Search Engines, Large Language Models or Both? Evaluating Information Seeking Strategies for Answering Health Questions

Marcos Fernández-Pichel[a,∗], Juan C. Pichel[a] and David E. Losada[a]

[a]*Centro Singular de Investigación en Tecnoloxías Intelixentes (CiTIUS), Universidade de Santiago (USC), Rúa de Jenaro de la Fuente s/n, Santiago de Compostela, 15705, Spain*

ARTICLE INFO

ABSTRACT

Search engines have traditionally served as primary tools for information seeking. However, the new Large Language Models (LLMs) have recently demonstrated remarkable capabilities in multiple tasks and, specifically, their adoption as question answering systems is becoming increasingly prevalent. It is expected that LLM-based conversational systems and traditional web engines will continue to coexist in the future, supporting end users in various ways. But there is a need for more scientific research on the effectiveness of both types of systems in facilitating accurate information seeking. In this study, we focus on their merits in answering health questions. We conducted an extensive study comparing different web search engines, LLMs and retrieval-augmented (RAG) approaches. Our research reveals intriguing conclusions. For example, we observed that the quality of webpages potentially responding to a health question does not decline as we navigate further down the ranked lists. However, according to our evaluation, web engines are less accurate than LLMs in finding correct answers to health questions. On the other hand, LLMs are quite sensitive to the input prompts, and we also found out that RAG leads to highly effective information seeking methods.

1. Introduction

With the recent advancements in Natural Language Processing (NLP), Large Language Models have become major players in numerous Information Access tasks (Longpre et al., 2023; Bubeck et al., 2023; Touvron et al., 2023). The release of ChatGPT in November 2022 has been a game-changer globally, marking a significant milestone and revolutionising many sectors. One of the outstanding features of current conversational AIs stands on their ability to generate coherent and human-like text, which has garnered attention and excitement among practitioners, researchers and the general public. This breakthrough has precipitated a transformative shift in the orientation of information access research towards LLMs, their potential applications and the interconnection between LLMs and other computer-based tools.

But the emergence and global adoption of advanced LLMs has sparked the urgent need to explore and understand their capacities and knowledge acquisition attributes. Some research studies have focused on the capabilities of these models under specific language understanding and reasoning benchmarks (Jiang et al., 2020; Liang et al., 2022; Chang et al., 2024). However, with the rapid embrace of generative language models, another significant shift in information access has occurred. The conversational paradigm has gained traction, enabling more interactive and user-friendly search experiences (Mao et al., 2023; Friedman et al., 2023; Polak and Morgan, 2023; O'Leary, 2022); and many citizens currently turn to conversational AIs for consulting multiple types of information needs.

However, the role of traditional web search engines (SEs) in answering user-submitted queries is far from being relegated. For example, SEs are still the main reference for numerous information seeking tasks. As stated in the annual Digital News Report conducted by the University of Oxford, "search engines continue to grow as a direct access to news websites" (Reuters Insitute, University of Oxford, 2023). Thus, it is expected that both conversational language models and traditional search engines will continue to coexist in the foreseeable future. Furthermore, there is extensive research on how to improve SEs with LLM-based elements and also on how to enhance LLMs' generations with search-based results.

In line with these developments, it is essential to acknowledge the significance of accurate medical information. Fox and colleagues demonstrated that 80% of Internet users seek medical advice online (Fox, 2011). However, the presence

---

∗Corresponding author

✉ marcosfernandez.pichel@usc.es (M. Fernández-Pichel)
ORCID(s): 0000-0002-6560-9832 (M. Fernández-Pichel)

of misinformation in search results can be severely harmful, especially when it concerns health-related content (Sharma et al., 2019). Medical hoaxes, miracle diets or advice given by unqualified sources are prevalent in online media, and the potential consequences of health misinformation can result in personal damage (Pogacar et al., 2017; Vigdor, 2020). We are not aware of specific statistics or research studies on the number of users using conversational AIs for health information seeking. Still, popular models such as ChatGPT have millions of users worldwide and, thus, it is crucial to put their answers under scrutiny. This is of paramount importance in a critical domain such as health.

Summing up, there is a pressing need to evaluate the abilities of classic and new information access tools in answering medical questions and, as a matter of fact, there is a lack of comprehensive studies in the literature that compare the effectiveness of LLMs with that of traditional SEs in the context of health information seeking. Moreover, since the performance of LLMs is highly dependent on the input prompt (Jiang et al., 2020; Brown et al., 2020; Liu et al., 2023), understanding their effectiveness with different types of prompts is of utmost importance. It is also crucial to explore the effect of combining both classes of tools and, for example, explore the behaviour of LLMs when prompted with medical questions together with related search results. This paper aims to contribute towards filling these gaps by conducting a thorough study aimed to answer the following research questions:

- **RQ1**. To what extent do search engines retrieve results that help to answer medical questions? Does the correctness of the information provided drop as we go down in the search engine result page?

- **RQ2**. Are LLMs reliable in providing accurate medical answers? How do different models compare in terms of their effectiveness in providing responses to medical questions?

- **RQ3**. Does the given context influence the capabilities of LLMs in providing the right answers? Do these models exhibit improvements when presented with a few in-context examples?

- **RQ4**. Do LLMs improve their performance when fed with web retrieval results?

Our comparison of SEs included experiments with Google, Bing, Yahoo! and DuckDuckGo. We found that Bing seems to be the most solid choice. Our evaluation also suggests that extracting the answers from the top ranked webpage often produces good results. This is good news, as web users are known to be reluctant to inspect many items from the search engine result pages (SERPs). Moreover, LLMs show an overall good performance but they are still very sensitive to the input prompt and, in some instances, they provide highly inaccurate responses. Finally, augmenting LLMs with retrieval results from SEs looks very promising and, according to our experiments, even the smallest LLMs can achieve state-of-the-art performance if fed with appropriate retrieval evidence. These findings contribute to a better understanding of health information seeking with the new conversational AIs and with the traditional search engines.

2. Related work

This paper is related to several scientific areas. The advances in information credibility and correctness, particularly in the area of health information, are relevant to our research and are discussed in Section 2.1. Section 2.2 reviews the most significant literature in health information access with traditional search engines and with the new LLMs. Finally, in Section 2.3 we review the main trends in retrieval augmented LLMs.

2.1. Health Information Credibility and Correctness

Credibility represents a subjective perception of the extent to which information from a webpage or other source can be trusted (Fogg, 1999). Previous research has extensively studied how online information credibility is established (McKnight and Kacmar, 2007; Ginsca et al., 2015; Kakol et al., 2017; Bodaghi et al., 2023). For instance, Viviani and Pasi (2017) presented a thorough survey on existing credibility determination methods, showing that health misinformation poses a socially relevant problem. The work by Matthews et al. (2003) showed that a corpus about alternative cancer treatments contained 90% of documents with at least one false claim. On the other hand, Sondhi and his peers presented an automatic approach, using classic supervised learning technology, for medical reliability classification of webpages (Sondhi et al., 2012). Fernández-Pichel et al. (2021b) successfully reproduced Sondhi's study and further applied it to new collections, thereby demonstrating its potential for generalisation. The same team also performed a comparison between traditional learning methods and neural-based approaches solutions for health-related misinformation detection (Fernández-Pichel et al., 2021a), concluding that traditional learning approaches still constitute a robust baseline for some prediction tasks.

Under similar supervised learning settings, other studies focused their efforts on how different characteristics or features influence credibility. Zhao et al. (2021) proposed a novel health misinformation detection model using both "central-level" features (e.g., the topics discussed) and the so-called "peripheral-level" features (including linguistic and sentiment features, and user behavioural features). Their approach was validated on a real-world dataset, obtaining 85% of accuracy in health misinformation detection. This study showed that behavioural features are more discriminative than linguistic features in detecting untrustworthy contents. In (Griffiths et al., 2005), the authors demonstrated that network-based features alone, such as those based on PageRank, do not suffice to determine the reliability of online content.

Another line of work specifically focused on end-users and their perceptions of credibility. Seminal work by Fogg (2003) proposed the prominence-interpretation theory, which helps to determine which website cues influence the perception of credibility. Other studies showed that perceived credibility depends on reading skills (Hahnel et al., 2018). Similarly, Liao and Fu (2014) studied age differences in credibility judgements. On the other hand, Schwarz and Morris (2011) focused on how to augment a search engine result page to improve medical credibility judgements. In (Fernández-Pichel et al., 2024), a user study with 1,000 participants was performed to analyse people's perceptions to the credibility of online health information. Among their main findings, these authors showed that individuals tend to overestimate the credibility of low quality sites. In the literature, related aspects, such as credibility, trustworthiness, and correctness have been explored. Our research particularly focuses on the correctness of responses for health questions provided by web search systems and conversational AIs.

Some teams focused their efforts on designing solutions that estimate the correctness of medical information. For instance, Pradeep et al. (2021) presented Vera, a transformer-based ranker that achieved state-of-the-art results in health misinformation detection tasks. This model was fine-tuned with assessments from the TREC 2019 Decision Track and it achieved the best results in the TREC 2022 Health Misinformation Track. This shared-data task fosters the development of systems that promote credible and correct documents over misinformation. However, systems like Vera take a health question *and its correct response* as an input and then search for harmful or helpful documents. These technological tools could support, for example, moderation services for online platforms. However, the need of pairs of questions and correct responses represents a limitation and, thus, the creators of Vera also conducted research on how to automatically infer the correct response for a health question (Pradeep and Lin, 2024). Specifically, the authors evaluated two different approaches: i) using LLMs under different settings (zero-, few-shot and chain-of-thought prompting), and ii) using the ranking produced by Vera and averaging the decisions extracted from the top 50 retrieved results. Their results suggest that the LLM-based approach (powered by GPT-4) outperformed the rest of strategies. Our comparison of SEs and LLMs is related to this study, but we compare multiple commercial search engines (while Vera is not a tool that is widely available to the public) and, furthermore, we do not only evaluate GPT models but consider six LLMs of different families. Our study, therefore, is reflective of the type of answers that the general public might get when interacting with popular information access tools. Additionally, we compare the answering capabilities of the LLMs with and without search results provided by the web search engines. This is an aspect that has not received enough attention in health information seeking. We also evaluate the effectiveness of the answers provided as we go down in the ranked lists of the SERPs.

2.2. Search Engines and Large Language Models in Health Information Access

Web search is widely used to obtain health advice and effective medical information retrieval has attracted the attention of the scientific community over the years (Fox, 2011). For instance, Baujard et al. (1998) conducted a seminal study focusing on the main trends in medical information retrieval and proposed an agent to perform effective medical information discovery. Similarly, Bin and Lun (2001) addressed the difficulty of finding relevant medical information in the web. These authors suggested that general-purpose search engines constitute a strong baseline for health information retrieval but also designed an agent-based system that was able to outperform SE baselines. Soldaini et al. (2016) studied the influence of injecting expert medical vocabulary to the query. Their results showed an increase around 7-12% in correct results for the modified queries. However, effective information retrieval (i.e., topicality) is not enough. As Pogacar et al. (2017) have shown, on-topic incorrect results can severely bias people's decisions. The need for new algorithmic solutions able to provide accurate results for health queries has motivated the development of shared-tasks that promote the identification of correct and credible information over misinformation, such as the TREC Health Misinformation (HM) Track (Clarke et al., 2020, 2021) or the CLEF eHealth initiative (Goeuriot et al., 2020; Kelly et al., 2019; Suominen et al., 2018). In this study, we are specifically interested in analysing how search results and LLMs' responses can help to answer binary health questions. Thus, we assess here the extent to which

search results or LLM completions provide misleading responses with respect to the established medical consensus for each health topic.

