns and generate novel responses to queries. Large language models (LLM), such as ChatGPT, are a form of generative AI wherein the software is trained on extensive textual datasets and can generate a response to prompts and questions(1).

AI in general and LLMs in particular are in a period of exponential growth, and researchers are exploring their utility to perform tasks with variable results(1, 2, 3, 4, 5). Previous studies have shown how these tools can help to advance research (4); one area where there is potential to realize efficiencies is in the creation of literature reviews and syntheses. The pace of scientific publication has been rapidly expanding,(6) and AI tools may provide a useful starting point and substantial time savings by automating some elements of a literature search. However, there is little research that compares the results generated using AI with results generated by skilled human researchers.

The purpose of this study is to conduct a literature review using OpenAI's ChatGPT-4 model ("GPT-4") and then conduct a comparative analysis against a review conducted by human researchers.

The way researchers utilize these tools and optimize the language used to generate a response from ChatGPT, known as prompt engineering, directly impacts the quality of results.(7) Clear, concise, neutral, structured, and specific prompts reduce the model's tendency to respond with generic or off-topic responses, as well as generating an unsubstantiated/false response also termed an AI hallucination.(8) Therefore, in order to conduct this study, we have introduced an approach to prompt engineering that may assist researchers who wish to use GPT-4 or other LLMs to generate literature reviews.

Method

We started with a completed literature review exploring the factors influencing the relational dynamics between physician-patient that motivate patients to file medico-legal complaints against physicians (citation is removed due to blind review but will be added). Using this review as a reference standard, we then tasked Open AI's GPT-4 model (training data updated in September 2021) with producing a literature review on the same topic. Subsequently, we compared the results generated by GPT-4 and the literature review conducted by human experts. It should be noted that while GPT-4 was used to generate a literature review and make suggestions for the manuscript title, it was not used to write this manuscript.

Human literature review

The first author conducted a traditional literature review to identify what factors affect relationships between physicians and patients. They utilized a systematic approach to ensure transparency and reproducibility. The review included a mix of studies, and assessed both qualitative and quantitative data together through thematic analysis (9). With the help of a research librarian, they developed a search strategy using medical subject headings (MeSH terms), keywords and key phrases for a single database (Ovid MEDLINE) to identify articles related to physician-patient relationships. The search strategy was calibrated to identify articles that were most relevant to the research question, rather than prioritizing an approach that would capture every potentially relevant paper (Please see Appendix D).

