

Article

Not peer-reviewed version

---

# Retrieval Augmented Generation for Large Language Models in Healthcare: A Systematic Review

---

[Lameck Mbangula Amugongo](#) <sup>\*</sup>, [Pietro Mascheroni](#), [Steven Geoffrey Brooks](#), Stefan Doering, [Jan Seidel](#)

Posted Date: 11 July 2024

doi: [10.20944/preprints202407.0876.v1](https://doi.org/10.20944/preprints202407.0876.v1)

Keywords: Retrieval Augmented Generation; Large Language Models; Healthcare; Evaluation; AI ethics



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

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

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

## Article

# Retrieval Augmented Generation for Large Language Models in healthcare: A Systematic Review

Lameck Mbangula Amugongo <sup>1\*</sup>, Pietro Mascheroni <sup>1</sup>, Steven Brooks <sup>2</sup>, Stefan Doering <sup>1</sup> and Jan Seidel <sup>1</sup>

<sup>1</sup> Biostatistics and Data Sciences Department, Boehringer Ingelheim Pharma GmbH & Co. KG, Birkendorfer Str. 65, Biberach an der Riß, 88400, Germany

<sup>2</sup> Biostatistics and Data Sciences Department, Boehringer Ingelheim (China) Investment Co., Ltd, Shanghai, China

\* Correspondence: mbangula\_lameck.amugongo@boehringer-ingelheim.com

**Abstract:** Large Language Models (LLMs) have demonstrated promising capabilities to solve complex tasks in critical sectors such as healthcare. However, LLMs are limited by their training data which is often outdated, the tendency to generate inaccurate ("hallucinated") content and a lack of transparency in the content they generate. To address these limitations, retrieval augmented generation (RAG) grounds the responses of LLMs by exposing them to external knowledge sources. However, in the healthcare domain there is currently a lack of systematic understanding of which datasets, RAG methodologies and evaluation frameworks are available. This review aims to bridge this gap by assessing RAG-based approaches employed by LLMs in healthcare, focusing on the different steps of retrieval, augmentation and generation. Additionally, we identify the limitations, strengths and gaps in the existing literature. Our synthesis shows that proprietary models such as GPT-3.5/4 are the most used for RAG applications in healthcare. Also, there is a lack of standardized evaluation frameworks for RAG-based applications. In addition, the majority of the studies do not assess or address ethical considerations related to RAG in healthcare. It is important to account for ethical challenges that are inherent when AI systems are implemented in the clinical setting. Lastly, we highlight the need for further research and development to ensure responsible and effective adoption of RAG in the medical domain.

**Keywords:** retrieval augmented generation; large language models; healthcare; evaluation; ai ethics

## 1. Introduction

Large Language Models (LLMs) have revolutionised natural language processing (NLP) tasks in various domains, including healthcare. For example, models such as Generative Pre-trained Transformers (GPT) [1,2], LLaMA[3] and Gemini [4], have shown impressive capabilities in generating coherent and contextually relevant text. However, their application in healthcare is hampered by critical limitations, such as the propensity to generate inaccurate or nonsensical information [5]. This issue is often referred to as "model hallucinations" [6] and methodologies for its mitigation are still an active area of research [7].

In healthcare, several LLMs have been customised to aid in different medical tasks. Models such as BioBERT [8] and ClinicalBERT [9] have been proposed, leveraging the power of Bidirectional Encoder Representations from Transformers (BERT) [10]. These models are developed through fine-tuning using biomedical texts with the aim to improve contextual language comprehension within the medical domain. However, they occasionally encounter challenges when dealing with contextual data. To address this contextual need in medicine, Med-Palm was introduced demonstrating good performance in retrieving clinical knowledge and excelling in decision-making on several clinical tasks [11]. However, Med-Palm could not outperform human clinicians, generated bias and returned harmful answers.

To address the aforementioned limitations, a novel approach called Retrieval-Augmented Generation (RAG) was proposed to expose the model to external knowledge sources [12]. RAG combines the power of LLMs with the ability to retrieve relevant information from external knowledge sources, such as medical databases, literature repositories, or expert systems. Briefly, the RAG process involves

retrieving relevant information from the knowledge source, and then using the relevant information to generate a response to answer the question. By incorporating a retrieval step, RAG leverages on in-context learning [13] to reduce hallucinations and enhance the transparency of the sources from which the LLM completion is generated. This is particularly important in healthcare, a knowledge-intensive domain that requires accurate, up-to-date, and domain-specific information [14]. In addition, by incorporating up-to-date clinical data and reliable medical sources such as clinical guidelines into LLMs, the latter can offer more personalised patient advice, quicker diagnostic and treatment suggestions and significantly enhance patient outcomes [15].

Despite the growth in RAG related research, we only came across a few review studies that outline the state-of-the-art in RAG [16] and methodologies for retrievers and generators [17]. To the best of our knowledge, there is no comprehensive review on methodologies and application of RAG for LLMs in the healthcare domain.

This review aims to fill this knowledge gap, providing a systematic analysis of RAG techniques in the medical setting. We examine different architectures and evaluation frameworks, and explore the potential benefits and challenges associated with the integration of retrieval-based methods. Finally, we propose future research directions and comment on open issues of current RAG implementations. Our contributions:

1. Provide a systematic review of RAG-based methodologies applied in the medical domain. Therefore contextualising the scope of RAG approaches in healthcare.
2. Provide an overview of evaluation methods, including metrics used to evaluate the performance of RAG pipelines.
3. Discuss ethical concerns associated with RAG pipelines in critical sectors such as healthcare.
4. Provide insights for future research directions in RAG-based applications.

## 2. Materials and Methods

Many researchers have proposed RAG as a way to provide LLMs with up-to-date and user-specific information, not available as part of the LLM's pre-trained knowledge (also known as "grounding the LLM") [18–24]. Our goal is to integrate the existing literature and evaluate the state-of-the-art RAG techniques used in healthcare. Thus, conducting a systematic literature review is a promising method to explore our objective. Moreover, our aim is to enhance the current knowledge about RAG for LLMs in healthcare. We intend to achieve this by employing a systematic and transparent methodology that produces reproducible results. For this, we employed the Preferred Reporting Items for Systematic reviews and Meta-Analyses (PRISMA) [25]. All studies included met the inclusion criteria shown in Table 1.

**Table 1.** Summary of inclusion criteria used in this study.

Criteria	Definition
Languages of articles	English
Years of publication	2020 - 2024
Solutions considered	<ul style="list-style-type: none"> <li>• Retrieval-Augmented Generation (RAG)</li> <li>• Large Language Models (LLMs)</li> </ul>

As illustrated in Table 1, the initial criteria used to identify articles, includes 1) only articles available in the English language; 2) articles published between January 2020 and March 2024, including archives with an arxivid or medrxiv id; 3) only papers proposing RAG-based methods applied in the medical domain. Articles that did not meet the above criteria were excluded. In addition, we also excluded articles that met the criteria below.

### *Exclusion Criteria*

- Review articles, including surveys, comprehensive reviews or systematic reviews.

- Papers for which the full text is not available.
- Short conference papers.