With the impressive recent development of LLMs, interest in assessing the correctness of the health-related AI-based completions has escalated. For example, Chervenak et al. (2023) demonstrated ChatGPT's abilities to answer fertility questions. Duong and Solomon (2023) compared the performance of Large Language Models against humans for answering multiple-choice questions of human genetics. Similarly, Holmes et al. (2023) conducted a comparative study of LLMs knowledge on a highly specialised topic, radiation oncology physics. They concluded that OpenAI's models outperformed all others. In (Jahan et al., 2023; Hamidi and Roberts, 2023; Samaan et al., 2023), the authors conducted studies on the role of LLMs for biomedical tasks, patient-specific EHR questions, and bariatric surgery topics, respectively. With a broader perspective, Johnson et al. (2023) involved physicians in a thorough evaluation of ChatGPT's accuracy in answering medical queries, and other researchers (Thirunavukarasu et al., 2023) evaluated ChatGPT's ability with the Applied Knowledge Test (AKT) of the Membership of the Royal College of General Practitioners, demonstrating performance close to human experts. All of the aforementioned studies are restricted to a single model, usually ChatGPT, and/or to a specific medical area. Fernández-Pichel et al. (2024) evaluated several LLMs of different nature with general health questions. This team examined the influence that different prompts and in-context examples have on the effectiveness of the LLMs' output. In our study, we go one step further and present a thorough comparison between LLMs and traditional web search engines and, additionally, test the combination of LLMs and SEs through retrieval-augmented generation (RAG) strategies.

2.3. Retrieval-Augmented Generation (RAG)

Several previous studies focused on how to exploit retrieval evidence to enhance the generative responses supplied by AIs (Li et al., 2022). Asai et al. (2023) provided a detailed overview of the advances in retrieval augmented LMs and explored some of their applications. Guu et al. (2020) proposed REALM, an augmented language model with a module that retrieves evidence from a large textual corpus such as Wikipedia. In (Borgeaud et al., 2022), the authors proposed RETRO, an architecture consisting of a language model which benefits from results obtained from a large database of trillions of tokens.

Lazaridou et al. (2022) conducted an evaluation on different few shot prompting strategies. They compared a method that provided examples of questions and correct answers (*closed book* strategy) against a method that fed the model with supporting evidence for the answer (*open book* strategy). The injected evidence came from an offline labelled collection or from relevant passages obtained from Google's search API. Their experiments showed that both types of open book strategies improved the closed book method. Izacard and Grave (2021) augmented a T5 model for an open domain question answering task. For retrieval of relevant passages, these authors used an offline collection and BM25 or DPR as search methods. The main conclusion was that augmentation is beneficial and that improvements are noticeable up to a maximum of 100 recovered passages. Asai et al. (2021) conducted a similar study, but they argued that injecting off-topic evidence might be counterproductive. Thus, their solution included an evidentiality estimation layer that predicts whether the passage provides actual evidence to answer the question.

Other studies focused on helping language models to interact with search engines. For instance, Nakano et al. (2021) developed a text-based web-browsing environment that can be exploited by a GPT-3 fine-tuned model. Thoppilan et al. (2022) showed that using external APIs (such as those supported by information retrieval systems) significantly improves groundedness, which is defined as the extent to which a generated response contains claims that can be validated against a known source. Shuster et al. (2022) proposed an architecture that first searches for evidence against a search engine, then selects the most relevant sentences from the retrieved documents, and finally, generates a response.

In the health domain, Li et al. (2023) fine-tuned a Llama model with medical conversations and also injected medical evidence extracted from Wikipedia and other medical sources. We contribute to this recent line of work by assessing the effect of including search engine's evidence for the generation of correct health answers. In Section 6.2.4, we detail these experiments, oriented to prompt several LLMs with evidence retrieved from Google's top results for a given health question.

3. Binary Question Answering with Search Engines and Large Language Models

In this research, we focus on estimating the ability of web search engines and conversational AIs to provide correct health responses. To that end, we selected a binary health question answering (QA) task as our reference for evaluation. This type of information needs are prevalent, as many users usually search for specific advice, e.g. "*Can X (treatment)*
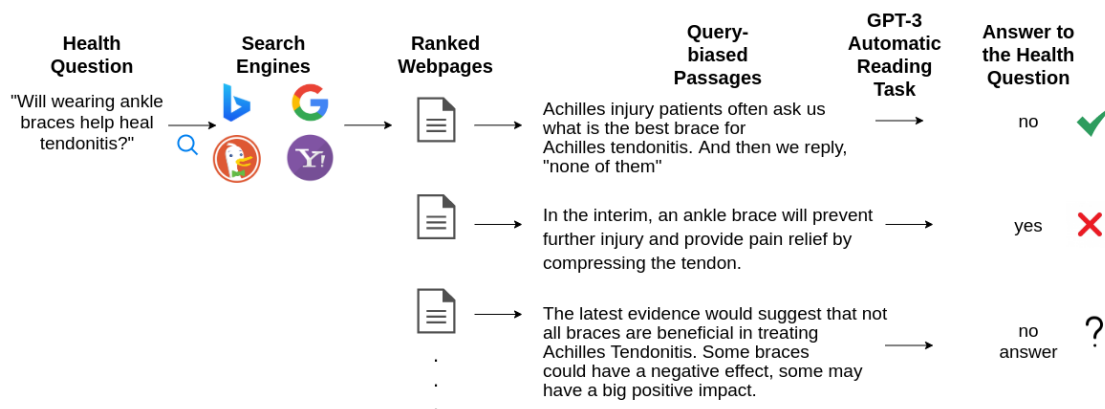
Figure 1: Estimating responses for health questions from the search engines results. The passages shown are real passages extracted from the top search results for the question "Will wearing ankle braces help heal tendonitis?".

cure Y (disease)?". This binary setting facilitates the quantitative assessment of the systems' responses. Automatic systems must provide the correct answer to these health questions and, as shown below, the target or reference response comes from the established medical consensus for each search topic.

For evaluating the responses provided by the SE, each health question was submitted to the engine. Next, we automatically extracted from each retrieved webpage the most relevant excerpt that answers the question, more details are provided in Section 5. After the passages were chosen, it was necessary to ascertain whether they provided an affirmative, negative, or non-responsive answer to the health question. To that end, we exploited GPT-3 model's capabilities to perform reading comprehension. Brown et al. (2020) evaluated the ability of this model on different reading comprehension settings. The model achieved remarkable results, for example, F1 of 85% in the CoQA dataset. Dijkstra et al. (2022) also demonstrated this model's ability to perform reading comprehension tasks in the educational domain. In our case, we defined a prompt that included the passage and health question, and specifically asked the model to answer the question based on the provided passage (without resorting to its internal knowledge)[1]. This estimation process arguably produces the sequence of answers that a search engine user would obtain from inspecting the SE results. Figure 1 illustrates the whole process, while specific experimental details are reported in Section 5.

Our interest here is not only to analyse the overall effectiveness of web search systems, but also to study the quality of the responses as we go deeper into the ranked lists. To that end, we propose two models of user behaviour that simulate alternative forms of inspection of the retrieved results. The *lazy user model* represents a user who stops inspecting results when presented with the first entry that gives a yes/no response to the user's question[2]. This user, therefore, sticks to the first answer found and does not spend time searching for contrasting evidence. The second model, referred to as *diligent user model*, represents a user who traverses the ranking from the top position and stops after finding three responses. We assume that this user takes a decision about his health question based on majority voting from the three provided responses. Web users are known to be reluctant to explore many search results and these two models represent two rational forms of exploration in the quest of the response for the health question. We leave the study of additional user models, including click models for web search (Chuklin et al., 2015), for future work. In Appendix A, we provide the pseudo-code for both user models.

For assessing conversational AIs, we submitted the health question to several generative models and obtained their completions. We worked under different settings, including zero- and few-shot strategies (providing examples of correct responses to other health questions), and different types of prompts. We forced the LLMs to provide a 'yes' or 'no' output. Specific details about the models and the configuration of the experiments can be found in Section 5

In this study, we also performed an "online" retrieval augmented generation (RAG) experiment (Asai et al., 2023). There are several approaches for augmenting generative language models with retrieved evidence. Following (Guu et al., 2020), we injected textual chunks in the input layer of the LLM. In essence, we prompted different large language models with passages extracted from the top entries of the SERPs. In this way, our experimental setup allows the

---

[1]Prompt: *<Passage>. Based on the previous text, answer 'yes', 'no' or 'no answer provided' to the following question: <health question>*
[2]Note that relevant passages might not respond to the user's question and, thus, the lazy user model does not necessarily stop at the top 1 entry.

```xml
<topic>
  <number>1234</number>
  <query>dexamethasone croup</query>
  <description>Is dexamethasone a good treatment for croup?</description>
  <narrative>Croup is an infection of the upper airway and causes swelling,
    which obstructs breathing and leads to a barking cough. As one kind of
    corticosteroids, dexamethasone can weaken the immune response and
    therefore mitigate symptoms such as swelling. A very useful document
    would discuss the effectiveness of dexamethasone for croup, i.e. a very
    useful document specifically addresses or answers the search topic's
    question. A useful document would provide information that would help
    a user make a decision about treating croup with dexamethasone, and
    may discuss either separately or jointly: croup, recommended treatments
    for croup, the pros and cons of dexamethasone, etc.</narrative>
  <disclaimer>We do not claim to be providing medical advice, and medical
    decisions should never be made based on the stance we have chosen.
    Consult a medical doctor for professional advice.</disclaimer>
  <stance>helpful</stance>
  <evidence>https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5804741/</evidence>
</topic>
```

Figure 2: Example of a Topic from the TREC 2021 Health Misinformation Track

comparison of three classes of outputs (SE alone, LLM alone, and a RAG variant that injects real-time SE results into the input of the LLM).

4. Health Questions

To obtain a solid set of health-related search topics we leveraged the data created under the TREC Health Misinformation (HM) Track. This is a three-year shared-task that aimed to foster the development of systems capable of detecting false health information, thereby empowering individuals to make health-related choices grounded on reliable and factual information (Clarke et al., 2020, 2021). These datasets contain health-related queries posed as questions (for instance, "*Does wearing masks prevent COVID-19?*") and their correct responses (yes/no). Each topic represents a searcher who is looking for information that is useful for making a "yes" or "no" decision regarding a health-related information need. The binary ground truth field represents the best understanding of current medical practice (gathered by the task organisers when building the collection). Figure 2 shows an example of a topic (the question is stored into the description field while the binary response is stored into the stance field[3]).