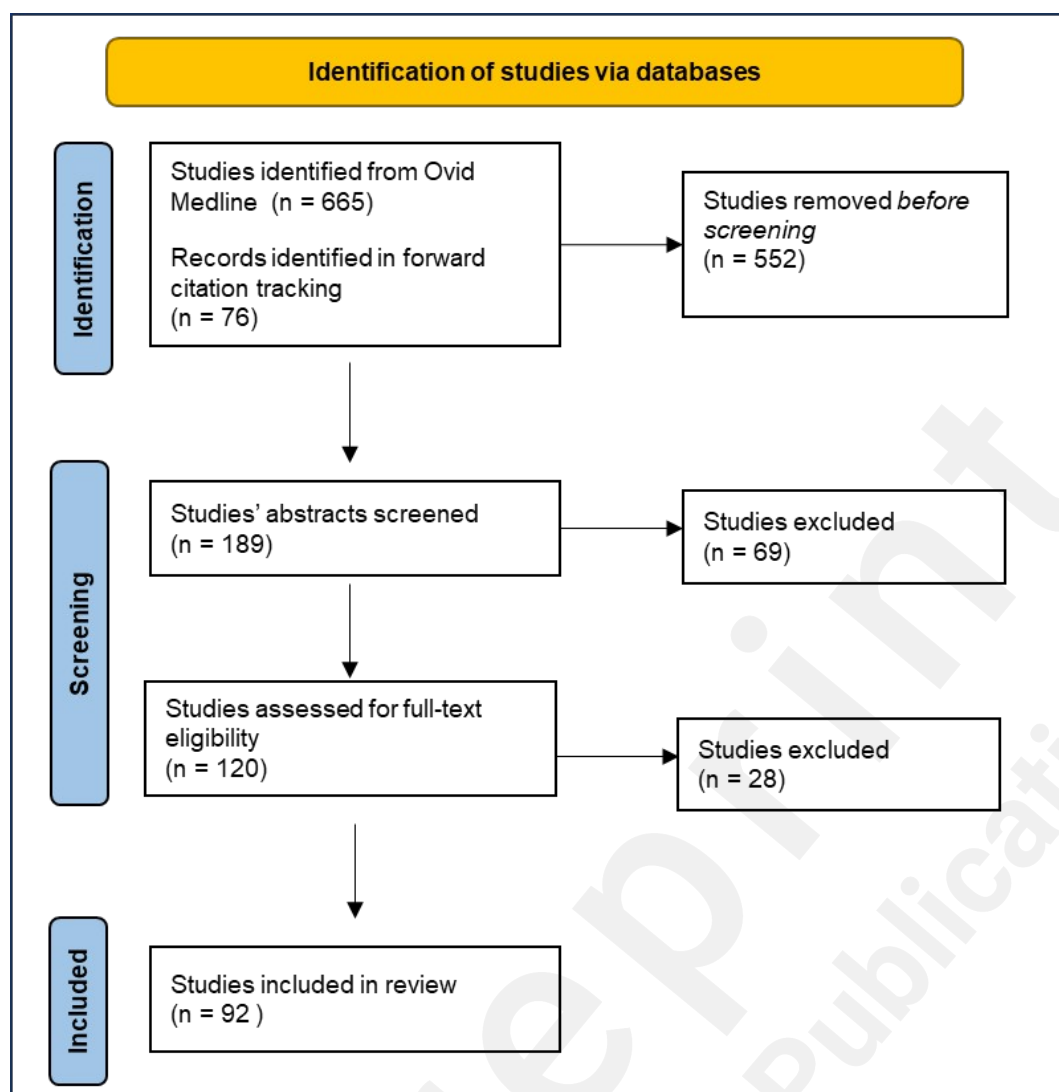
Subsequently, the librarian screened titles and abstracts, and then the main author screened full text articles for inclusion against pre-defined eligibility criteria. Articles had to be empirical research studies or literature reviews that discussed relational factors between physicians and patients that affected patient satisfaction and medico-legal complaints. Studies were excluded if they were not based on empirical research (e.g. editorials, commentaries, reports) or if they were unrelated to the research question (Please see Table 1 and Figure 1)

Table 1 displays the eligibility criteria to identify relevant studies.

Inclusion criteria	<hr/> <ul style="list-style-type: none">The study described an empirical research study or a literature reviewThe study focused on or described relational factors between physicians and patients impacting patients' satisfactionThe record focused on or described the relationship between patients' satisfaction and medico-legal risk against physiciansThe study focused on or described medico-legal complaints against physicians caused by relational problems between patients and physicians <hr/>
Exclusion criteria	<hr/> <ul style="list-style-type: none">The study was not empirical research, e.g., editorials, commentaries, reports, etc.The study contained no explicit mention of physicians and patients relationshipThe study was not related to either patients' satisfaction or medico-legal risk against physicians <hr/>

We used a thematic analysis approach to review and synthesize the included manuscripts to identify the relational factors that influence patient satisfaction and medico-legal complaints and reported the findings in a published narrative review (citation is removed due to the blind review process)..

Figure 1. Overview of article screening and inclusion into the study



AI prompt engineering for the literature search

Previous work suggested that the use of single prompts may not be very effective for complex tasks. (2) We began our process with a single prompt for the literature search (detailed in Appendix A) and the results were clearly inadequate, confirming these findings. Consequently, we developed a series of prompts in an iterative sequential format. This approach operated on the premise that GPT-4 would benefit from incremental and interactive guidance to yield optimal results. In this approach, the researcher designed sequential prompts based on the assessment of the previous responses generated by GPT-4, starting from a general prompt, and designing subsequent prompts to refine the output towards the desired form.