### 2.1. Search Technique

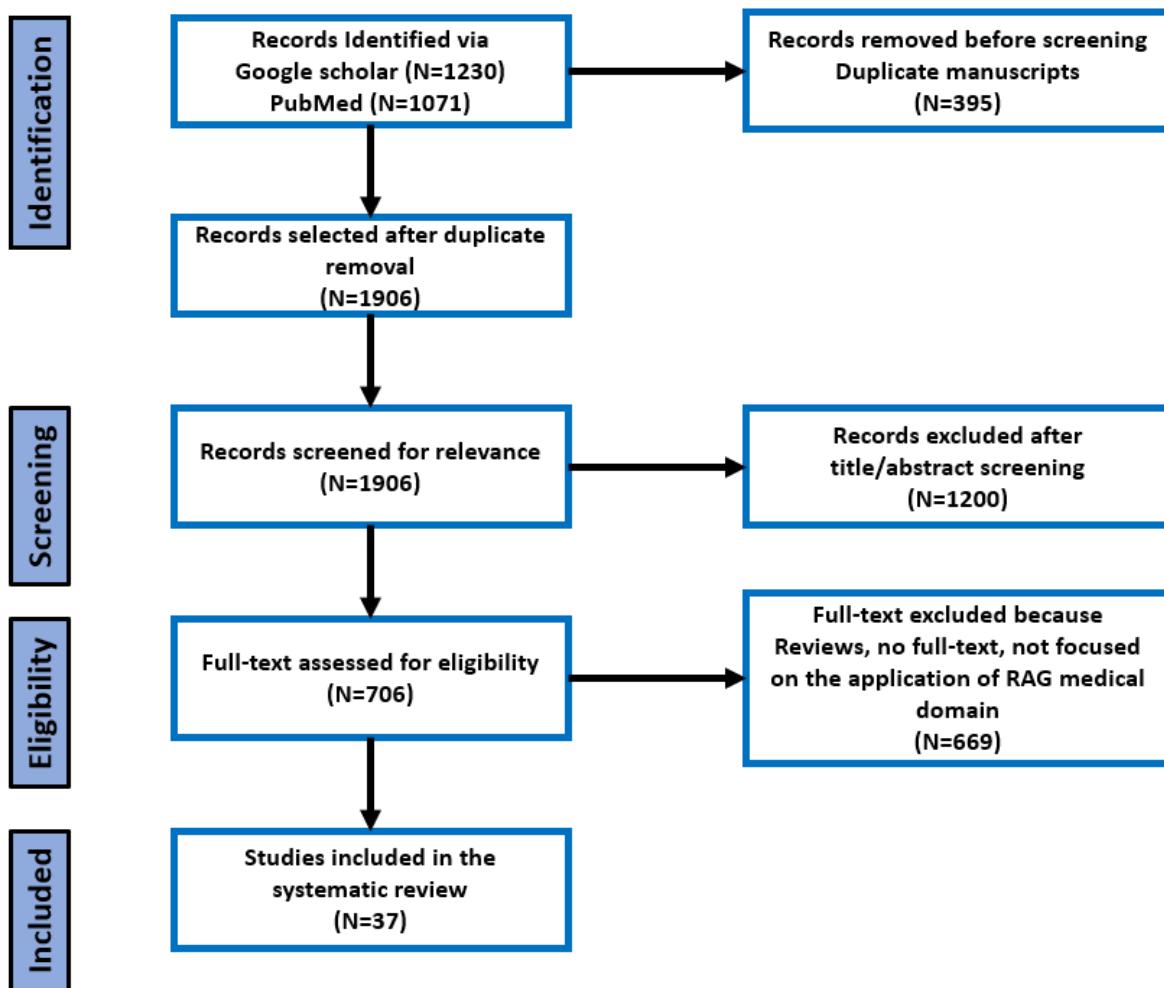
First, we carried out a scoping review using Google Scholar and PubMed to identify and retrieve articles that proposed the application of RAG in the medical domain. The fields considered in the search included the title, abstract, and the article itself. The search terms used are available in Table 2. We used specific search terms to retrieve more relevant articles.

**Table 2.** The keywords used to query the selected databases.

Database	Search keywords
PubMed	(large language models OR LLMs OR "transformer models" OR "Generative AI") AND (healthcare OR medicine OR medical) AND (retrieval OR augmented OR generation OR grounded)
Google Scholar	("Large Language Models" OR LLMs OR Transformer Models" OR "Generative Models") AND (Retrieval-Augmented Generation OR grounding) AND (healthcare OR medical OR medicine)

### 2.2. Study Selection

After identification, the articles were imported in the ReadCube (Digital Science & Research Solutions Inc, Cambridge, MA 02139, USA) literature management software to create a database of references. We used a three-step process for the selection of articles to be included in this study: relevance of the title, relevance of the abstract and finally relevance of the full-text [26]. This process ensured that only papers that met our eligibility criteria were reviewed. Figure 1 illustrates the process used for screening and determining the eligibility and exclusion criteria.



**Figure 1.** PRISMA [25] workflow applied to the search, identification and selection of the studies that were included in the systematic review.

### 2.3. Data Extraction

For each of the selected studies, we extracted the following information: 1) LLMs, 2) embedding (a numerical representation of information), 3) pre-retrieval, 4) post-retrieval, 5) advanced methodologies and 6) outcomes. In addition, we critically evaluated the technique used to assess the performance of the RAG-based application in the medical domain, including ethical concerns such as privacy, safety, robustness, bias, and trust (explainability/interpretability). Lastly, we performed analyses on data extracted from the papers surveyed in this study.

## 3. Results

### 3.1. Included Studies

We selected and included 37 studies between 2020 – 2024. The articles were selected from 2,301 articles retrieved from Google Scholar and PubMed after multiple exclusion steps. Lastly, the compressive review includes studies that employed innovative RAG-based approaches addressing questions such as “what information to retrieve”, “when to retrieve” and “how to use the retrieved information” to ground LLMs in the medical domain.

### 3.2. Datasets

Relevant knowledge sources are essential for LLMs to generate clinically correct responses to users' queries. Several datasets have been proposed in the literature to augment the responses of LLMs

in healthcare. Generally, existing retrieval datasets are divided into two categories: question-answering (QA) and information retrieval. As seen in Table 3, most of the datasets used in the studies we reviewed are aimed at QA for medical dialogue. QA dataset provide a short concise answers. On the other hand, information retrieval datasets are geared towards extracting and presenting relevant information to the user's query (generally extracted from large datasets) [27]. The major sources of the dataset used in the studies we surveyed include Web, PubMed, and Unified Medical Language System (UMLS).

Among the available datasets, MedDialog [28] contains conversations and utterances between patients and doctors in both English and Chinese, sourced from two websites, namely healthcaremagic.com and iclinic.com. MedDialog dataset covers 96 different diseases. In another study, a name entity dataset called MedDG was created from 17,000 conversations related to 12 types of common gastrointestinal diseases, collected from an online health consultation community [29]. Besides QA datasets, some researchers have curated information-retrieval datasets.

**Table 3.** Detailed information datasets used in retrieval to augment responses of LLMs in the medical domain.

Category	Author	Domain	Dataset	#Q in Train	#label in train	#q in dev	#q in test	#instance
Question Answering	Tsatsaronis et al. [30]	Biomedical	BioASQ	3,743	35,285		497	15,559,157
	Chen et al. [28]	Clinical	MedDialog					EN: 257,454
	Liu et al. [29]	Gastrointestinal	MedDG	14,864	-	2,000	1,000	C: 1,145,231
	Abacha et al. [31]	Biomedical	LiveQA	634	-	-	104	
	Zakka et al. [24]	Biomedical	ClinicalQA				130	
	Lozano et al. [32]	Biomedical	PubMedRS-200				200	
	Jin et al. [33]	Biomedical	PubMedQA		1,000		1,000	
	Jin et al. [34]	Biomedical	MedQA	USMLE 10,178 MCMLE 27,400 TWMLE 11,298	1272 3425 1412	1273 3426 1413	12,723 34,251 14,123	
	Ma et al. [35]	Orthodontic	MD-QA					59,642
	Chen et al. [36]	10 pediatric diseases	IMCS-21	-	-	-	-	4,116
	Zeng[37]	Biomedical	MMCU-Medical					2,819
Information Retrieval	Boteva et al. [38]	Biomedical	NFCorpus	5,922	110,575	24	323	3,633
	Roberts et al. [39]	COVID-19	TREC- COVID-19	-	-	-	-	
	Johnson et al. [40]	Radiology	MIMIC-CXR	-	-	-	-	Img: 377,110 Txt: 227,927
	Ramesh et al. [41]	Radiology	Adapted MIMIC-CXR	-	-	-	-	226,759

\* Abbreviations: Biomedical semantic indexing and Question Answering (BioASQ), Chinese (C), English (EN), Images (Img), United States Medical Licensing Examination (USMLE), Text (Txt). #q represents the number of queries and #instance the number of texts in the dataset.