The TREC HM 2020 dataset was restricted to questions related to COVID-19, while topics from TREC HM 2021 and 2022 collections covered a wide range of health topics. The 2020 questions were disclosed in the middle of 2020, which raises the possibility that they could have been included in the pre-training phase of some Large Language Models. Such inclusion could potentially give them an advantage over generative models that were trained earlier. The 2021 dataset was made available in mid-July 2021 and, thus, it could have been used for training newer models such as ChatGPT and GPT-4, but not for GPT-3 (its training ended earlier). The 2022 dataset, on the other hand, was released after all the LLMs had been created. This configuration of health topics therefore shapes an assorted evaluation with questions created and released at different dates. In our analysis we pay special attention to the topic set creation dates, the knowledge cutoff dates of each LLM, and make additional experiments to estimate memorisation (see Section 6.2.5). In the case of search engines, their constant crawling and indexing activities allow them to have access to continuously updated information and, of course, this distinctive feature of SEs needs to be taken into account when analysing results.

Summing up, the selection of health topics integrates a diverse set of binary health-related questions, which pose different levels of difficulty to the models based on their exposure to such data and the level of specificity of the information needs.

---

[3]The binary ground truth was encoded as "helpful" or "unhelpful".

5. Experimental Setup

We evaluated four well-known search engines: *Google*, *Bing*, *Yahoo*, and *Duckduckgo*. We used a scraping tool[4] on the organic search results to collect the top retrieved webpages. We gathered the top 20 ranked entries since users rarely go beyond the second page of results and, thus, the extraction focuses on webpages that have some chance of being inspected by a standard user. Then, we obtained the raw content of the webpage, converted it to workable plain text, and split it into passages. To obtain the most relevant passage, the health question and each passage were embedded into a vectorial contextual representation using MonoT5 (Pradeep et al., 2023), which is a highly effective model fine-tuned for passage retrieval. This constitutes a well-established approach for the automatic extraction of excerpts answering user queries (Nogueira and Cho, 2019). Rosa et al. (2022) demonstrated the effectiveness of this model for several retrieval tasks.

In some cases, the retrieval results do not provide an answer to the question. In our accounting of correct answers across ranking positions these cases are recorded as failures, as they do not respond with the correct answer. It is important, though, to bear in mind that there is an important distinction between a non-answer and an incorrect answer. In the future, we will further delve into this issue but, in the current study, we focus on analysing the relative trends of correct responses. Furthermore, the user models discussed above assume that users skip non-answers and, thus, our user model evaluation constitutes a natural complement to the report on the proportion of correct answers across ranking positions.

We also evaluated six LLMs of different nature and architectures (closed and open source). Fine-tuned models specifically trained for health or clinical topics, such as ChatDoctor (Yunxiang et al., 2023) or BioBERT (Lee et al., 2020), were not included in our comparative study because we focus here on systems that are widely available to non-specialised users. We are specifically interested in evaluating LLMs that are currently used by millions of individuals worldwide. The models included in our comparative study are:

- GPT-3, also known as *text-davinci-002* (d-002). It is a model with a decoder-only structure with 175 billion parameters. Its training corpus is extensive, including diverse sources and the entirety of Wikipedia, with information current up to June 2021.

- *text-davinci-003* (d-003), the subsequent iteration of GPT models, built upon its predecessor by incorporating InstructGPT methodologies (Ouyang et al., 2022). This model has been refined through reinforcement learning with human feedback (RLHF) and has the same training data timeline as d-002.

- *ChatGPT* represents an evolution of OpenAI models towards a more dialogue-oriented and user-friendly behaviour. Its knowledge cutoff goes up to September 2021. For our experiments we used gpt-3.5-turbo version (its snapshot from June 2023).

- *GPT-4*, another conversational agent by OpenAI that represents a significant leap forward, outperforming ChatGPT in various complex tasks that require human-like reasoning, such as passing academic examinations (OpenAI, 2023). Its training data is up-to-date as of September 2021.

- *Flan T5* (FT5), a sequence-to-sequence model from Google. It underwent fine-tuning with a diverse array of instruction-based datasets, as described in (Longpre et al., 2023). The data for this model was sourced from the "Flan 2022" open-source repository, which includes a wide range of contents collected up until 2022. For our experimentation, we used the flan-t5-xl version.

- *Llama2*, a language model developed by Meta AI. It was trained with more than 1 million human annotations of conversational data. Its training data goes up to September 2022, but its fine-tuning also includes data up to July 2023. For these experiments, we used the llama-13b-chat version.

The first four models were tested through OpenAI's official Python API[5], while Flan T5[6] and Llama[7] were tested through their Hugging Face implementations. We set the models' temperature to 0, with the intention of minimising the randomness or creativity of the completions.

---

[4] https://github.com/tasos-py/Search-Engines-Scraper
[5] https://openai.com/blog/openai-api
[6] https://huggingface.co/google/flan-t5-xl
[7] https://huggingface.co/TheBloke/Llama-2-13B-chat-GGUF. This is a quantised version of the model to run in CPU.

A pivotal part of this research consists of determining the effectiveness of these models for answering health questions under different input conditions or contexts. Having in mind that online users are known to be reluctant to have complex interactions with automated systems, we first tested the LLMs' effectiveness when responding to non-expert individuals who send the question and give little or no context at all:

- **no-context prompt**: a context formed only with the medical question, i.e. "*Can Vitamin D cure COVID-19?*".

- **non-expert prompt**: The text "*I am a non-expert user searching for medical advice online*" is added before the health question. This prompt intends to be representative of a regular end user searching for medical advice.

As a next step, we also tested more elaborated prompts and, additionally, assessed the influence of including in-context examples. It is unlikely that a regular user opts for these complex strategies but, still, they can help to further understand and exploit the models' internal knowledge:

- **expert prompt**: The text "*We are a committee of leading scientific experts and medical doctors reviewing the latest and highest quality of research from PubMED. For each question, we have chosen an answer, either 'yes' or 'no', based on our best understanding of current medical practice and literature.*" followed the corresponding medical question. This contextual instructions were produced by Waterloo's team in their participation in the TREC 2022 HM track (Pradeep and Lin, 2024). The core idea is to guide the LLM towards reputed sources.

To automatically record responses, we forced the models to answer only with "yes" or "no" tokens. In contrast to the scenario with search engines, there can therefore be no unanswered questions here. More complex prompt engineering techniques, like Chain-of-Thought (CoT) (Wei et al., 2022), are left for future work.

For the RAG experiments, we fed some LLMs with relevant passages obtained from the search results produced by Google. We injected passages from Google's top 5 results and we clearly instructed the LLM to compare the provided evidence with its internal knowledge prior to providing a definitive answer[8]. The extraction of relevant passages from Google's top 5 results was also supported by MonoT5.

6. Results

In Section 6.1 we report the effectiveness of different search engines in providing correct responses and we also analyse user experience following the two user models explained in Section 3. In Section 6.2 we evaluate the performance of LLMs under zero- and few-shot settings. Here, we also propose a taxonomy of errors made by LLMs, make an error analysis, present results for retrieval-augmented strategies and, finally, estimate the effects of memorisation or data leakage during LLMs' training process.

6.1. Search Engines

Using the methodology explained in Section 3, we evaluated the effectiveness of the different search engines in providing the correct answer to the health questions. Figure 3 shows the results for the different collections. Each data point in the graph represents the proportion of correct responses seen up to that position[9]. For example, the fourth point in each line (position=4) represents the proportion of correct responses obtained after inspecting the top 4 results.

From the results, we observe no significant drop in performance as we move down the rankings, which speaks well of the SEs' retrieval capabilities. Regarding different engines, Bing seems a solid choice across all datasets. The 2022 topics produced the best results. A possible explanation for this outcome is that the 2022 collection contains less specialised health questions, and it might be easier to retrieve correct answers from the web.

We also wanted to analyse the extent to which retrieval results provide an actual answer to the health questions. Ranked lists contain off-topic results and, additionally, many on-topic pages do not provide a clear response to the user's request. For each of the 50 health questions available in each collection, we processed its top retrieved results to identify whether they answered the health question[10]. For example, for Google's rankings and top 1 position we computed the number of questions that had an actual response. The same count was obtained for each position (up to

---

[8]Prompt for RAG: *Provide an answer to the question using the provided evidence and contrasting it with your internal knowledge. Evidence: <Evidence from search engine>. Question: <query>. Your answer:*

[9]Remember that for this evaluation we consider as non-correct answers the entries providing no response.

[10]This identification was automatic, supported by the same passage extraction approach and reading comprehension technology described above.
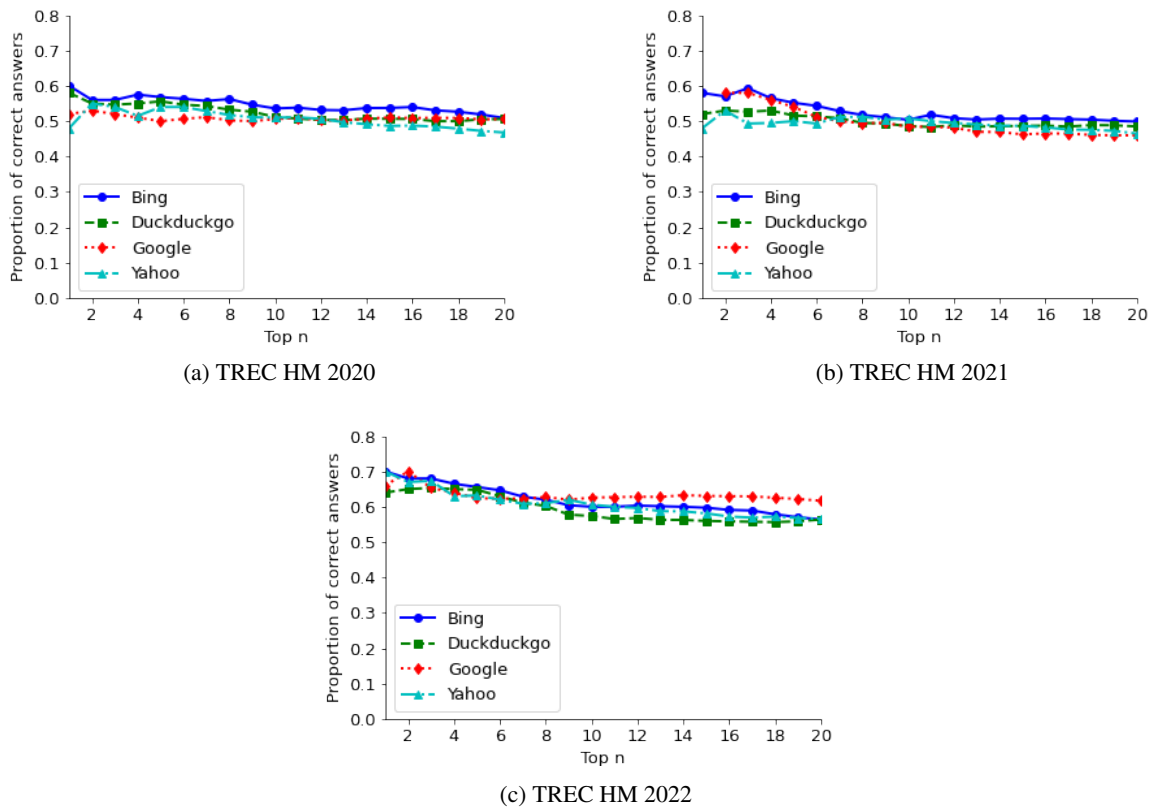
(a) TREC HM 2020