An initial series of prompts were used to explore GPT-4's breadth of knowledge about the factors impacting relationships between physicians and patients. The first prompt was general, simple, and short, asking GPT-4 to list relevant factors to the subject matter. Since we did not know all the relevant factors related to the topic, subsequent prompts were designed to ask for more factors to extend the list of factors and reinforce previous instructions, while specifying the desired tone. Further prompts extended the factors and ensured the validity of their content by introducing additional criteria such as the number of sentences and asking for precise references. (Please see Table 2 & Appendix B)

Table 2

Step	Prompts
1	Write a literature review on the relational problems between physicians and patients that lead to medico-legal complaints against physicians - from a health service researcher perspective - and provide precise references
2	Please add at least 15 more factors related to relational problems between physicians and patients that lead to medico-legal complaints against physicians to the list considering the sensitivity, precision, and accuracy of information
3	You provided me with 21 relational factors between physicians and patients that contribute to the likelihood of filing a medico-legal case against a physician. Please write an elaborated, scientific and accurate description for each factor that include at least 15 sentences, and provide at least two real and precise references that support your arguments

In order to replicate the format of the literature review done by human researchers, the researcher who had conducted the literature review explored a series of prompts to guide GPT-4 through a more in-depth exploration of the identified relational factors. They began by prompting GPT to suggest evidence-based ways to improve each relational factor (e.g., “please also describe how to address communication issues using methods derived from scientific publications and research”), which were then evaluated. In cases where the proposed strategy was deemed unsuitable, they either recommended a specific alternative to replace the initial strategy or asked ChatGPT-4 to generate a different one. If the new strategy met the criteria, they instructed GPT-4 to incorporate it into the written description. For example, when asked about communication issues, GPT-4 first suggested the Four Habit Model(10), which they evaluated to be somewhat out of date. With further prompting, GPT-4 suggested newer strategies to improve communication between physicians and patients, such as the Teach-Back Method and the Shared Decision-Making Model(11), which they then instructed GPT-4 to incorporate in the description. They were able to make these adjustments because they utilized ChatGPT search while armed with subject matter expertise and an understanding of the available literature for this topic. They leveraged this knowledge to refine the approach to prompt engineering during the process. (Please see Table 3 and Appendix C).

Table 3

Step	Prompts
1	Please also describe how to address communication issues using methods derived from scientific publications and research
2	Is the four habit model the most cited and most recent paper on how to address communication problems? Can you please find a balance between the most cited research papers and the most recent ones, when trying to find references to explain the problem and to address the problem?
3	Please explain the teach back model and shared decision-making in communication

 issues using relevant references

Comparison of human vs AI literature reviews

To the best of our knowledge there are no validated tools or checklists to compare human and AI literature reviews. Therefore, we chose to compare the reviews subjectively with respect to the accuracy, response time, comprehensibility, consistency, breadth and depth of knowledge, contextual understanding, and transparency of the outputs. The criteria are defined as follows:

- **Accuracy:** We defined the accuracy of the outcome as the percentage of correct responses.
- **Response Time:** We defined response time as the time it took to conduct the review, including identifying factors and demonstrating what they are.
- **Consistency:** We defined consistency as degree of reliability and stability in the results of a study when it is repeated under similar conditions to ensure it can be replicated.
- **Breadth of knowledge:** We defined the breadth of knowledge as the extent and range of information one has access to across various subjects and disciplines.
- **Contextual understanding:** Contextual understanding refers to the ability to comprehend the meaning and relevance of information within its specific context. It goes beyond just the literal meaning of words, but also the nuances and implications shaped by the situation and the specific circumstances under which the communication occurs(12).

Results

Exploring literature search using AI (GPT-4 model)

Utilizing iterative prompts and starting from a general prompt, GPT-4 initially generated 6 relational factors. After being repeatedly asked for more factors, it became evident that GPT-4 began to produce unrelated factors after reaching 21, likely due to hallucination. Overall, GPT-4 generated a list of 21 relational factors and provided 54 references. (Please see Table 3). Of these 21 factors, 3 (14%) were identified as irrelevant. Thirteen (24%) of the references were identified as somewhat related to the topic, but not particularly strong in their relevance, and 4 (7.5%) were identified as irrelevant.

Table 3. Factors affecting the relationships between physicians and patients that lead to patient dissatisfaction and medico-legal complaints as identified in human- and AI-led literature reviews.