For example, [41] curated a dataset containing 226,759 reports derived from the MIMIC-CXR dataset of X-ray images and radiology report text. Another retrieval dataset is the TREC-COVID-19 dataset, curated using data from medical library searches, MedlinePlus logs, and Twitter posts on COVID by high-profile researchers [39]. Guo et al. [42] presented a dataset linking scientific abstracts to expert-authored lay language summaries. The dataset, named CELLS, was generated with the aim of improving lay language generation models and includes 62,886 source-target pairs from 12 journals.

Benchmark datasets have been released to evaluate the retrieval abilities of LLMs. The Biomedical QA (BioASQ) dataset was proposed through a series of competitions as a benchmark dataset to assess systems and methodologies for large scale medical semantic indexing and QA tasks [30]. Other evaluation datasets for QA tasks include MedMCQA [43], PubMedQA [33] and MedQA [34]. Benchmark datasets such as MedMCQA, PubMedQA and alike do not include broad medical knowledge, and thus lack the detail required for real world clinical applications. To address this limitation, MultiMedQA was created incorporating seven medical QA datasets, including six existing datasets, namely: MedQA, MedMCQA, PubMedQA, LiveQA 2017, MedicationQA [44] and MMLU [45] clinical topics, and a new dataset comprising of the most searched medical questions on the internet.

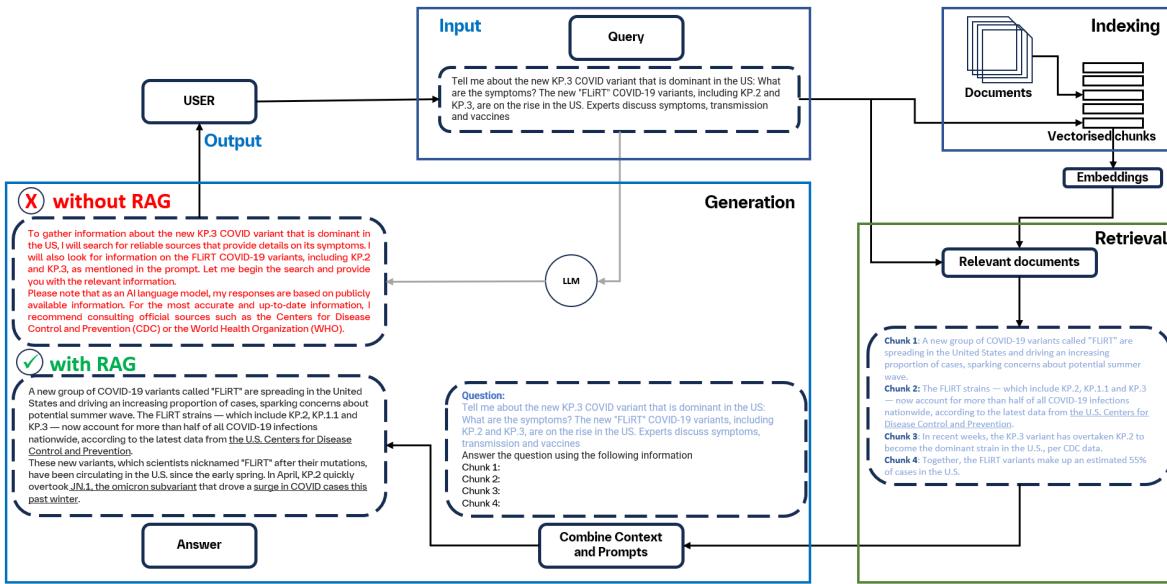
Although MultiMedQA is a useful benchmark, it does not capture the actual clinical scenarios and workflows followed by clinicians. To address this limitation, Zakka et al. [24] curated ClinicalQA as a benchmark dataset containing open-ended questions for different medical specialties, including treatment guidelines recommendations. The authors in [46] introduce MIRAGE, a benchmark for biomedical settings consisting of five commonly used datasets for medical QA. In particular, MMLU-Med, MedQA-US and MedMCQA are included to represent examination QA datasets, while PubMedQA and BioASQ-Y/N account for research QA settings. Note that all tasks in MIRAGE consist of multi-choice questions, and accuracy and its standard deviation are the default evaluation metrics.

The majority of the existing datasets have insufficient medical labels. For instance, most datasets only provide a single label, e.g. medical entities, which are not detailed enough to represent the condition and intention of a patient in the real world. Moreover, the existing annotated datasets are limited in scale, usually consisting of only a few hundred dialogues [36]. To address these issues, Chen et al. [36] proposed an extensive medical dialogue dataset with multi-level fine-grained annotations comprising five separate tasks which include named entity recognition (NER), dialogue act classification, symptom label inference, medical report generation, and diagnosis-oriented dialogue policy.

Most datasets released to date are curated in English, making it difficult to obtain a non-English dataset. However, recently Chinese datasets such MedDG [29] and MedDialog-CN [28] have been proposed

### 3.3. RAG Overview

RAG aims to ground the responses of LLMs to provide more factual and truthful responses and reduce hallucination. This is achieved by including new knowledge from external sources. As illustrated in the example of Figure 2, a user asks a question about a new COVID-19 variant ("Tell me about the new KP.3 COVID variant that is dominant in the US: What are the symptoms? The new "FLiRT" COVID-19 variants, including KP.2 and KP.3, are on the rise in the US. Experts discuss symptoms, transmission and vaccines."). An LLM such as ChatGPT will not be able to provide information on recent events because responses from LLMs are time-constrained by the data they are trained on (which is, in the best cases, a few months old). RAG helps LLMs overcome this limitation by retrieving information from up-to-date knowledge sources. In our case, the search algorithm will retrieve a collection of articles related to the prompted virus. Then, retrieved articles together with the prompt are used by the LLM model to generate an informed response.



**Figure 2.** A schematic illustration of the basic RAG workflow applied to answering a user question.

Adapted from [16].

As seen in Figure 2, RAG workflow comprises three important steps. The first step involves splitting documents into distinct segments, and vector indices are created using an encoder model (in the case of dense retrieval). This process is called indexing. After that, segments are located and retrieved based on their vector similarity with the query and indexed segments. The last step involves the model generating a response based on the context derived from the fetched segments and query. These steps constitute the core structure of the RAG process, reinforcing its capabilities in information retrieval and context-aware response generation. RAG leverages a synergistic approach that integrates information retrieval and in-context learning to enhance the performance of an LLM. With RAG, the LLM performance is contextually bolstered without performing computationally expensive model retraining or fine-tuning. This nuanced approach makes RAG practical and relatively easy to implement, making it a popular choice for building conversational AI tools, especially in critical sectors such as healthcare. A recent review outlined the progression of RAG technologies and their impact on tasks requiring extensive knowledge [16]. Three developmental approaches within the RAG framework were outlined: Naive, Advanced, and Modular RAG, each representing a successive improvement over the one before. In the context of this review, we discuss the RAG techniques used in healthcare by grouping them into these three categories.

### 3.4. Naive RAG

Naive RAG is one of the earliest techniques employed to ground LLMs to generate text using relevant information retrieved from external knowledge sources [16]. It is the simplest form of RAG without the sophistication needed to handle complex queries. The Naive RAG process follows the mentioned three steps: indexing, retrieval and generation. Implementations of Naive RAG in the medical domain are discussed below.

Ge et al. [47] used text embeddings to transform guidelines and guidance documents for liver diseases and conditions. Then, they converted user queries into embeddings in real-time using the text-embedding-ada-002 model and performed a search over a vector database to find matches for the embeddings before generating a response using GPT-3.5-turbo or GPT-4-32k. Their results show that they were able to generate more specific answers compared to general ChatGPT using GPT-3.5.