(b) TREC HM 2021

(c) TREC HM 2022

Figure 3: Proportion of correct answers obtained from the search engines after inspecting the top n ranked results.
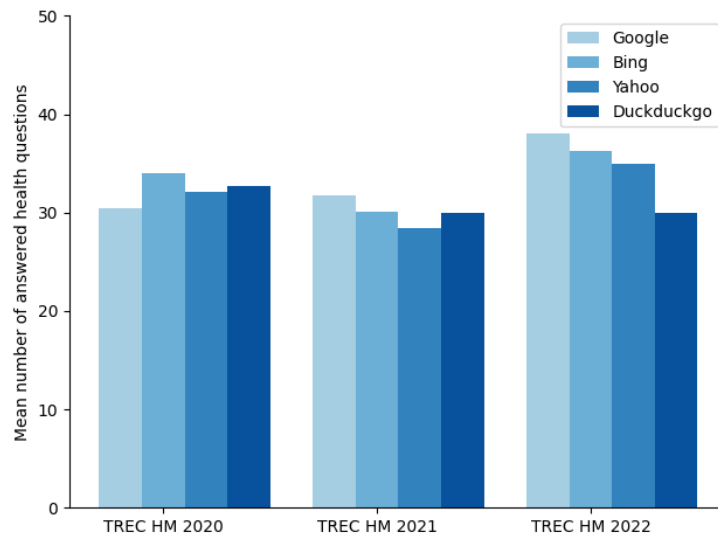

Figure 4: SE answering scores. Average number of health questions (out of 50) that were answered at the SE's top 20.

rank #20) and, finally, a global *SE answering score* was obtained by averaging the number of questions answered at each position. These scores, plotted in Figure 4, are much lower than 50, reflecting that many retrieval results do not
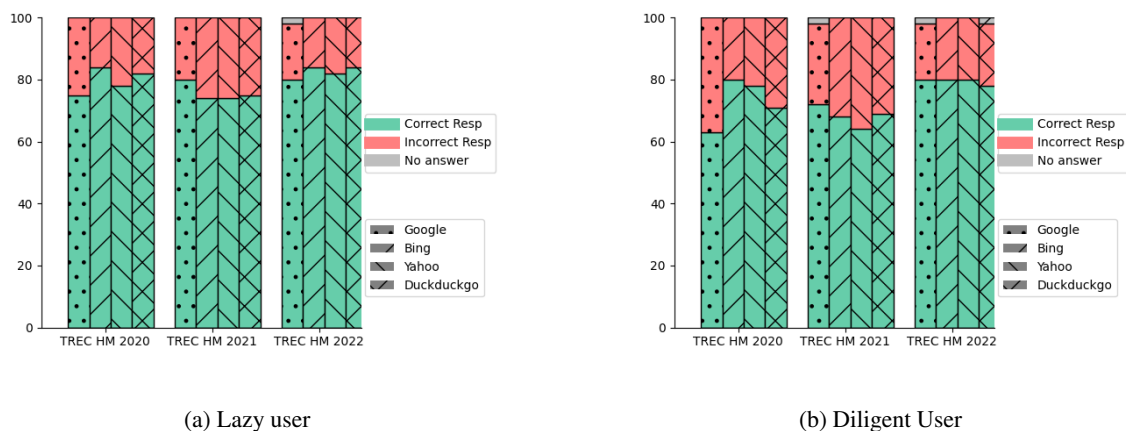
(a) Lazy user

(b) Diligent User

Figure 5: Lazy and Diligent User Models. Percentages of correct, incorrect and no-answer responses.

provide a clear yes/no recommendation to the health question. According to this experiment, Bing is the engine whose rankings provided more answers for 2020 questions (around 30 questions answered) and, from Figure 3(a) we infer that Bing's responses tend to be better than those provided by the three other SEs. For 2021, Google and DuckDuckgo supplied responses to more questions but at the cost of providing more incorrect responses (Figure 3(b)). Finally, for 2022 questions, Google provided many responses and, according to Figure 3(c), it was not inferior to the other SEs. Note also that, on average, between 15 and 20 questions did not receive a response at any given ranked position. This suggests that restricting the analysis to a few top ranked webpages, which is the common behaviour of end users, often results in a failure to meet these health information needs.

The proportion of correct answers is low (60%-70%), but the good news is that this is because many ranked results did not provide an answer. In fact, if we do not count the entries that do not provide a response as failures then the proportion of correct answers rises to 80%-90% for all search engines. Although high, these values still reflect a worrying percentage of 10-15% of incorrect answers.

6.1.1. User Behaviour Models

Let us now evaluate the search results using the two user behaviour models (*lazy* and *diligent*) described in Section 3. In Figure 5, we report the percentage of correct responses, incorrect responses and no answers for the three collections[11]. We also report in Table 1 the effort required to make a decision, measured as the mean number of results a user needs to inspect before reaching the stopping criterion.

Observe that the lazy behaviour produces better results with less effort compared with the diligent behaviour. Indeed, the additional effort spent by the diligent user does not translate into a lower percentage of incorrect responses. This might sound surprising, but it somehow reinforces the confidence in the search engine top result (the top answer suffices). In fact, our results suggest that the diligent user, who goes deeper in the ranking to find additional responses, would make poorer health-related decisions.

6.2. Large Language Models

Next, we evaluate the performance of LLMs for our binary question answering task under zero- and few-shot settings. We also propose a taxonomy for the errors that these models tend to produce and perform an error analysis based on this taxonomy. Finally, the results for retrieval-augmented strategies are presented, along with an estimation of the effects of memorisation during the LLMs' training process.

6.2.1. Zero-shot evaluation

Figure 6 plots the proportion of correct answers for the three prompting strategies ("no context", "non-expert" and "expert"). Llama2 and text-davinci-003 emerge as the top performers for the TREC HM 2020 and 2021 datasets. When

---

[11] Note that in a few instances the user exhausted the ranked list and found no response (grey area in Figure 5). These cases are more frequent in the diligent user scenario because this model needs three answers to make an informed decision.

	TREC HM 2020		TREC HM 2021		TREC HM 2022	
	Lazy	*Diligent*	*Lazy*	*Diligent*	*Lazy*	*Diligent*
Google	1.7	5.2	1.7	5.2	2	4.9
Bing	1.8	5.4	1.5	4.3	1.6	4.6
Yahoo	1.6	5.5	1.6	5.2	1.7	4.4
Duckduckgo	1.7	5.1	1.4	4.9	1.7	4.6

Table 1
Effort required to make a decision by each user behaviour model. The effort is measured as the mean number of results inspected.

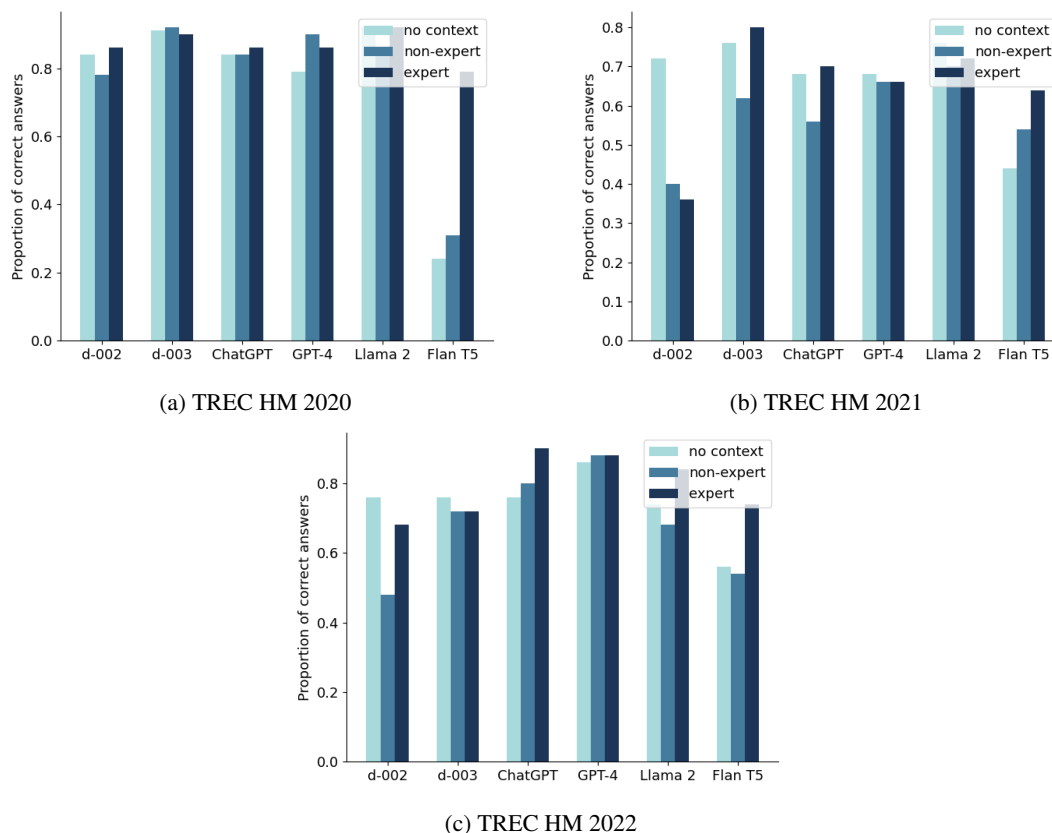

(a) TREC HM 2020

(b) TREC HM 2021

(c) TREC HM 2022

Figure 6: Zero-shot experiments. For each LLM, each bar represents the proportion of correct answers for the three prompting strategies ("no context", "non-expert" and "expert").

it comes to the TREC HM 2022 data, GPT-4 and ChatGPT stand out. Additionally, there are variances in performance linked to the use of different prompts. In general, the "expert" context is the most effective. We believe this effectiveness is due to the incorporation of key phrases like "research from PubMed" or "medical practice and literature", which guide the model towards more credible knowledge sources.

Although the models generally show relative stability, their performance still varies based on the inputs provided. For example, FlanT5 and text-davinci-002 are highly sensitive to the type of prompt. This raises concerns, as a model's proportion of correct answers can drop from 90% to 75% or even lower numbers. While the overall performance levels are high, these inconsistencies are troubling. Even when using the most reliable prompt ("expert"), there are still concerning situations. For instance, GPT-4's performance drops to 66% on the 2021 dataset.