Human literature review	AI-generated literature review, following iterative prompts #1-#3
Communication <ul style="list-style-type: none"> • Understanding patients' concerns and expectations • Clarity of communication • Information sharing and transparency • Tone and attitude 	1. Communication Issues 2. Lack of Informed Consent 3. Perceived Negligence or Incompetence 4. Mismatched Expectations 5. Perceived Lack of Care or Empathy 6. Systemic Issues 7. Cultural and Language Barriers 8. Failure to Follow-Up 9. Breakdown in Continuity of Care 10. Patient Autonomy Disregarded 11. Trust Erosion 12. Financial Conflicts of Interest 13. Power Dynamics 14. Failure to Respect Confidentiality 15. Inadequate Documentation * 16. Unaddressed Patient Concerns 17. Provider Burnout
Individual characteristics <ul style="list-style-type: none"> • Physician characteristics • Patient characteristics 	
Perceived care and empathy	
Healthcare system and policies	

<ul style="list-style-type: none"> • Navigating the healthcare system • Wait times • Resource constraints 	18. Poor Coordination Among Care Teams * 19. Patient's Previous Negative Experiences 20. High Patient Expectations 21. Medical Complexity *
--	--

* Factors indicated with an asterisk were identified by GPT-4, but were judged to be inaccurate by human researchers.

GPT-4 demonstrated an impressive ability to retrieve a breadth of information; however, our assessment showed that this information could be superficial, requiring an in-depth investigation to ensure its reliability and validity. Since we were uncertain how many relevant factors are related to the topic, we had prompted GPT-4 to extend the list of relevant factors. We also observed that GPT-4 will not communicate to its user when the topic has been saturated and when to stop asking for more information. For example, when we pushed it to go beyond contributing to the relational problems between physicians and patients, GPT-4 provided 30 factors, but the additional factors were increasingly irrelevant or obviously incorrect.

Although the description provided by GPT-4 for each factor was initially short and concise, with prompting the elaborations for each factor became more detailed and comprehensive. Additionally, we noted that GPT-4 initially displayed limitations in adhering to prompted numerical guidelines, such as requesting a specific sentence count, word limit, or number of references, but it started to better follow the instructions when they were reinforced in subsequent prompts.

Our findings showed that GPT-4 can offer relevant responses to questions; however, there were instances where more precise, suitable, or applicable alternative answers existed. For example, when tasked with suggesting a mitigation strategy for communication issues between physicians and patients, GPT-4's initial recommendation was the Four Habits Model. However, upon deeper expert analysis, the researchers determined that the Teach-Back Method and Shared Decision-Making model were more fitting for the review. This underscores that initial responses from GPT-4, though relevant, may require further evaluation to determine their optimal relevance and applicability.


Human literature review

A total of 120 articles were identified for review. Title and abstract screening against the eligibility criteria yielded 113 papers that were directly relevant to our objectives, of which 92 were included for full text reading and analysis. Two researchers (MM, JHF) reviewed the included articles and identified factors that affected the physician-patient relationship in ways that contributed to patient (dis)satisfaction and potential medico-legal complaints. These factors were sorted into the themes and subthemes displayed in table 3.

Comparison of human vs AI literature reviews

While both reviews identified factors influencing the relational problems between physicians and patients, there were important differences. In the human-led literature search, we used a semi-structured approach to find relevant references, then conducted a thematic analysis to group the factors into themes and convey the concepts clearly to the target audience. In contrast, GPT-4 used a proprietary search algorithm to explore the web and find relevant articles and identify relevant factors. Also, it only followed the instructions to list the factors, so there was no synthesis or grouping of the factors. A qualitative comparison of the categories below can be found in Figure 1.

Figure 1. Qualitative Comparison of Reviews Conducted by GPT-4 vs. Human Researchers



Circle Sizes (Large to Small) Qualitatively Represent *Differences* in Criteria Between GPT-4 and Human Researchers; Not Intended for Precise Measurement.

Accuracy

Of the 21 relational factors that were produced by GPT-4, 18 (86%) were assessed to be accurate based on our subjective assessment (please see Table 3 and Appendix B). As noted above, GPT-4 will continue to suggest factors based on the user's prompting, so the accuracy rate would decline if the user kept asking for additional factors.