In another study, [48] presented ChatENT, a platform for question and answer over otolaryngology–head and neck surgery data. For the development of ChatENT, the authors curated a knowledge source from open source access and indexed it in a vector database. Then, they implemented a Naive

RAG architecture using GPT-4 as the LLM to generate responses. With ChatENT they demonstrated consistent responses and fewer hallucinations compared to ChatGPT4.0.

Zakka et al. [24] proposed Almanac, an LLM augmented with clinical knowledge from a vector database. They used a dataset containing 130 clinical cases curated by 5 experts with different specialties and certifications. The results showed that Almanac provided more accurate, complete, factual and safe responses to clinical questions. In another study, GPT-3.5 and GPT-4 models were compared to a custom retrieval-augmentation (RetA) LLM [49]. The results showed that both GPT-3.5 and GPT-4 generated more hallucinated responses than the RetA model in all 19 cases used for evaluation.

Thompson et al. [50] implemented a pipeline for zero-shot disease phenotyping over a collection of electronic health records (EHRs). Their method consisted in enriching the context of a PaLM 2-based LLM by retrieving text snippets from the patients' clinical records. The retrieval step was performed using regular expressions (regex) generated with the support of an expert physician. Then, a MapReduce technique was implemented to supply the final query to the LLM using only a selection of the complete snippet sets retrieved by the regex. The authors tested several prompting strategies and were able to obtain an improved performance in pulmonary hypertension phenotyping compared to the decision-rule approach devised by expert physicians (F1-score = 0.75 vs F1-score = 0.62, respectively).

The authors in [46] performed a systematic evaluation of naive RAG over a set of medical QA datasets from both examination and research areas. They compared the output of chain-of-thought prompting with and without external context from 5 different medical and general corpora. Inclusion of external knowledge increased GPT-4 average accuracy over multi-choice questions from 73.44% to 79.97%, whereas the average accuracy of GPT-3.5 and Mixtral were improved from 60.69% to 71.57% and from 61.42% to 69.48%, respectively. For questions in which related references can be found in PubMed, RAG strongly improved the performance of LLaMA2-70B (from 42.20% to 50.40%), leading to an accuracy that is close to the corresponding fine-tuned model over the medical domain. In addition, it was shown that by combining different retrievers using rank fusion algorithms leads to improved performances across various medical QA tasks. The authors investigated the impact of the number of retrieved chunks on the accuracy of the responses. They found that for accuracy of response, the optimal number of retrieved chunks varies with the QA dataset: for PubMed-based QA datasets, fewer chunks provided better accuracy, while for examination-based QA datasets, more retrieved chunks the LLM provided a better response.

Guo et al. [42] presented work in which text-to-summary generation is performed. They evaluated the ability of language models to generate lay summaries from scientific abstracts. The authors proposed a series of custom-made language models based on the BART architecture, enriching the text generation with retrieved information from external knowledge sources. They made use of both dense retrieval and definition-based retrieval, with corpora including PubMed and Wikipedia. Finally, the authors compared the output of the custom models to general pre-trained LLMs such as GPT-4 and LLaMA-2. They showed a trade-off between the integration of external information and understandability of the generated output, and validated their observations with both automated and human-based evaluation.

In general, the reviewed literature shows that Naive RAG systems perform better than foundational models alone. Different authors emphasize the importance to tune parameters such as the number of chunks and the retrieval mode to improve the model outcomes. Also, it is found that some Naive RAG approaches may be limited by low retrieval quality, redundancy, and generation errors. To overcome these challenges, advanced RAG methods have been proposed, which use techniques such as chunk optimisation, metadata integration, indexing structure, and context fusion. A review of the latter approaches is provided in the next section.

**Table 4.** A detailed list of different RAG methods used in the surveyed studies.

Authors	LLMs	Embedding	Pre-retrieval	Post-retrieval	Adv. Meth.	Outcomes
Al Ghadban et al. [19]	GPT-4	NA	Chunking	SS MMR	One-shot learning.	Acc: #141 (79%). Adeq: #49 (35%) Promising role RAG and LLMs in medical education.
Chen et al. [21]	GPT-4	NA	OIS	SS	Zero-shot CoT	Improved accuracy with RAG compared foundational models.
Chen et al. [22]	LLaMA-2-7B-chat	all-MiniLM-L6-v2	Chunking	SS	FT	FT+RAG provided best performance.
Gao et al. [51]	T5[52] GPT-3.5-T		OIS KGs	Path ranker. Aggregated	Zero-shot with path prompts	Improved diagnosis performance.
Ge et al. [47]	GPT-3.5-T GPT-4	ada-002	EDG	SS	Prompting strategy	7/10 completely correct with RAG.
Guo et al. [42]	LLaMA-2 GPT-4	BERT	Alignment optimisation	RR Summary	RALL	Improved generation performance and interpretability.
Jeong et al. [23]	Self-BioRAG		Chunking	RR	Critic & LLMS Generator	7.2% absolute improvement over SOTA with 7B or less.
Jiang et al. [53]	GPT-3.5 Baichuan13B-chat	BGE[54]	OIS KGs	RR	Query expansion CoK Noise filtering	Superior performance with RAG. F1: 4.62% better than baseline.
Jin et al. [55]	GPT-3.5-T GPT-4	ada-002	Chunking	SS	Integration with XGBoost	F1: 0.762. Acc: 83.3%. RAG surpasses the performance traditional methods.
Kang et al. [56]	GPT-3.5-T GPT-4	None	Create TOC	Truncation Summary	Use LLM for retrieval & generation	Improved retrieval capabilities. Score: 5.5.
Ke et al. [57]	GPT-4 GPT-3.5-T LLaMA-2 7B LLaMA-2 13B	ada-002	EDG Chunking	SS		Acc: 91.4%. RAG the performance comparable to human evaluators. Faster decision-making.
Li et al. [58]	LLaMA-7B	NA	EDG Chunking	RR	FT + wikipedia retrieval	RAG improved acc & efficiency F1 score: 0.84. Reduced workload.

**Table 4 Continued:** A detailed list of different RAG methods used in the surveyed studies.

Long et al. [48]	GPT-4	ada-002	Chunking	SS	Knowledge specific database	Improved performance with RAG over base models.
Lozano et al. [32]	GPT-3.5-T GPT-4	NA	Relevance Classifier	Synthesis Summary	Online search	Improvement with RAG over base models.
Manathunga & Illangasekara[59]	GPT-3.5-T	ada-002	Chunking	SS Summary	Embedding visualisation	Benefits of RAG in retrieving information quickly from large knowledge bases.
Markey et al. [60]	GPT-4	NA	Chunking	SS	Online search.	Potential for GenAI-powered clinical writing.
Miao et al. [61]	GPT-4				FT + CoT	RAG provides specialised & accurate medical advice for nephrology practices
Murugan et al. [62]	GPT-4	ada-002	MR	SS MMR	Prompt engineering Guardrails	Improvements with RAG.
Neupane et al. [63]	GPT-3.5-T Mistral-7B -Instruct	ada-002	Structured context. Chunking Manual indexing. Auto-merging retrieval.	Contextual compression	Online search	Efficacy in generating relevant responses. GPT-3.5-T: 0.93 & Mistral-7B: 0.92.
Ong et al. [64]	GPT-4 Gemini Pro 1.0 Med-PaLM-2	ada-002 bgeSEnv1.5[65]		SS	LLM vs "copilot"	RAG-LLM outperformed LLM alone.
Parmanto et al. [66]	LLaMA-2-7B Falcon-7B GPT 3.5-T	BGE[54]	Chunking	SS	FT	RAG + FT best results.
Quidwai & Lagana[67]	Mistral-7B -Instruct	bgeSEnv1.5	Indexing Chunking	SS	Pubmed dataset curation	Improved accuracy over base models.
Ranjit et al. [68]	davinci-003 GPT-3.5-T GPT-4	ALBEF[69]		Compression	Coupling to vision model	RAG achieved better outcomes BERTScore: 0.2865. Semb: 0.4026.
Rau et al. [70]	GPT-3.5-T GPT-4	ada-002	Chunking	SS	Visual interface	Superior performance with RAG. Time & cost savings.