The difficulty levels of the three datasets vary. The 2020 health questions (centered on COVID-19) seem easier for the large language models (LLMs). This could be due to the models' previous exposure to such health questions during

prompt	d002				d003			
	0-shot	1-shot	2-shot	3-shot	0-shot	1-shot	2-shot	3-shot
no-context	0.76	0.7	**0.78**	**0.78**	0.76	**0.86**	**0.86**	**0.86**
non-expert	0.48	**0.64**	**0.74**	**0.76**	0.72	**0.82**	**0.82**	**0.82**
expert	0.68	**0.74***	**0.76***	**0.78***	0.72	**0.82***	**0.84***	**0.84***

prompt	FT5				ChatGPT			
	0-shot	1-shot	2-shot	3-shot	0-shot	1-shot	2-shot	3-shot
no-context	0.56	**0.66**	**0.64**	**0.7**	0.76	**0.82**	**0.88**	**0.84**
non-expert	0.54	**0.68***	**0.66***	**0.64**	0.8	0.8	**0.88**	**0.86**
expert	0.74	**0.68***	0.72	0.72	0.9	0.84	0.88	0.88

prompt	Llama2				GPT-4			
	0-shot	1-shot	2-shot	3-shot	0-shot	1-shot	2-shot	3-shot
no-context	0.74	**0.7***	**0.84**	**0.76**	0.86	0.84	0.86	0.86
non-expert	0.68	**0.72**	**0.74**	**0.62***	0.86	0.86	**0.88**	**0.88**
expert	0.84	**0.64***	**0.76**	**0.6***	0.88	0.88	**0.92**	**0.9**

Table 2
Few-shot experiments. Proportion of correct answers of each model-prompt combination with up to three shot samples. For each row, if a few shot instance surpasses the 0-shot case then the few shot case is marked in bold and the symbol * marks those cases where McNemar's test ($\alpha = .05$) found a significant difference between both variants.

their extensive training. This hypothesis will be further investigated in Section 6.2.5. Another possible reason could be the significant relevance of COVID-19 as a topic, potentially encouraging specialized data curation processes.

We ran McNemar's test to ascertain the significance of performance differences between the leading models. The comparison between ChatGPT and GPT-4 showed no significant difference in 7 out of 9 cases (3 datasets x 3 prompts). The comparison between ChatGPT and Llama2 revealed no significant difference in 7 out of 9 cases, and the comparison between GPT-4 and Llama2 showed no significant differences at all. The pairwise comparisons of d-003 vs ChatGPT, d-003 vs Llama2, and d-003 vs GPT-4 produced a higher frequency of statistically significant outcomes, but, still, a majority of compared cases resulted in no significant differences.

Comparing the 0-shot results achieved by LLMs (Figure 6) with those achieved by SEs (Figure 3) we can observe that, in general, LLMs produce a higher proportion of correct answers. Even the user behaviour evaluation (Figure 5), which simulates users skipping non-answers from SE results, yields effectiveness statistics that are inferior to those obtained with LLMs. This suggests that the huge amounts of training data of the LLMs and their advanced reasoning capabilities are key advantages compared to the extraction of responses from a few top ranked search results.

6.2.2. Few-shot evaluation

To examine the impact of demonstrations on LLMs, we made additional tests with the health questions from TREC HM 2022. We sent each question to the models preceded by one to three demonstrations taken from TREC HM 2021. To this end, we randomly selected three pairs of *(medical question, correct answer)* from the 2021 dataset to serve as in-context examples and investigated their influence[12]. Previous studies have indicated that a small set of in-context examples suffices for instructing the LLMs (Liu et al., 2023).

Table 2 reports the proportion of correct answers when providing varying number of demonstrations (from 1-3). The impact of demonstrations varies significantly across different models. Specifically, both versions of GPT-3 (davinci-002 and davinci-003) and FlanT5 showed a clearly positive effect with in-context examples. For these models, the inclusion of some in-context variants resulted in statistically significant improvements. Conversely, the models that performed best in the zero-shot setting did not experience any noticeable benefit from the addition of demonstrations. When it comes to the type of prompts, the "expert" variant gained the most from the inclusion of few-shot examples. Furthermore, these experiments suggest that using one demonstration is good enough (adding more than one does not consistently improve performance).

---

[12] The selected question-answer pairs were: ("Will wearing an ankle brace help heal achilles tendonitis?", "No"), ("Does yoga help manage asthma?", "Yes"), and ("Is starving a fever effective?", "No")
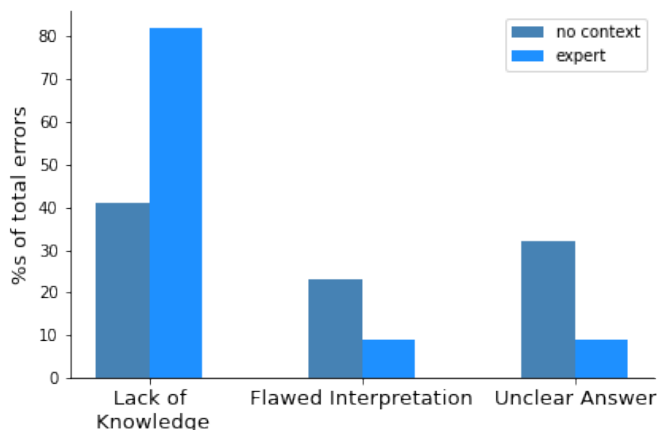
Figure 7: Percentages of the different type of errors for the "no context" and "expert" prompting strategies.

6.2.3. Error analysis

In order to gain a deeper understanding of the performance of the LLMs, we scrutinised the instances in which all models were unsuccessful in delivering a correct response. This analysis aimed at clarifying the reasons for such low effectiveness. Specifically, we sent these "low performance questions" again to the models but, this time, we did not set a limit to the token output. In this way, we could manually analyse the explanations provided and tried to understand the reasoning of the models. This examination was conducted with the most effective prompt ("expert") and with the "no context" prompt, which arguably mirrors the type of input submitted by a lay user.

For the TREC HM 2020 collection, 8% of the questions had a incorrect answer produced by all models using the "no context" prompt (and 6% of the questions had incorrect responses by all models with the "expert" prompt). The TREC HM 2021 collection had 12% of the questions with wrong answers by all models (for both prompts), while the TREC HM 2022 collection had 4% of questions where all LLMs were incorrect (with both prompts). These figures confirm that the TREC HM 2021 collection is the most challenging, with a higher percentage of errors. These numbers also reinforce the idea that providing no context is suboptimal, compared to using the expert prompt. And, regardless of the LLM, there are some low performance questions that pose serious difficulties to the models.

After manually reviewing the models' outputs, we were able to categorise the errors into a taxonomy that represents the most frequent health advice mistakes, namely:

- **Incorrect understanding of current medical consensus**: Sometimes, the models provided responses that directly contradict the medical consensus. For example, in response to the question "Can Hydroxychloroquine worsen COVID-19?", ChatGPT answered "no, there is no evidence that hydroxychloroquine worsens covid-19..." which contradicts medical evidence[13].

- **Misinterpretation of the question**: In this class of errors, the LLMs misunderstood the question. For instance, in response to "Can bleach prevent COVID-19?" the models produced completions such as "No, bleach should not be ingested...". However, the correct interpretation of this question is that the use of bleach for surface disinfection can indeed prevent COVID-19. A human would hardly interpret the question in the way that the LLM did.

- **Ambiguous answer**: This category includes responses where models did not provide a clear answer. These responses cannot be considered correct, but the LLM's output could arguably be useful. For example, "sit-ups can be both beneficial and harmful, depending on your individual circumstances and the way you perform the exercise...".

Figure 7 reports the percentage of each type of error for the different prompting conditions. In summary, the most common class of error for both prompting strategies is a lack of understanding about medical consensus. This is

---

[13]FDA cautions against use of hydroxychloroquine for COVID-19

(a) TREC HM 2020

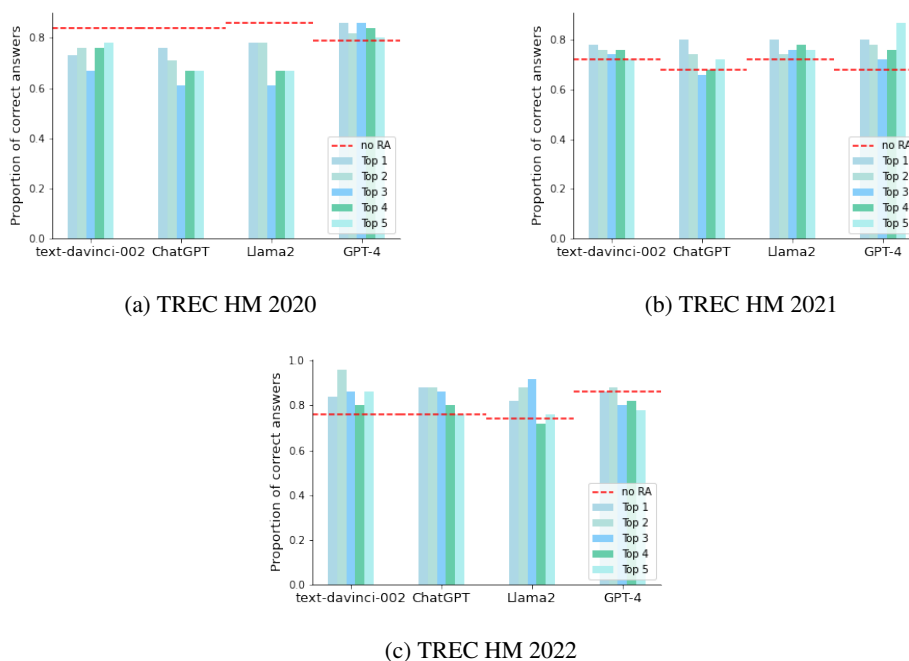(b) TREC HM 2021

(c) TREC HM 2022

Figure 8: Retrieval-augmented LLMs. Proportion of correct answers using the "no context" prompting strategy. Each "Top n" bar depicts the performance obtained by feeding the n-th result from Google.

worrying because this is the most harmful type of error. It also appears that providing "expert" context reduces potential misinterpretations of the questions and prevents ambiguous answers, but it also leads to more errors regarding current medical consensus.

6.2.4. Retrieval Augmented Language Models

As argued above, a stimulating line of research consists of combining the potential of LLMs and SEs. It is interesting to explore the effect of including extracts from SE results into the LLMs' instructions. To reduce the monetary cost and computational load of these experiments, we decided to extract the passages from a single engine (Google). The same passages from Google were fed to four LLMs: text-davinci-002, ChatGPT, GPT-4 and Llama2, under two prompting strategies ("no context" and "expert").

Figure 8 depicts the results of augmenting LLMs with each of the passages from Google's top 5 for the "no context" prompting strategy. For the 2020 questions, the LLMs do not seem to require these additional passages and, actually, three of the LLMs got worst results when presented with additional evidence. For the two other datasets, the LLMs seem to improve their performance with textual evidence from the search engine. It must be noticed that, in some cases (e.g., 2022), the provided passages make that the least sophisticated model, text-davinci-002, becomes comparable or even superior to more recent models, such as GPT-4. We perceive this as an important outcome of our study since it demonstrates that lighter models can achieve state-of-the-art performance when provided with additional evidence.