In contrast, experienced human researchers have the nuanced judgement to identify the relevant factors and eliminate the ones that are not relevant to the subject matter. Typically, they can offer a coherent rationale to justify their identification of a factor as either relevant or irrelevant to the subject matter. In the review conducted by researchers, all the identified factors are considered quite relevant to the subject matter and their relevance is supported by scientific evidence.

Response Time

The AI model generated results within seconds, and the entire series of experiments and prompts were conducted over a few days. The human-led literature review was not conducted as a time trial and occurred as part of a researcher's regular activities over the course of several months. Had the review been conducted explicitly for this study, it would have required substantially more time for human researchers to search the literature, read and comprehend the papers, and produce results, compared to GPT-4. Our evaluation indicated that OpenAI's GPT-4 model exhibited a significant marked advantage in response time.

Consistency

In general, the GPT-4 model produced reliable responses to prompts, but similar prompts could sometimes result in variable outputs. We observed that shorter and more precise prompts were more

likely to yield consistent results, whereas complexity and length in prompts led to more variability in outputs. When conducting literature reviews, human researchers produce fairly consistent results when they have adequate resources (e.g. access to a skilled health research librarian for literature search strategies) and follow established techniques (e.g. PRISMA for systematic reviews) (9).

Breadth and depth of knowledge

Our experiments demonstrated a considerable breadth of knowledge within the GPT-4 model, significantly surpassing that of human researchers. This was particularly evident when the model almost instantly generated an extensive list of contributing factors to relational problems between physicians and patients, as well as a comprehensive list of potential mitigation strategies for each factor.

While breadth of knowledge is valuable when conducting literature reviews, synthesizing the information derived from such a review requires a depth knowledge and the ability to apply, analyze, and evaluate information related to that topic. This is an area where the GPT-4 model fell short, and a human researcher with experience in a specific subject area may have an advantage.

Contextual understanding

While LLMs are non-sentient and do not understand meaning in a traditional sense, our experiments revealed that GPT-4 was able to produce outputs that included a satisfactory level of contextual information to allow readers to understand and link key concepts. For example, through iterative prompting, the software was able to produce a list of physician-patient relational issues that included factors as varied as power dynamics, provider burnout, medical complexity, and cultural and language barriers. This level of context was improved by iterative feedback and prompting, providing expanded definitions and additional references (12). However, given the fact that GPT-4 started to hallucinate when asked to generate more factors, we concluded that it did not have a deep contextual understanding to stop generating meaningless outcomes. On the other hand, human researchers possess an understanding of meaning that consistently resulted in superior proficiency in interpreting and responding to nuanced contextual elements in this literature search which would prevent such errors.

Transparency

Another area where human researchers have an advantage is transparency. Human researchers can describe their literature search methods; state and rationalize eligibility criteria; explain the inclusion or exclusion of various articles; describe the approaches used in synthesis; and answer specific questions about their methods. There is significantly less transparency in the way that LLMs process prompts, collect information, and generate outputs at this time. Even when prompted to explain how it completed its literature review, GPT-4 will explain broadly that it drew upon diverse training data but can't provide a full list of the relevant resources it's reviewed, and so the backend review process is almost hidden.

Discussion

Many researchers are considering how AI tools can support their research. As with any new technology, there is a spectrum of uptake from 'early adopters' to 'stubborn resistors.' This paper explored how a widely available LLM tool, GPT-4, conducts literature reviews and compared the generated outcomes with a similar review conducted by human researchers.

We found that human-generated literature reviews were more transparent, more consistent, and accurate, as long as the literature review was approached systematically and the researcher had sufficient experience and expertise in the subject area. In contrast, GPT-4 generated results were much faster, provided an impressive breadth of content, and were reasonably accurate. We also found that the model was often inconsistent in its outputs, and at times generated irrelevant information, especially if forced to generate a certain *number* of factors.