**Table 4 Continued:** A detailed list of different RAG methods used in the surveyed studies.

Russe et al. [71]	GPT 3.5-T GPT-4	ada-002	Chunking	SS	Prompting strategy	Acc: GPT 3.5-T: 57% GPT-4: 83%.
Shi et al. [72]	GPT-3.5-T	MP-Net	Chunking MR	SS	ReAct architecture	Improved performance with RAG over baseline models.
Soman et al. [73]	LLaMA-2-13B GPT-3.5-T GPT-4 Prometheus	PubMedBert MiniLM	KGs	Similarity Context pruning.		KG-RAG enhanced performance.
Soong et al. [49]	GPT-3.5-T GPT-4	ada-002	Chunking	Summary		Improved performance with RAG over baseline models.
Thompson et al. [50]	Bison-001	NA	Token splitter. Regex Adding metadata Chunking	Map Reduce	Regex + LLM aggregation	RAG-LLM outperformed rule-based method. F1: 0.75.
Unlu et al. [74]	GPT-4	ada-002		SS	Iterative prompting	Potential to improve efficiency and reduce costs. Acc: 92.7%.
Vaid et al. [75]	GPT-3.5 GPT-4 Gemini Pro LLaMA-2-70B Mixtral-8x7B LLaMA-2-13B				CoT prompting. Agents	RAG with GPT-4 best performance.
Wang et al. [76]	GPT-3.5-T GPT-4		QO HR	Knowledge self-refiner	LLM-aided pre- and post-retrieval	RAG outperform baseline models.
Wang et al. [77]	LLaMA-2-7B LLaMA-2-13B GPT 3.5-T	ColBERT	QO	RR	JMLR	Demonstrate potential of joint IR & LLM training.
Wornow et al. [78]	GPT-4 LLaMA-2-70B Mixtral-8x7B LLaMA-2-7B	MiniLM BGE	Chunking	SS	Compared zero-shot and retrieval.	RAG with GPT-4 beats SOTA in zero-shot.
Yu, Guo and Sano [79]	LLaMA-2-7B GPT-3.5	ada-002	Chunking	SS	Feature extraction from ECG	RAG outperform few-shot approach.
Zakka et al. [24]	text-davinci-003	ada-002	Chunking	Similarity threshold	Adversarial prompting	RAG-LLM outperform ChatGPT.

**Table 4 Continued:** A detailed list of different RAG methods used in the surveyed studies.

Ziletti and D'Ambrosi[80]	GPT-3.5-T	bgeSEnv1.5	Entity masking	SS EN	Text-to-SQL	GPT-4-T best accuracy and executability.
	GPT-4-T					
	Gemini Pro 1.0					
	Claude 2.1					
	Mixtral-8x7B					
	Mixtral-Medium					

<sup>1</sup> \* Accuracy (Acc); Advanced Methodologies (Adv. Meth.); OpenAI's text-embedding-ada-002 model (ada-002); ALign the image and text representations BEfore Fusing (ALBEF); Accuracy (ACC); BAAI general embedding (bge); bge-small-en-v1.5 (bgeSEnv1.5); Chain-of-Knowledge (CoK); Chain-of-Thought (CoT); Text-davinci-003 (davinci-003); Desnse X Retrieval (DXR), Dense Passage Retriever (DPR); Enhancing data granularity (EDG); Entity normalisation (EN); Fine-tuning (FT); Frequency (Freq); Joint Medical LLM and Retrieval Training (JMLR); In-context learning (ICL), Image-Text Contrastive learning (ITC); Hybrid Retrieval (HR); Hypothesis Knowledge Graph Enhanced (HyKGE), Knowledge graphs (KGs); Maximal Marginal Relevance (MMR); Mixed Retrieval (MR); N/A (Not Available); Retrieval-Augmented Language Modelling (RALM); Retrieval Augmented Generation (RAG); Regular expression (Regex); Re-ranking (RR); Retrieval-Augmented Lay Language (RALL); Optimising Index Structure (OIS); Query Optimisation (QO); Similarity Search (SS); State-of-the-art (SOTA); Table of Contents (TOC); Turbo (T).

### 3.5. Advanced RAG

Advanced RAG employs specialized pre-retrieval and post-retrieval mechanisms with the aim to address challenges faced by Naive RAG such as failure to retrieve all relevant information, problem of integrating context from chunks retrieved and generating answers using irrelevant context. This is achieved by employing different strategies, including pre-retrieval, retrieval and post-retrieval. Pre-retrieval methods aim to optimize data indexing and can be achieved using different strategies such as improving data granularity, optimizing indexing structures, adding metadata, optimizing alignment and using mixed retrieval [16].

During the retrieval phase, dense retrieval-based RAG systems use embedding models to identify suitable context by computing the similarity between the prompt and the chunks. The retrieval step can be optimized by filtering the retrieved chunks using a threshold on their similarity with the user query. For example, Quidwai and Lagana [67] developed a system with a sophisticated retrieval component which allowed for the use of a predefined threshold to determine when insufficient relevant information were retrieved, after which the system can respond "Sorry, I could not find relevant information to complete your request." In this way, they reduced the generation of misleading or false information.

Presenting all retrieved content can introduce noise, shift attention from important content and may exceed the context window limit (number of tokens or words) of the LLM. To overcome the context window limits of LLMs and focus on crucial information, the post-retrieval process combines the query with the relevant context from knowledge sources before feeding it to the LLM. An important step is to reorder the retrieved information so that the most relevant content is closer to the prompt. This idea has been applied in frameworks like LlamaIndex [81] and LangChain [82].

Rau et al. [70] used LlamaIndex and GPT-3.5 to create a context aware chatbot grounded in a specialised knowledge base containing vectorised American College of Radiology (ACR) appropriateness criteria documents and compared its performance to GPT-3.5 and GPT-4. Out of 50 case files, they found that their context-based accGPT gave the most accurate and consistent advice that matched the ACR criteria for "usually appropriate" imaging decisions, in contrast to generic chatbots and radiology experts. Similarly, Russe et al. [71] used LlamaIndex as an interface between the external knowledge and a context-aware LLM (FraCChat). They extracted text information from radiology documents using the GPTVectorIndex function, which divides the content into smaller chunks (up to 512 tokens), converting chunks into data nodes. The data nodes were then encoded and stored as a dictionary-like data structure and then used in the answer creation. Using 100 radiology cases, FraCChat performed better on classification of fractures compared to generic chatbots, achieved 57% and 83% correct full "Arbeitsgemeinschaft Osteosynthesefragen" codes with GPT-3.5-Turbo and GPT-4 respectively.

A more recent study has shown that "noise" (documents not directly relevant to the query) can impact the performance of RAG systems - some models such LLaMA-2 and Phi-2 perform better when irrelevant documents are positioned far from the query [83]. Yang et al. [84] explored techniques for dealing with noise, namely: Bio-Epidemiology-NER, direct and indirect extraction to identify medical terminology augment an LLM using UMLS knowledge base. Using GPT-3.5 augmented with UMLS information, they created a trustable and explainable medical chatbot supported by factual knowledge. However, the extracted terminologies are not always related to the question asked, producing incomplete answers. Wang, Ma and Chen [76] used a large dataset containing high quality medical textbooks as an external knowledge source, combined with multiple retrievers to improve LLM performance in generating high quality content.