On the other hand, Figure 9 shows the behaviour of the retrieval augmented language models with the "expert" prompt. Tendencies remain similar, but we can highlight a larger increase in performance for text-davinci-002 in the 2021 collection. Its performance grows from 0.36 up to values above 0.70 with retrieval augmentation. Still, these experiments are not conclusive about the circumstances in which a LLM benefits from retrieval evidence. From our results, we can conclude that there seems not to be a consistent, clearly positive effect from prompting the LLMs with individual passages from the top 5 results. As a side note, we also evaluated the influence of augmenting with multiple passages. For example, we made tests with the first top 3 passages concatenated but, again, we obtained a mixed bag-of-results. In the future, it would be interesting to further explore RAG variants and the interactions among LLM complexity, types of prompts, size and variety of retrieval results, and types of health questions.
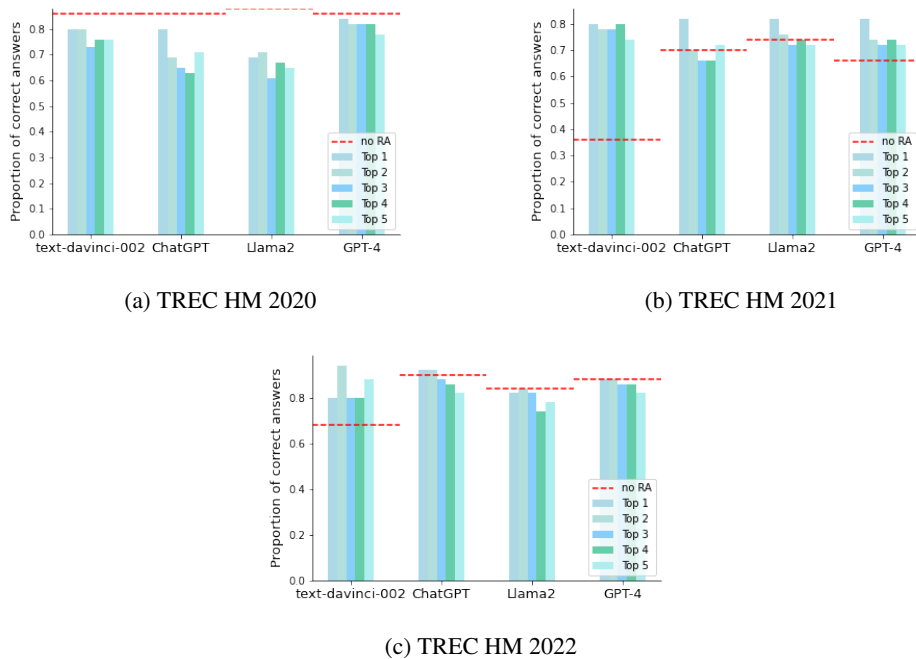
(a) TREC HM 2020

(b) TREC HM 2021

(c) TREC HM 2022

Figure 9: Retrieval-augmented LLMs. Proportion of correct answers using the "expert" prompting strategy. Each "Top n" bar depicts the performance obtained by feeding the n-th result from Google.

6.2.5. Memorisation

LLMs have demonstrated impressive performance across numerous language-related tasks. However, in some instances, this performance may be attributed to the inclusion of ground truth data from the evaluated benchmarks within the training datasets of the LLMs. This issue becomes particularly problematic with proprietary models that do not disclose their training data. Thus, there is no straightforward way to verify the sources of the training data. A fair assessment of these models must test their ability to generalise beyond the training data. A system that merely replicates the answer from an existing ground truth file should not be deemed as intelligent. A truly intelligent system learns from the training data and subsequently makes appropriate inferences to answer new questions. Drawing a parallel with education, a student who had access to the exam answers should fail, while a student who studied all the relevant material and provided correct answers should pass.

Memorisation is a thriving area of research that aims to determine whether a NLP benchmark was included in the pre-training process of a LLM (Nori et al., 2023; Golchin and Surdeanu, 2023; Magar and Schwartz, 2022; Sianz et al., 2023). As part of our investigation, we conducted memorisation estimation experiments to further validate the LLMs' ability to accurately answer medical and health-related questions.

We employed the heuristics proposed by Golchin and Surdeanu (2023), which were validated under controlled contamination experiments. Essentially, the method involves: i) prompting the model with a *general instruction* (with no information about the benchmark or specific identifiers of the task or split), ii) prompting the model with a *guided instruction* (which identifies the benchmark and task), and iii) comparing the responses against the actual ground truth text. If the model's output from the guided prompt is more similar to the ground truth text than the model's output from the general prompt, this suggests that the model may have been exposed to the benchmark during its training. A statistical significance test that compares the two means of similarity can thus identify those cases where we can suspect that the model ingested the benchmark during its pre-training. Observe that this is a high precision but low recall method. This means that if the method estimates that the model ingested the dataset, we can assert with some certainty that it did. However, if the heuristic does not provide evidence, we cannot definitively state that the model did not ingest it.

In our experiment, the general prompt was: "*Complete the narrative field based on the query, question and answer fields such that the narrative provides an explanation for the answer to the given question. Query: {query}, Question:*

Model	Version	TREC HM 2020			TREC HM 2021		
		Levenshtein	*BLEURT*	*ROUGE*	*Levenshtein*	*BLEURT*	*ROUGE*
ChatGPT	General	0.45	0.41	0.27	0.46	0.48	0.29
	Guided	0.30	0.39	0.10	0.38	0.47	0.23
GPT-4	General	0.45	0.42	0.26	0.45	0.46	0.29
	Guided	0.44	0.50*	0.24	0.42	0.51*	0.25
Llama2	General	0.42	0.37	0.23	0.44	0.43	0.28
	Guided	0.43*	0.35	0.24	0.45	0.40	0.27

Model	Version	TREC HM 2022		
		Levenshtein	*BLEURT*	*ROUGE*
ChatGPT	General	0.48	0.41	0.26
	Guided	0.49	0.42	0.27
GPT-4	General	0.47	0.41	0.24
	Guided	0.47	0.43	0.25
Llama2	General	0.41	0.32	0.14
	Guided	0.40	0.25	0.15

Table 3
Results of the memorisation experiments for different models and datasets. For each guided vs general comparison, the symbol * marks those cases where the guided completion surpassed the general completion and Wilcoxon test ($\alpha = .05$) found a significant difference between both variants.

{question}, Answer: {answer}, Narrative:". The guided prompt was: "*You are provided with the query, question and answer fields of a topic from the TREC {year} Health Misinformation topic set. Complete the narrative field of the topic as exactly appeared in the dataset. Only rely in the original form of the topic in the dataset to complete the narrative field. Query: {query}, Question: {question}, Answer: {answer}, Narrative:*". Note that the latter prompt explicitly mentions the name of the dataset and compels the model to generate the completion in its original format.

A pair of texts is thus obtained from these two completions, which are then compared against the narrative text from the actual benchmark. To calculate the overlaps (*general completion vs golden truth* and *guided completion vs golden truth*), we used the same metrics as in Golchin and Surdeanu (2023): BLEURT to gauge lexical similarity (Sellam et al. (2020)) and ROUGE-L (Lin and Och (2004)) to assess semantic relevance. In addition, we also calculated the Levenshtein distance, which quantifies the number of character permutations required to convert the completion into the original data. This comparative analysis was carried out for each available topic (i.e., each TREC question generated two completions and we report here the average similarity across all topics).

We performed this analysis for ChatGPT, GPT-4 and Llama2, as shown in Table 3. These estimated memorisation scores provide some indication that GPT-4 may have ingested TREC HM 2020 and TREC HM 2021 datasets, where we found statistically significant improvements for the guided completion compared to the general one (in terms of semantic similarity). The Levenshtein metric also suggests that Llama2 may have been trained with the TREC HM 2020 collection. For ChatGPT, we found no signs of having memorised any benchmark. ChatGPT performs on par with GPT-4 and Llama2 in the TREC HM 2020 and 2021 collections (under the zero-shot setting) and, in fact, there is no statistical difference between ChatGPT and these two models. This seems to suggest that ChatGPT's solid results in the TREC HM 2020 collection are not due to contamination effects but originate from the unique characteristics of the COVID-19 topics. Moreover, no model appears to have ingested the TREC HM 2022 collection and, still, many models performed effectively for this dataset. Overall, these findings highlight the ability of the LLMs to successfully transfer the knowledge gained during training and generate precise answers to health and medical questions. Although some specific instances of question-answer pairs might have been available to some of the models, our memorisation experiments suggest that in most of the cases the LLMs' responses do not derive from accessing the TREC topic and its ground truth answer.

7. Discussion

Search engines are the classic information access tools to retrieve contents from online sources. The first goal of our study was to estimate their capacity in providing correct answers to binary health questions (**RQ1**). We have evaluated four popular SEs and observed that the percentage of correct answers found in the SERPs is in the range of 50% to 70%. These figures are low and might sound highly concerning. However, this outcome is partially explained by the presence of many results that do not provide an answer. By focusing on the retrieved pages that provide an unequivocal answer to the reference health questions, we obtained much higher proportions of correct responses. Still, SE companies have room for improvement, as many top ranked webpages do actually contain harmful health recommendations (10%-15% of wrong answers). This is a natural consequence of the open and unmoderated nature of the web and we encourage developers of search technologies to further advance in the removal of low quality contents from their indexes and SERPs.

Regarding the presence of correct and incorrect answers over the ranked positions, the quality of the responses does not seem to decrease as we go deeper in the rankings (at least for the top 20 results). Our results also show that Bing seems to be the most solid choice among the four SEs. As part of the analysis, we also modelled two different search user behaviours, lazy and diligent. We found out that lazy behaviour, based on making decisions from the first observed response, is not inferior to a more diligent method based on acquiring and comparing three responses to the health questions.

Our next goal was to determine whether LLMs are reliable for providing accurate medical answers (**RQ2**). Our results suggest that, in general, the most capable LLMs generate better answers compared with those extracted from top webpages ranked by SEs. It seems that the extensive training data of the LLMs, coupled with their superior reasoning abilities, offer significant advantages over the extraction of responses from a handful of top-ranked search results. Among the largest models (GPT-4, ChatGPT, LLama2), we found no clear winner; and we observed poor performance from models such as FlanT5. Despite showing an overall good performance, there are still some barriers to the adoption of LLMs. For instance, under some circumstances, LLMs provide more than 30% of incorrect results.

Another concerning outcome is that the quality of the LLMs' completions in response to health questions was highly sensitive to the input prompt (**RQ3**). We found that some input prompts, which guide the models towards reputed sources, are much more effective than basic prompts (or prompts that give no context at all). But lay users would hardly resort to sophisticated prompts or complex interactions with the LLMs. This suggests that the future adoption of LLMs to support QA in the health domain would need to wrap the user's questions into automatically extended contexts (or, alternatively, design health-oriented assistants that guide the AIs towards high quality knowledge).

We also conducted a thorough error analysis and demonstrated that, even with the most sophisticated prompts, LLMs made errors due to a lack of medical knowledge. To validate our findings, we also conducted a set of experiments to demonstrate that LLMs are not "cheating" when answering health questions. Our memorisation experiments suggest that, at least for some of the question sets, there is no evidence that the LLMs saw the question-answer pairs during training and, thus, it appears that their answers come from their general knowledge about the health topic.