One of the fundamental differences between the literature review generated by GPT-4 and humans

was in terms of contextual understanding. We attribute this difference to one often-cited limitation of LLMs, their status as so-called ‘stochastic parrots’(13) that use statistical probabilities of which word is most likely to be next, rather than understanding meaning. With prompting, GPT-4 rapidly produced an extensive list of factors that affect the relationship between physicians and patients that appeared very relevant. However, a deeper examination by experts identified inaccurate outputs amidst accurate ones. This underscores the necessity of expert evaluation in discerning the nuanced veracity of the information generated by GPT-4.

In fact, in our study, we identified two potential scenarios where researchers might encounter challenges while working with GPT-4. Firstly, effective communication with the model, specifically through adept prompt engineering, is crucial. Inadequate or improper prompting, particularly for complex tasks like conducting a literature review, leads to unsatisfactory results (detailed in Appendix A). Secondly, novice researchers, unfamiliar with a specific field, might employ effective prompting techniques and obtain a broad array of information. This breadth of knowledge can be initially impressive, yet it is important to recognize that the generated content may include errors or inaccurate information. It is for this reason that researchers must carefully review the results to identify and correct potential inaccuracies. The importance of expert oversight in evaluating the reliability of GPT-4 generated content is clear.

To address the first challenge, this paper introduced an iterative algorithm to effectively search the literature. We suggested an approach to prompt engineering that uses directive iterative prompts to guide GPT-4 to develop a literature review for researchers. This structured approach includes two phases. In the initial phase, researchers are advised to formulate a sequence of prompts that are broad yet precise, progressively becoming more specific. This approach should be designed to incrementally introduce and reinforce instructions, guiding GPT-4 towards generating an output that offers a thorough and comprehensive perspective on a particular subject. In the second phase, the researcher can independently query elements, concepts or factors identified in the first phase to explore these in greater detail. At all phases of the process, the researcher’s own understanding of the subject will shape the prompts and can drastically improve GPT-4’s literature review, suggesting relevant ideas and references while guiding the software away from outdated or incorrect concepts.

For overcoming the second challenge, we suggest approaching GPT-4 as a research assistant who possesses limited contextual expertise and occasionally synthesizes responses entirely. This requires substantial insight and knowledge from the researcher to diligently guard against the so-called ‘hallucinations’ of the software. Such vigilance is crucial, as GPT-4 can produce convincing yet entirely fabricated content and references (2, 14).

For this reason, it seems that GPT-4 might be a more useful tool for experienced researchers looking for wide surveys on a particular topic. The human researcher’s knowledge and expertise in a specific area allows them to develop appropriate prompts; iterate with the software to refine the outputs; introduce relevant frameworks and key references; and ultimately guide the process towards the desired output with a clear-eyed understanding of the limitations of what is produced. However, it can also offer different benefits to other audiences including more novice researchers. Leveraging its extensive knowledge base and inhuman quickness, GPT-4 can help newcomers to familiarize themselves with the domain under review. The software acts as an information assistant, offering a wide spectrum of knowledge within a defined domain. In addition, for researchers who have few resources and/or constrained schedules, it can be utilized to facilitate the literature review process by offering a robust preliminary draft outline, encompassing key concepts that serve as foundational building blocks. Other studies have explored the potential use of GPT-4 and other LLMs for research tasks such as scholarly writing (2, 15), medical writing(14, 16, 17) and systematic reviews(18). Still, the rapid improvement in generative AI software has also spurred a rapid growth in concerns, such as those related to the ethics of ChatGPT as a co-author(19) or the potential for it to be used to disseminate misinformation and promote plagiarism.(4) As with any nascent technology,

transparency around its use will be essential, and caution is, perhaps, warranted.

Overall, our study clearly demonstrates the potential utility of GPT-4, a LLM, in supporting the conduct of literature reviews, particularly when an iterative feedback approach to prompt engineering is employed. The software successfully reviewed the literature, identified several factors relevant to the subject matter, and was able to respond to prompts requesting additional detail and references. In some instances, and for some researchers, the benefits of using GPT-4 for a literature review (including good breadth of knowledge, reasonable accuracy, and an impressive response time) outweigh the identified shortcomings (including some inconsistency, some inaccuracy, and less depth of knowledge). We suggest that our structured approach to prompt engineering may serve as a model for researchers looking to integrate generative AI into their literature searches. Given the detailed assessment of the generated outcomes with human-led reviews, we recommend approaching these models as an assistant rather than a wise professor; researchers relying on GPT-4 to provide them with a full and nuanced understanding of a complex or rapidly-evolving subject do so at their own peril.