To improve retrieval and reasoning, researchers are investigating incorporating commonsense knowledge graphs (KGs) with dialogue systems in the medical domain. Two conversational models, MedKgConv[85] and MED[86] for medical dialogue generation were proposed, utilising multi-head attention and knowledge-aware neural conversation respectively. MedKgConv uses a BioBERT encoder to encode conversation history, which is then processed through Quick-UMLS to extract knowledge graphs for reasoning, and a BioBERT decoder for response generation. MED, on the other hand,

encodes augmented KGs alongside patient conversations using LLMs, enhancing it through medical entity annotation in a semi-supervised manner. Both models demonstrated improved performance on MedDialog and Covid datasets, with MedKgConv showing an increased 3.3% in F1 score and 6.3% in BLEU-4, and MED outperforming BioBERT by 8.5 points on F1 and 9.3 points on BLEU-4. These results underscore the effectiveness of integrating dialogue and graph-based knowledge in generating medical dialogues. The aforementioned KG approaches are implemented for QA tasks, where KGs contain structured information used as context for predicting the answer. As such, they could have limited versatility.

Soman et al. [73] proposed a context-aware prompt framework that adeptly retrieves biomedical context from the Scalable Precision Medicine Open Knowledge Engine (SPOKE) [87]. Using their KG-RAG approach, the performance of LLaMA-2 was significantly improved, showing a 71% improvement from the baseline.

In the realm of clinical development, advanced RAG techniques have been explored for clinical trial patient matching [78,88] and to accurately identify and report on inclusion and exclusion criteria for a clinical trial [74]. For instance, Jin et al. [88] employed aggregated ranking to perform patient-trial-matching using clinical notes. Similarly, a recent study used two retrieval pipelines, first selecting the top-k most pertinent segments from the patients' notes, and then using top-k segments as a prompt input to assess the LLM [78]. Their results demonstrated that it is possible to reduce processes that typically takes an hour-per-patient to a matter of seconds.

Another study used GPT-4 enabled with clinical notes through RAG for clinical trial screening [74]. The authors used LangChain's recursive chunking to divide patient notes into segments to preserve the context. To optimise the retrieval, they employed Facebook's AI Similarity Search (FAISS) [89]. They showed that using GPT-4 with RAG to screen patients for clinical trials can improve efficiency and reduce costs.

In summary, Advanced RAG systems are shown to improve over using only foundational models or Naive RAG approaches. The reviewed literature highlights a trade-off between complexity and performance gains, with more complex RAG implementations providing better outcomes at the cost of more application blocks to implement and maintain. In addition to that, some approaches stress the importance of ranking of the retrieved chunks, or the use of multiple retrievers to improve extraction of the required information from the knowledge source. With respect to the latter component of RAG, some authors employ particular implementations of the knowledge source (e.g., knowledge graphs). The improvements in response generation compared to more standard approaches depend on the nature of the information included in the knowledge source and the complexity of the user queries.

### 3.6. Modular RAG

Modular RAG incorporates several techniques and modules from advanced RAG, allowing more customisation and optimization of the RAG system, as well as integration of methods to improve different functions [16]. For example, modular RAG can include a search module for similarity retrieval, which can improve the quality and diversity of the retrieved content, and apply a fine-tuning approach in the retriever, which can adapt the retriever to the specific domain or task [90].

Wang, Ma and Chen [76] presented an LLM augmented by medical knowledge, using modules consisting of hybrid retrievers, query augmentation and an LLM reader. Each module enhanced a particular task, e.g., the query augmentation module improved prompts for effective medical information retrieval and the LLM reader module provided medical context to the question. They reported an improvement in response accuracy ranging between 11.4% to 13.2% on open-medical QA tasks, as compared to the GPT-4-Turbo without RAG. Despite using a smaller dataset, they showed that using medical textbooks as a knowledge source outperformed Wikipedia in the medical domain. This highlights the importance of context and quality information in specialised domains such as healthcare.

A recent study by Jin et al. [55] proposed a framework that integrates advanced techniques such as large-scale feature extraction combined with RAG, accurate scoring of features from medical knowledge using LlamaIndex, and XGBoost [91] algorithm for prediction. By employing this modular approach, they showed improved prediction of potential diseases, surpassing the capabilities of GPT-3.5, GPT-4, and fine tuning LLaMA-2.

In another study, a new approach to perform RAG is presented, eliminating the reliance on vector embedding by employing direct and flexible retrieval using natural language prompts [56]. The authors used an LLM to handle the step of document retrieval to response generation without needing a vector database and indexing, simplifying the RAG process. They showed the performance of prompt-RAG through a QA GPT-based chatbot using Korean medicine documents. They showed that the novel prompt-RAG achieved good results, outperforming ChatGPT and RAG-based models which used traditional vector embedding. Specifically, based on ratings from three doctors, prompt-RAG scored better in relevance and informativeness, similarly on readability. On the downside, the response time was significantly slower.

When user queries contain limited context information, the retriever may be unable to retrieve relevant documents from the knowledge sources. A method that uses hypothetical outputs generated from user queries has been proposed [18], improving performance in zero-shot scenarios. Compared to knowledge stored in unstructured documents such as in the portable document format (PDF), KGs are more ideal for RAG because of the easiness in accessing relations among knowledge items [92]. Ongoing explorations are focused on designing the best strategy to extract information from KGs and to facilitate interaction between LLMs and KGs. For instance, [53] presented a Hypothesis Knowledge Graph Enhanced (HyKGE) framework to improve the generation of responses from LLMs. The HyKGE comprises 4 modules, namely: hypothesis output, Named Entity Recognition (NER), KG retrieval, and noise knowledge filtering. Experimenting on two medical QA tasks, the authors demonstrated a 4.62% improvement compared to baseline in F1 score with their modular HyKGE framework. HyKGE also was able to address challenges such as poor accuracy and interpretability, and showcased potential application in the field of medicine.

As a general trend, researchers are now exploring the implications of having multiple LLMs jointly working together. For instance, Lozano et al. [32] proposed a RetA LLM system called Clinfo.ai consisting of 4 LLMs, jointly forming an LLM chain. They employed the first LLM to perform an index search on either PubMed or Semantic Scholars, an LLM for relevance classification, an LLM for article summarisation using a user query and the fourth LLM using task-specification prompts to guide the LLM output. Another study proposed a RAG-based method employing multiple LLM agents for feature extraction, prompt preparation and augmented model inference [79]. The authors evaluated their approach using two datasets for diagnosing arrhythmia and sleep apnea. The findings suggest that their zero-shot strategy not only outperforms previous methods that use a few-shot LLM, but also comparable to supervised techniques trained on large datasets. Finally, agents are increasingly being used to improve capabilities of LLMs in solving complex problems by sharing tasks across multiple agents [93]. Multi-agents systems have the potential to improve the capabilities of LLMs in solving complex problems in the medical domain, by splitting tasks into multiple subtasks and assigning specific LLM agents to each of them.

In summary, approaches that implement the Modular RAG framework are characterized by more complex pre- and post-retrieval steps, showcasing reformulation of the user query or LLM readers. From the trend in the more recent literature, we envision a wider and wider adoption of agentic frameworks in RAG systems. LLM-based agents [94] allow for a redefinition of complex tasks into simpler ones, with dedicated "reasoning engines" to tackle them. This has the advantage of casting the original problem into a modular approach, with the related pros and cons in terms of managing maintainance and complexity.

To conclude this section, we provide a summary of the reviewed literature, encompassing the different RAG implementations, in Table 4. Overall, the reviewed implementations of RAG in healthcare

follow the three paradigms of RAG, namely Naive, Advanced, and Modular RAG, as described in [16]. Across each RAG implementation, researchers have proposed different solutions to improve the retrieval and generation steps. In the next section, we will introduce the common evaluation criteria that are being adopted to assess performance of RAG-based applications.