We also discovered that retrieval-augmented generation is promising for answering health questions (**RQ4**). We demonstrated that smaller LLMs reach the level of superior models when grounded in evidence provided by a search engine. This opens up the debate on whether it is worthwhile to persist in the generation of massive and computationally demanding models, or alternatively, we can direct our efforts towards leveraging lighter models enriched with search evidence.

We are aware that our study presents some limitations. For instance, the automatic extraction of answers from specific passages for evaluating search engines can be prone to error. However, the passage retrieval stage and the reading comprehension stage were both based on state-of-the-art technologies that were effectively tested elsewhere. We do not claim that the process is error-free but we are confident about the robustness of the trends found. In any case, in the near future we intend to further validate these findings by involving human evaluators to manually annotate the correctness of the search results. This is however a costly process that is not exempt from problems. Moreover, our current experiments did not consider the effect of personalisation. Retrieval results are known to be dependent on multiple user factors (e.g., geolocalization, user preferences, etc) and, thus, it will be important to study the relative quality of the medical responses taking into account geographical factors or other user-dependent variables.

About LLM evaluation, we are aware that using proprietary models poses the difficulty of reproducing these experiments. Their architectural design is unknown, the training data is not disclosed and, often, these are black-box models in constant evolution. However, we must acknowledge that these models are currently being used by millions of

people worldwide and, thus, it is of paramount importance to assess their answering capabilities in the health domain. Note that SEs can also be regarded as closed systems whose core elements (e.g., ranking algorithms) are largely unknown and, additionally, their indexes undergo continuous updates. Our experiments and comparison, therefore, give a specific picture of LLMs and SEs at a certain point in time. The forthcoming evolution of these systems will call for new trials and experiments to supplement the present study. In any case, we facilitate the reproducibility of our experiments by providing the code[14], the generated outputs, and all the dates of execution of the experiments[15].

Another limitation affects the type of information needs. In these experiments, we restricted ourselves to binary question answering. This was a practical decision because we then were dealing with yes/no responses whose correctness can be automatically assessed. This represents a valuable first step towards understanding the relative effectiveness of SEs and LLMs in the health domain. In the near future, we would like to explore other classes of information needs that require more elaborated answers.

8. Conclusions

In this paper, we have addressed a socially relevant problem: the ability of current information access tools to support information seeking in the health domain. Nowadays, the utilisation of conversational AIs for information access is on a steady rise. In our study, we have analysed their utility in providing support for health-related questions and, in particular, we have compared their performance with that of traditional web search systems.

In general, LLMs tend to produce better answers, compared with those obtained from inspecting a few top ranked results. This can be attributed to several factors, including the presence of non-relevant contents in the SERPs, the unmoderated nature of the web, and limitations derived from extracting responses from a small set of highly ranked webpages. LLMs get leverage from their massive amounts of training data and knowledge acquisition capabilities but, still, generative models sometimes provide concerning advice that goes against the medical consensus. We have also assessed an "online" retrieval-augmented generation strategy that prompts the LLMs with evidence retrieved from the search engines. This strategy produced solid results and we believe this may instigate a fascinating new trajectory of research that we aim to investigate further.

As future work, we will also consider more sophisticated prompting strategies, such as chain-of-thought (CoT) prompting. We also intend to explore alternative in-context learning methods. We will also consider more sophisticated retrieval-augmented strategies such as those based on soft incorporation in the middle layers of the models (Borgeaud et al., 2022).

References

Asai, A., Gardner, M., Hajishirzi, H., 2021. Evidentiality-guided generation for knowledge-intensive NLP tasks. arXiv preprint arXiv:2112.08688 .

Asai, A., Min, S., Zhong, Z., Chen, D., 2023. Retrieval-based language models and applications, in: Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 6: Tutorial Abstracts), pp. 41–46.

Baujard, O., Baujard, V., Aurel, S., Boyer, C., Appel, R., 1998. Trends in medical information retrieval on internet. Computers in Biology and Medicine 28, 589–601.

Bin, L., Lun, K., 2001. The retrieval effectiveness of medical information on the web. International journal of medical informatics 62, 155–163.

Bodaghi, A., Schmitt, K.A., Watine, P., Fung, B.C., 2023. A literature review on detecting, verifying, and mitigating online misinformation. IEEE Transactions on Computational Social Systems .

Borgeaud, S., Mensch, A., Hoffmann, J., Cai, T., Rutherford, E., Millican, K., Van Den Driessche, G.B., Lespiau, J.B., Damoc, B., Clark, A., et al., 2022. Improving language models by retrieving from trillions of tokens, in: International conference on machine learning, PMLR. pp. 2206–2240.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al., 2020. Language models are few-shot learners. Advances in neural information processing systems 33, 1877–1901.

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y.T., Li, Y., Lundberg, S., et al., 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. arXiv preprint arXiv:2303.12712 .

Chang, Y., Wang, X., Wang, J., Wu, Y., Yang, L., Zhu, K., Chen, H., Yi, X., Wang, C., Wang, Y., et al., 2024. A survey on evaluation of large language models. ACM Transactions on Intelligent Systems and Technology 15, 1–45.

Chervenak, J., Lieman, H., Blanco-Breindel, M., Jindal, S., 2023. The promise and peril of using a large language model to obtain clinical information: ChatGPT performs strongly as a fertility counseling tool with limitations. Fertility and Sterility .

Chuklin, A., Markov, I., de Rijke, M., 2015. Click Models for Web Search. Springer International Publishing. URL: https://doi.org/10.1007/978-3-031-02294-4, doi:10.1007/978-3-031-02294-4.

Clarke, C., Maistro, M., Smucker, M., 2021. Overview of the trec 2021 health misinformation track, in: Proceedings of the Thirtieth Text REtrieval Conference, TREC.

---

[14]https://github.com/MarcosFP97/llm-binary-health-qa
[15]These experiments were run between September and December 2023.

Clarke, C., Maistro, M., Smucker, M., Zuccon, G., 2020. Overview of the trec 2020 health misinformation track, in: Proceedings of the Twenty-Nine Text REtrieval Conference, TREC, pp. 16–19.

Dijkstra, R., Genç, Z., Kayal, S., Kamps, J., et al., 2022. Reading comprehension quiz generation using generative pre-trained transformers., in: iTextbooks@ AIED, pp. 4–17.

Duong, D., Solomon, B.D., 2023. Analysis of large-language model versus human performance for genetics questions. European Journal of Human Genetics , 1–3.

Fernández-Pichel, M., Bink, M., Losada, D.E., Elsweiler, D., 2024. A user study on people's perception to the credibility of online health information. Proceedings of ROMCIR .

Fernández-Pichel, M., Losada, D.E., Pichel, J.C., Elsweiler, D., 2021a. Comparing traditional and neural approaches for detecting health-related misinformation, in: Int. Conf. of the Cross-Language Evaluation Forum for European Languages, Springer. pp. 78–90.

Fernández-Pichel, M., Losada, D.E., Pichel, J.C., Elsweiler, D., 2021b. Reliability prediction for health-related content: a replicability study, in: European Conference on Information Retrieval, Springer. pp. 47–61.

Fernández-Pichel, M., Losada, D.E., Pichel, J.C., 2024. Large language models for binary health-related question answering: A zero- and few-shot evaluation, in: Proceedings of the 24th International Conference on Computational Science, pp. 325–339.

Fogg, B.J., 1999. Persuasive technologie301398. Communications of the ACM 42, 26–29.

Fogg, B.J., 2003. Prominence-interpretation theory: Explaining how people assess credibility online, in: CHI'03 extended abstracts on human factors in computing systems, pp. 722–723.

Fox, S., 2011. Health topics: 80% of internet users look for health information online. Pew Internet & American Life Project.

Friedman, L., Ahuja, S., Allen, D., Tan, T., Sidahmed, H., Long, C., Xie, J., Schubiner, G., Patel, A., Lara, H., et al., 2023. Leveraging large language models in conversational recommender systems. arXiv preprint arXiv:2305.07961 .

Ginsca, A.L., Popescu, A., Lupu, M., et al., 2015. Credibility in information retrieval. Foundations and Trends in Information Retrieval 9, 355–475.

Goeuriot, L., Suominen, H., Kelly, L., Miranda-Escalada, A., Krallinger, M., Liu, Z., Pasi, G., Gonzalez Saez, G., Viviani, M., Xu, C., 2020. Overview of the clef ehealth evaluation lab 2020, in: International Conference of the Cross-Language Evaluation Forum for European Languages, Springer. pp. 255–271.

Golchin, S., Surdeanu, M., 2023. Time travel in llms: Tracing data contamination in large language models. arXiv preprint arXiv:2308.08493 .

Griffiths, K.M., Tang, T.T., Hawking, D., Christensen, H., 2005. Automated assessment of the quality of depression websites. Journal of Medical Internet Research 7, e59.

Guu, K., Lee, K., Tung, Z., Pasupat, P., Chang, M., 2020. Retrieval augmented language model pre-training, in: International Conference on Machine Learning, PMLR. pp. 3929–3938.

Hahnel, C., Goldhammer, F., Kröhne, U., Naumann, J., 2018. The role of reading skills in the evaluation of online information gathered from search engine environments. Computers in Human Behavior 78, 223–234.

Hamidi, A., Roberts, K., 2023. Evaluation of ai chatbots for patient-specific ehr questions. arXiv preprint arXiv:2306.02549 .

Holmes, J., Liu, Z., Zhang, L., Ding, Y., Sio, T.T., McGee, L.A., Ashman, J.B., Li, X., Liu, T., Shen, J., et al., 2023. Evaluating large language models on a highly-specialized topic, radiation oncology physics. arXiv preprint arXiv:2304.01938 .

Izacard, G., Grave, É., 2021. Leveraging passage retrieval with generative models for open domain question answering, in: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, pp. 874–880.

Jahan, I., Laskar, M.T.R., Peng, C., Huang, J., 2023. Evaluation of chatgpt on biomedical tasks: A zero-shot comparison with fine-tuned generative transformers. arXiv preprint arXiv:2306.04504 .

Jiang, Z., Xu, F.F., Araki, J., Neubig, G., 2020. How can we know what language models know? Transactions of the Association for Computational Linguistics 8, 423–438.

Johnson, D., Goodman, R., Patrinely, J., Stone, C., Zimmerman, E., Donald, R., Chang, S., Berkowitz, S., Finn, A., Jahangir, E., et al., 2023. Assessing the accuracy and reliability of ai-generated medical responses: an evaluation of the chat-gpt model .

Kakol, M., Nielek, R., Wierzbicki, A., 2017. Understanding and predicting web content credibility using the content credibility corpus. Information Processing & Management 53, 1043–1061.

Kelly, L., Suominen, H., Goeuriot, L., Neves, M., Kanoulas, E., Li, D., Azzopardi, L., Spijker, R., Zuccon, G., Scells, H., et al., 2019. Overview of the clef ehealth evaluation lab 2019, in: Experimental IR Meets Multilinguality, Multimodality, and Interaction: 10th International Conference of the CLEF Association, CLEF 2019, Lugano, Switzerland, September 9–12, 2019, Proceedings 10, Springer. pp. 322–339.