Limitations and future research

Our study has some limitations. Given the iterative nature of our approach to prompting GPT-4, we did not predefine our prompts or methods, and the researcher leading the prompts (MM) had extensive experience in the subject area; these factors undoubtedly influenced our prompts and thus our outcomes. Our approach to comparing the human- and AI-led literature reviews was subjective, exploratory and qualitative.

We acknowledge the limitations posed by utilizing a single database and employing a human-conducted review as the comparison standard. However, the opaque nature of ChatGPT's search strategy presents challenges in directly comparing search methodologies. These aspects are critical for interpreting our findings and suggest avenues for future research. Additionally, while we have detailed GPT-4's prompt strategies in the appendix, the proprietary and evolving nature of its algorithm limits a comprehensive methodological comparison. Future research should examine AI capabilities in detecting emerging trends and gaps, enhancing our understanding of its utility and constraints in academic research.

In our methodology for the human literature search, we employed thematic analysis, a subjective process influenced by the researchers' expertise and perspectives. We highlight the inherent subjectivity of thematic analysis as a key limitation. Similarly, our review of ChatGPT's capability to conduct literature reviews acknowledges the qualitative and subjective nature of this evaluation. Our aim was to offer insights and guidance for researchers interested in leveraging AI tools like ChatGPT in their research endeavors.

Our study's methodology involved the same researcher in both conducting the human literature review and guiding the AI, as well as participating in the team that evaluated the outcomes. While this was intended to leverage the researcher's subject expertise, it introduces a potential bias, as the researcher was not blinded to the results of the human review during the AI evaluation. This could influence the assessment and interpretation of the AI-generated content. Future studies might consider a more diversified evaluation team to further mitigate bias and enhance the objectivity of the findings.

This study is limited to an in-depth examination of OpenAI's ChatGPT-4 model, providing a detailed understanding of this specific tool's capabilities and limitations in conducting literature reviews on a particular topic. While this focus allows for a precise evaluation of GPT-4, we acknowledge that this technology is evolving very fast, and it may not reflect the performance of other AI tools that are designed to handle similar tasks. Despite this limitation, our work shows the potential of AI to streamline the initial stages of literature reviews. To build on this foundation, future research should compare the effectiveness of various AI models across a broader range of topics, thereby enhancing our understanding of the general applicability of AI-assisted literature reviews.

Moreover, upcoming studies should focus on enhancing prompt engineering methods to further

leverage ChatGPT-4's capabilities in conducting literature reviews. Addressing identified limitations, such as improving the depth and contextual understanding in AI-generated reviews, is crucial. Expanding the training datasets of ChatGPT-4 to include more diverse and recent publications could potentially mitigate issues of relevance and accuracy. Additionally, investigating the role of AI in identifying emerging trends and gaps within specific research fields, particularly in health services and medical research, would provide valuable insights into the practical applications and limitations of AI in academic research.

Final notes and considerations

In incorporating AI like ChatGPT into academic research, ethical considerations are crucial. There is the potential for bias in AI outputs, reflecting the biases present in the training data. Ensuring transparency about how AI is used, including prompt selection and response interpretation, is vital for replicability and trust. Responsible use of AI requires acknowledging the limitations and not substituting it for human expertise. As AI technologies become more prevalent in research, it is essential to establish ethical guidelines that promote awareness of bias, transparency, and responsible usage. Integrating ChatGPT-4 and similar large language models (LLMs) into academic research could dramatically change how we conduct studies, particularly literature reviews. This technology could speed up our ability to study extensive fields, enabling quicker responses to new information or gaps in knowledge. However, it's crucial to remember that the depth of understanding and critical analysis, which are at the heart of academic work, cannot be fully replicated by AI.