### 3.7. Evaluation Metrics and Frameworks

In a RAG system, there are three important aspects: 1) the user prompt (that is, the user query), 2) the retrieval of the context related to the prompt, and 3) using the retrieved context and the user prompt to produce output. Because there are usually no reference answers to questions (i.e., *ground truth*), evaluations have been focused on quality aspects, separately evaluating the retriever and generator. Commonly used evaluation metrics for RAG-based applications include: 1) **Accuracy/Correctness** - e.g., by comparing generated text to ground truth or using human evaluators to assess the correctness of LLM outputs; 2) **Completeness** - the proportion of all retrieved context that is relevant to the user query; 3) **Faithfulness/Consistency** - the degree to which the LLM's output is grounded in the provided context and factually consistent (i.e., not an hallucination) [95]; 4) **Relevance** - whether the generated answer addresses the question asked by the user (*answer relevance*), and whether the retrieved context is relevant to the query (*context relevance*); 5) **Fluency** - the ability of the system to generate natural and easy-to-read text. The evaluation metrics employed in the surveyed studies are presented in Table 5.

To date, there is no harmonized evaluation approach, with different frameworks using different metrics. For example, Retrieval Augmented Generation Assessment (RAGAs) [95] is one of the most commonly used frameworks to evaluate RAG-based systems. It evaluates a RAG pipeline using four aspects: faithfulness, answer relevance, context relevance and context recall, combining these four aspects to generate a single score to measure the performance of the system. Other packages, such as continuous-eval, have proposed a combination of deterministic, semantic and LLM-based approaches to evaluate RAG pipelines [96]. In addition, TrueLens framework also provides metrics for objective evaluation of RAG pipelines [97]. Another evaluation framework is DeepEval [98], an open-source framework for evaluating LLMs, including RAG applications.

The previously described Clinfo.ai from Lazano et al. [32] introduced a novel dataset, PubMedRS-200, which contains question-answer pairs derived from systematic reviews. This dataset allows for automatic assessment of LLM performance in a RAG QA system. Their framework and benchmark dataset are openly accessible to promote reproducibility. In another study, a toolkit for evaluation of Naive RAGs was proposed [46]. Briefly, the authors considered different LLMs (both from general and biomedical domains), supplemented with five diverse corpora. They compared the performance on medical QA tasks using RAG over four different retrievers. As the tasks in the QA set required the choice between a given set of answers, they chose to use accuracy as the evaluation metric.

Guo et al. [42] presented an evaluation of LLM-generated lay summaries that comprises both automatic and human-based metrics. For evaluating the quality of the synthetic text generation to the target text, they employed ROUGE-L, BERTScore, BLUE and METEOR metrics. Notably, this list of metrics includes both lexicon-based metrics and embeddings similarity. Also, the authors used the Coleman-Liau readability score and word familiarity to assess the easiness of text readability. In addition, the authors trained a RoBERTa model [99] to generate a "Plainness Score", a metric that indicates how much a generated text is representative of the target one. In their effort to holistically evaluate the output of the language models, the authors also employed human reviewers for evaluating characteristics such as grammar, understandability, meaning preservation, correctness of key information and relevance of external information in the generated texts. The human evaluators gave their evaluation on a Likert scale.

Despite the progress achieved in evaluating RAG systems, two main limitations still remain: reliance on human evaluations, which hinders scalability and comparability, and the focus on output quality without considering the relevance of the information used. The challenge of selecting relevant

sources for summarisation is as significant as the summarisation itself. Therefore, a benchmark is needed for comprehensive evaluation of document selection and summary generation capabilities.

In summary, the evaluation landscape of RAG-based applications is still in an early stage, with several library-specific metrics put forward by different teams. As a general remark, there is a need to translate from pure word/gram-based metrics to more complex ones that can capture the generative nature of LLM outputs. Also, there is a suite of benchmarks and metrics developed for information retrieval that should be adapted and refined for the retrieval part of RAG systems.

### 3.8. Ethical Considerations

Healthcare is one of the most regulated industries, guided by principles such as bioethics [100] and various data regulations. For LLMs to be accepted and adopted in healthcare, ethical considerations surrounding their use such as bias, safety and hallucination need to be addressed.

Recently, RAG has shown the ability to boost the performance of LLMs by grounding LLM's responses using retrieved knowledge from external sources. While the benefits of RAG are numerous, its practical application in the medical domain also underscores ethical considerations given the critical nature of medical decision-making. Though RAG can help to constrain LLMs from generating incorrect output, a recent study has shown that even with RAG, LLMs may generate incorrect answers and explanations [47]. Any system that might suggest an incorrect treatment plan or diagnosis could have disastrous consequences for the patient. Therefore, the correctness of an LLM-based system is an ethical concern. As seen in Table 5, all studies included in this review evaluated the accuracy of their RAG pipeline, including outputs generated by the model.

Data leakage is another known problem for LLMs that has been extensively studied [101–103]. Researchers have proposed RAG as a safer approach to reduce LLMs' tendency to output memorised data from its training. However, as argued in [104], information from pre-training/fine-tuning datasets (from the LLM) and the retrieval dataset can be potentially exposed when using RAG. In the context of healthcare, a retrieval dataset may contain sensitive patient information such as medications, diagnoses, and personally identifiable information (PII). As such RAG-based LLM system need to ensure and evaluate safety and privacy. The majority of the studies we reviewed do not address privacy issues with the exception of [58] and [74]. Li et al. [58] reduced privacy risks by removing identifiable information about the patient or doctor from their retrieval dataset. Another study used a multi-layered approach to ensure data privacy and security when using Azure OpenAI for sensitive healthcare data [74]. Another important concern that cannot be overlooked when LLMs are applied in the healthcare is safety. Only three out of the thirty-seven studies we reviewed evaluated their models against intentional or unintentional harms. In one study, adversarial prompting [105] was used to evaluate the robustness and reliability of system outputs in unexpected situations [24]. Another study used preoperative guidelines to guide decision-making of LLMs, thus, improve patients safety [57]. Similarly, a technique that flags safety concerns was developed, demonstrating zero instances of alarming red flags during testing [48].

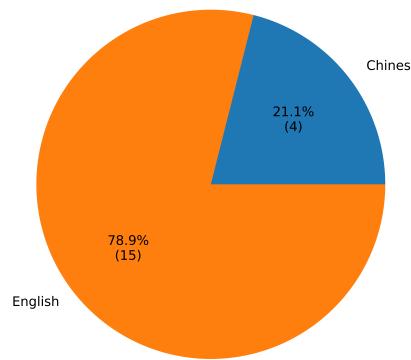
**Table 5.** Metrics used to evaluate RAG-based systems in the medical domain. Additionally, we assess whether ethical principles: privacy, safety, robustness (robust), bias, trust (explainability/interpretability) have been considered.