Lazaridou, A., Gribovskaya, E., Stokowiec, W., Grigorev, N., 2022. Internet-augmented language models through few-shot prompting for open-domain question answering. arXiv preprint arXiv:2203.05115 .

Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C.H., Kang, J., 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. Bioinformatics 36, 1234–1240.

Li, H., Su, Y., Cai, D., Wang, Y., Liu, L., 2022. A survey on retrieval-augmented text generation. arXiv preprint arXiv:2202.01110 .

Li, Y., Li, Z., Zhang, K., Dan, R., Jiang, S., Zhang, Y., 2023. Chatdoctor: A medical chat model fine-tuned on a large language model meta-ai (llama) using medical domain knowledge. Cureus 15.

Liang, P., Bommasani, R., Lee, T., Tsipras, D., Soylu, D., Yasunaga, M., Zhang, Y., Narayanan, D., Wu, Y., Kumar, A., et al., 2022. Holistic evaluation of language models. arXiv preprint arXiv:2211.09110 .

Liao, Q.V., Fu, W.T., 2014. Age differences in credibility judgments of online health information. ACM Transactions on Computer-Human Interaction (TOCHI) 21, 1–23.

Lin, C.Y., Och, F., 2004. Looking for a few good metrics: Rouge and its evaluation, in: Ntcir workshop.

Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., Neubig, G., 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. ACM Computing Surveys 55, 1–35.

Longpre, S., Hou, L., Vu, T., Webson, A., Chung, H.W., Tay, Y., Zhou, D., Le, Q.V., Zoph, B., Wei, J., et al., 2023. The flan collection: Designing data and methods for effective instruction tuning. arXiv preprint arXiv:2301.13688 .

Magar, I., Schwartz, R., 2022. Data contamination: From memorization to exploitation. arXiv preprint arXiv:2203.08242 .

Mao, K., Dou, Z., Chen, H., Mo, F., Qian, H., 2023. Large language models know your contextual search intent: A prompting framework for conversational search. arXiv preprint arXiv:2303.06573 .

Matthews, S.C., Camacho, A., Mills, P.J., Dimsdale, J.E., 2003. The internet for medical information about cancer: help or hindrance? Psychosomatics 44, 100–103.

McKnight, D.H., Kacmar, C.J., 2007. Factors and effects of information credibility, in: Proceedings of the ninth international conference on Electronic commerce, pp. 423–432.

Nakano, R., Hilton, J., Balaji, S., Wu, J., Ouyang, L., Kim, C., Hesse, C., Jain, S., Kosaraju, V., Saunders, W., et al., 2021. Webgpt: Browser-assisted question-answering with human feedback. arXiv preprint arXiv:2112.09332 .

Nogueira, R., Cho, K., 2019. Passage re-ranking with bert. arXiv preprint arXiv:1901.04085 .

Nori, H., King, N., McKinney, S.M., Carignan, D., Horvitz, E., 2023. Capabilities of gpt-4 on medical challenge problems. arXiv preprint arXiv:2303.13375 .

OpenAI, 2023. Gpt-4 technical report. arXiv:submit/4812508 .

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al., 2022. Training language models to follow instructions with human feedback. Advances in Neural Information Processing Systems 35, 27730–27744.

O'Leary, D.E., 2022. Massive data language models and conversational artificial intelligence: Emerging issues. Intelligent Systems in Accounting, Finance and Management 29, 182–198.

Pogacar, F.A., Ghenai, A., Smucker, M.D., Clarke, C.L., 2017. The positive and negative influence of search results on people's decisions about the efficacy of medical treatments, in: Proceedings of the ACM SIGIR Int. Conf. on Theory of Information Retrieval, pp. 209–216.

Polak, M.P., Morgan, D., 2023. Extracting accurate materials data from research papers with conversational language models and prompt engineering–example of chatgpt. arXiv preprint arXiv:2303.05352 .

Pradeep, R., Chen, H., Gu, L., Tamber, M.S., Lin, J., 2023. Pygaggle: A gaggle of resources for open-domain question answering, in: Advances in Information Retrieval: 45th European Conference on Information Retrieval, ECIR 2023, Dublin, Ireland, April 2–6, 2023, Proceedings, Part III, Springer. pp. 148–162.

Pradeep, R., Lin, J., 2024. Towards automated end-to-end health misinformation free search with a large language model, in: Goharian, N., Tonellotto, N., He, Y., Lipani, A., McDonald, G., Macdonald, C., Ounis, I. (Eds.), Advances in Information Retrieval, Springer Nature Switzerland, Cham. pp. 78–86.

Pradeep, R., Ma, X., Nogueira, R., Lin, J., 2021. Vera: Prediction techniques for reducing harmful misinformation in consumer health search, in: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 2066–2070.

Reuters Insitute, University of Oxford, 2023. Reuters digital news report 2023. URL: https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2023. [accessed February 13, 2024].

Rosa, G.M., Bonifacio, L., Jeronymo, V., Abonizio, H., Fadaee, M., Lotufo, R., Nogueira, R., 2022. No parameter left behind: How distillation and model size affect zero-shot retrieval. arXiv preprint arXiv:2206.02873 .

Samaan, J.S., Yeo, Y.H., Rajeev, N., Hawley, L., Abel, S., Ng, W.H., Srinivasan, N., Park, J., Burch, M., Watson, R., et al., 2023. Assessing the accuracy of responses by the language model chatgpt to questions regarding bariatric surgery. Obesity surgery , 1–7.

Schwarz, J., Morris, M., 2011. Augmenting web pages and search results to support credibility assessment, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 1245–1254.

Sellam, T., Das, D., Parikh, A.P., 2020. Bleurt: Learning robust metrics for text generation. arXiv preprint arXiv:2004.04696 .

Sharma, K., Qian, F., Jiang, H., Ruchansky, N., Zhang, M., Liu, Y., 2019. Combating fake news: A survey on identification and mitigation techniques. ACM Transactions on Intelligent Systems and Technology (TIST) 10, 1–42.

Shuster, K., Komeili, M., Adolphs, L., Roller, S., Szlam, A., Weston, J., 2022. Language models that seek for knowledge: Modular search & generation for dialogue and prompt completion. arXiv preprint arXiv:2203.13224 .

Sianz, O., Campos, J.A., García-Ferrero, I., Etxaniz, J., Agirre, E., 2023. Did chatgpt cheat on your test? URL: https://hitz-zentroa.github.io/lm-contamination/blog/. [accessed January 19, 2024].

Soldaini, L., Yates, A., Yom-Tov, E., Frieder, O., Goharian, N., 2016. Enhancing web search in the medical domain via query clarification. Information Retrieval Journal 19, 149–173.

Sondhi, P., Vydiswaran, V.V., Zhai, C., 2012. Reliability prediction of webpages in the medical domain, in: European Conference on Information Retrieval, Springer. pp. 219–231.

Suominen, H., Kelly, L., Goeuriot, L., Névéol, A., Ramadier, L., Robert, A., Kanoulas, E., Spijker, R., Azzopardi, L., Li, D., et al., 2018. Overview of the clef ehealth evaluation lab 2018, in: Experimental IR Meets Multilinguality, Multimodality, and Interaction: 9th International Conference of the CLEF Association, CLEF 2018, Avignon, France, September 10-14, 2018, Proceedings 9, Springer. pp. 286–301.

Thirunavukarasu, A.J., Hassan, R., Mahmood, S., Sanghera, R., Barzangi, K., El Mukashfi, M., Shah, S., 2023. Trialling a large language model (chatgpt) in general practice with the applied knowledge test: observational study demonstrating opportunities and limitations in primary care. JMIR Medical Education 9, e46599.

Thoppilan, R., De Freitas, D., Hall, J., Shazeer, N., Kulshreshtha, A., Cheng, H.T., Jin, A., Bos, T., Baker, L., Du, Y., et al., 2022. Lamda: Language models for dialog applications. arXiv preprint arXiv:2201.08239 .

Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al., 2023. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 .

Vigdor, N., 2020. Man fatally poisons himself while self-medicating for coronavirus, doctor says. URL: https://www.nytimes.com/2020/03/24/us/chloroquine-poisoning-coronavirus.html. [accessed June 9, 2022].

Viviani, M., Pasi, G., 2017. Credibility in social media: opinions, news, and health information—a survey. Wiley interdisciplinary reviews: Data mining and knowledge discovery 7, e1209.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., ichter, b., Xia, F., Chi, E., Le, Q.V., Zhou, D., 2022. Chain-of-thought prompting elicits reasoning in large language models, in: Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., Oh, A. (Eds.), Advances in Neural Information Processing Systems, Curran Associates, Inc.. pp. 24824–24837. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf.

Yunxiang, L., Zihan, L., Kai, Z., Ruilong, D., You, Z., 2023. Chatdoctor: A medical chat model fine-tuned on llama model using medical domain knowledge. arXiv preprint arXiv:2303.14070 .

Zhao, Y., Da, J., Yan, J., 2021. Detecting health misinformation in online health communities: Incorporating behavioral features into machine learning based approaches. Information Processing & Management 58, 102390.

A. Algorithms for User Behaviour

Algorithms 1 and 2 show the pseudo-code for the two user behaviour models:

Algorithm 1: Lazy User Model

Input: Ranked list of SERP results S, Query q, Correct Answer a
Output: Correct response or Incorrect response or No answer, and Effort Required
$effort \leftarrow 0$
foreach *entry s in S* **do**
  $effort \leftarrow effort + 1$
  if *s provides an answer to q* **then**
    if *s answer matches a* **then**
      | **return** *Correct response*, *effort*
    end
    else
      | **return** *Incorrect response*, *effort*
    end
  end
end
return *No answer*, *effort*

Algorithm 2: Diligent User Model

Input: Ranked list of SERP results S, Query q, Correct Answer a
Output: Correct response or Incorrect response or No answer, and Effort
$effort \leftarrow 0$
$correct\_responses \leftarrow 0$
$incorrect\_responses \leftarrow 0$
foreach *entry s in S* **do**
  $effort \leftarrow effort + 1$
  if *s provides an answer to q* **then**
    if *s answer matches a* **then**
      | $correct\_responses \leftarrow correct\_responses + 1$
    end
    else
      | $incorrect\_responses \leftarrow incorrect\_responses + 1$
    end
  end
  if $correct\_responses + incorrect\_responses = 3$ **then**
    if $correct\_responses \geq 2$ **then**
      | **return** *Correct response*, *effort*
    end
    else
      | **return** *Incorrect responses*, *effort*
    end
  end
end
if $correct\_responses \neq 0$ *and* $correct\_responses \leq 2$ *and* $incorrect\_responses = 0$ **then**
  | **return** *Correct response*, *effort*
end
else if $incorrect\_responses \neq 0$ *and* $incorrect\_responses \leq 2$ *and* $correct\_responses = 0$ **then**
  | **return** *Incorrect response*, *effort*
end
return *No answer*, *effort*   // In the case of a tie with only two answers (one correct response and one incorrect), we return "No answer"