The use of LLMs might also make research more accessible, allowing a wider range of voices to contribute to scholarly conversations. Yet, we must navigate this future carefully, paying close attention to ethical concerns like bias in AI outputs and maintaining transparency in AI's role in research processes. As we move forward, the challenge will be to harness AI's power to enhance our work while ensuring that the essence of research, critical thinking, depth of analysis, and human insight, remains at the forefront. The potential is vast, but so is our responsibility to use these tools wisely.

Conflict of Interest

The authors have no conflicts of interest to declare. All co-authors have seen and agree with the contents of the manuscript and there is no financial interest to report.

References

1. The AI writing on the wall. *Nature Machine Intelligence*. 2023;5(1):1-.
2. Lingard L. Writing with ChatGPT: An Illustration of its Capacity, Limitations & Implications for Academic Writers. *Perspect Med Educ*. 2023;12(1):261-70.
3. Pavlik JV. Collaborating With ChatGPT: Considering the Implications of Generative Artificial Intelligence for Journalism and Media Education. *Journalism & Mass Communication Educator*. 2023;78(1):84-93.
4. Van Noorden R, Perkel JM. AI and science: what 1,600 researchers think. *Nature*. 2023;621(7980):672-5.
5. Perkins M, Roe J. Academic publisher guidelines on AI usage: A ChatGPT supported thematic analysis. *F1000Research*. 2023;12:1398.
6. Landhuis E. Scientific literature: Information overload. *Nature*. 2016;535(7612):457-8.
7. Giray L. Prompt Engineering with ChatGPT: A Guide for Academic Writers. *Ann Biomed Eng*. 2023.
8. Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, et al. Survey of Hallucination in Natural Language Generation. *ACM Comput Surv*. 2023;55(12):Article 248.
9. Pluye P, Hong QN, Bush PL, Vedel I. Opening-up the definition of systematic literature review: the plurality of worldviews, methodologies and methods for reviews and syntheses. *Journal of Clinical Epidemiology*. 2016;73:2-5.
10. Frankel RM, Stein T. Getting the most out of the clinical encounter: the four habits model. *J Med Pract Manage*. 2001;16(4):184-91.
11. DeWalt DA, Broucksou KA, Hawk V, Brach C, Hink A, Rudd R, et al. Developing and testing the health literacy universal precautions toolkit. *Nurs Outlook*. 2011;59(2):85-94.
12. Gumperz JJ. Contextualization and understanding. *Rethinking context: Language as an interactive phenomenon*. Berkeley, Calif.: Department of Linguistics, University of California at Berkeley; 1988.
13. Bender EM, Gebru T, McMillan-Major A, Shmitchell S. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? ? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency; Virtual Event, Canada: Association for Computing Machinery*; 2021. p. 610-23.
14. Alkaisi H, McFarlane SI. Artificial Hallucinations in ChatGPT: Implications in Scientific Writing. *Cureus*. 2023;15(2):e35179.
15. Dergaa I, Chamari K, Zmijewski P, Ben Saad H. From human writing to artificial intelligence generated text: examining the prospects and potential threats of ChatGPT in academic writing. *Biol Sport*. 2023;40(2):615-22.
16. Benichou L. The role of using ChatGPT AI in writing medical scientific articles. *Journal of stomatology, oral and maxillofacial surgery*. 2023;124(5):101456.
17. Doyal AS, Sender D, Nanda M, Serrano RA. ChatGPT and Artificial Intelligence in Medical Writing: Concerns and Ethical Considerations. *Cureus*. 2023;15(8):e43292.
18. Qureshi R, Shaughnessy D, Gill KAR, Robinson KA, Li T, Agai E. Are ChatGPT and large language models "the answer" to bringing us closer to systematic review automation? *Syst Rev*. 2023;12(1):72.
19. De Angelis L, Baglivo F, Arzilli G, Privitera GP, Ferragina P, Tozzi AE, et al. ChatGPT and the rise of large language models: the new AI-driven infodemic threat in public health. *Front Public Health*. 2023;11:1166120.

Supplementary Files

Untitled.

URL: <http://asset.jmir.pub/assets/8b3d585ef25ff3c537f88620f710f68e.docx>

Multimedia Appendixes

Untitled.

URL: <http://asset.jmir.pub/assets/7e70a1eeaad685db1183c7d1d283ebed.docx>