Author	Accuracy (Correctness)	Complete	Faithfulness (Consistency)	Fluency	Relevance	Ethical Issues							
						Verify Source	Privacy	Safety	Robust	Bias	Trust	Eval	
Al Ghadban et al. [19]	✓	✓	✓	✗	✓	✓	✗	✗	✗	✓	✗	✗	Auto
Chen et al. [21]	✓	✓	✓	✓	✓	✓	✗	✓	✗	✓	✓	✓	Man
Chen et al. [22]	✓	✓	✓	✓	✓	✓	✗	✓	✗	✓	✓	✓	Man
Ge et al. [47]	✓	✗	✓	✗	✓	✓	✗	✗	✗	✗	✗	✗	Man
Guo et al. [42]	✓	✓	✗	✗	✓	✓	✓	✗	✗	✗	✗	✗	Both
Jin et al. [88]	✓	✗	✓	✗	✓	✓	✗	✗	✗	✗	✗	✓	Man
Jin et al. [55]	✓	✗	✗	✗	✓	✓	✗	✗	✗	✗	✗	✗	Auto
Kang et al. [56]	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	Man
Ke et al. [57]	✓	✗	✓	✗	✓	✓	✗	✗	✓	✗	✗	✗	Man
Li et al. [58]	✓	✓	✓	✗	✗	✓	✓	✓	✗	✗	✗	✗	Auto
Long et al. [48]	✓	✗	✓	✓	✗	✓	✓	✓	✗	✗	✗	✗	Man
Lozano et al. [32]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	Auto
Muragan et al. [62]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	Man
Neupane et al. [63]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	Auto
Parmanto et al. [66]	✓	✓	✓	✗	✗	✓	✓	✗	✗	✗	✗	✗	Both
Ranjit et al. [68]	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	Auto
Rau et al. [70]	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	Auto
Russe et al. [71]	✓	✓	✓	✗	✓	✓	✗	✗	✗	✗	✗	✗	Auto
Shi et al. [72]	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✓	Man
Soong et al. [49]	✓	✓	✓	✗	✗	✓	✓	✓	✗	✗	✗	✗	Man
Thompson et al. [50]	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	Auto
Unlu et al. [74]	✓	✓	✗	✓	✗	✗	✗	✗	✓	✗	✗	✗	Auto
Vaid et al. [75]	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	Man
Wang et al. [76]	✓	✓	✓	✓	✗	✗	✓	✓	✗	✗	✗	✗	Auto
Wornow et al. [78]	✓	✓	✓	✓	✗	✗	✗	✓	✗	✗	✗	✗	Auto
Yu, Guo and Sana [79]	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	Auto
Zekka et al. [24]	✓	✓	✗	✗	✗	✗	✓	✓	✗	✓	✗	✗	Man

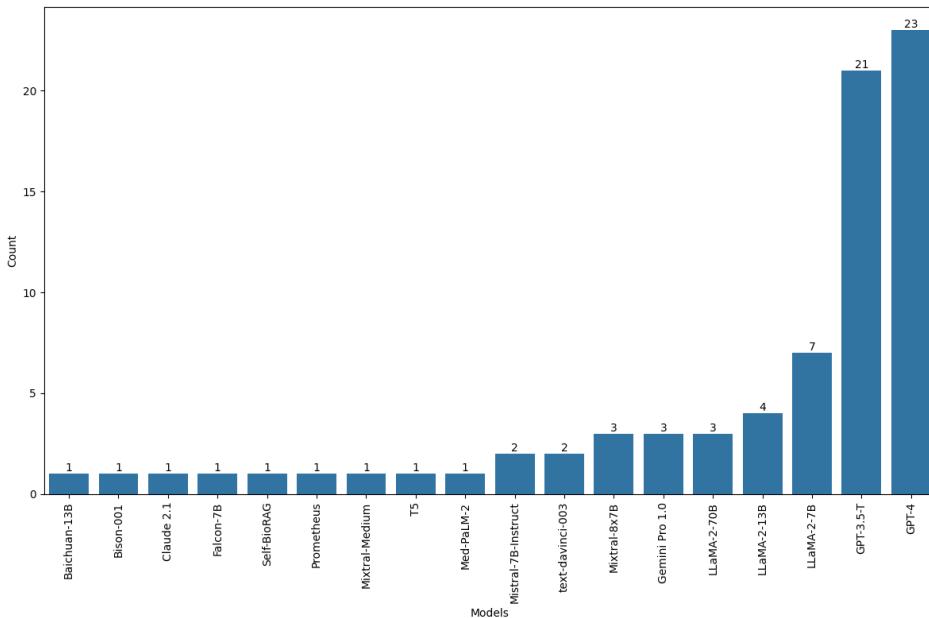
\* Evaluation (Eval) indicates the type of evaluation used: automated (Auto), manual (Man) or a combination of both approaches.

Bias is another concern that is extensively studied in AI due to its ability to amplify inequity. Surprisingly, most of the papers we examined do not assess the existence of bias in the responses generated by their RAG systems. In the studies we reviewed that evaluated bias, Chen et al. [21] utilised a physician to determine the potential bias in the answers generated by DocQA, a RAG-enabled LLM. They observed reduced bias content with RAG compared to an LLM alone. Though RAG substantially reduces the LLM's ability to use memorised data from model training [72], it does not avert bias completely, which is baked into the LLMs underlying training data.

Ethical considerations such as bias, privacy, hallucination and safety are of paramount importance and should be addressed when working with RAG-based LLMs in healthcare. In a nutshell, these concerns can be addressed by implementing robust data privacy measures, promoting transparency and accountability, mitigating bias, emphasizing human oversight, and promoting human-autonomy. From our review, we see that only a few articles tackle these issues, and we see here wide margin of improvement.



(a) The languages of the datasets used.



(b) Names of LLMs.

**Figure 3.** a) Dataset Languages. b) LLMs explored for retrieval-augmented generation in healthcare. Please note that some studies used more than one model, hence the total count of models is higher than the number of studies included in this review

### 3.9. Data Analysis

Figure 3a presents the distribution of languages of common retrieval datasets used for RAG in medical domain. It is evident that the majority (78.9%) of the datasets are in English, except for four datasets which are in Chinese. As seen in Figure 3b, the majority of the studies that employed RAG in the healthcare setting made use of proprietary models such as OpenAI's GPT-3.5/4 models. However, the use of these models raises privacy issues, especially when dealing with sensitive information such as patient data. In such cases, open-source models deployed in a controlled environment may be a suitable solution. However, it is worth noting that open-source models generally have lower performance compared to proprietary models and have a more limited context window [106]. We

also observed that the majority of the studies we reviewed made use of OpenAI's text-embedding-ada-002 model [107] for generating embeddings. Other embedding models used included BAAI general embedding (BGE) [54,65], and HuggingFace's all-MiniLM-L6-v2 [22]. Other studies used custom embedding models, as in [23]. A study by Kang et al. [56] demonstrated the feasibility of embeddings-free RAG in the medical domain. This suggests that one may not always need to utilise vector embeddings for successful RAG implementation.

The majority of the studies we surveyed used dense passage retrievers (DPR) as a retrieval method in their RAG architecture. Few studies have used custom retrievers. For pre-retrieval, strategies employed are enhancing data granularity, adding metadata, optimising indexing, mixed retrieval and alignment optimisation. This involves methods such as chunking, knowledge graphs, creating Table-of-Contents and entity masking. Finally, studies have explored different modalities to improve the performance of their RAG architectures, ranging from one-shot learning, chain-of-thought and prompting to more advanced techniques such as using LLMs as agents

#### 4. Discussion

We comprehensively reviewed recent advancements of RAG-based approaches for LLMs in healthcare. Our survey discusses broader applications of RAG-based approaches within the landscape of LLMs for healthcare, dividing RAG approaches into three paradigms: naive, advanced and modular RAG. Moreover, we outline evaluation frameworks, including objectives and metrics used to assess the performance of RAG-based LLMs for healthcare.

Our findings indicate that proprietary LLMs such as GPT-3.5/4 are the most commonly used for RAG, employed in 30 out of the 37 studies we reviewed (see Table 4). The dominance of proprietary models is not surprising, given their superior performance in tasks like "zero-shot reasoning" when compared to open-source models. For example, GPT-4 and Claude 2 have consistently outperformed their open-source counterparts in these tasks [106], illustrating the strong capabilities and potential for off-the-shelf LLMs in solving complex problems in medicine. Another key observation from our review is the language bias in the datasets used as external knowledge sources. We found that the majority of these datasets are in English, with only a few exceptions in Chinese. This language bias presents a challenge in evaluating the performance of RAG on non-English datasets, other than Chinese. The lack of representative datasets highlight a gap in the current research and underscores the need for more diverse language datasets. Representation is crucial for fairness and equity in medical AI systems [108,109].

We found that most RAG-based studies are primarily focused on optimising retrieval and generation, including techniques such as incorporating metadata, re-ranking, chunking, summarisation, and adapting prompts. However, we believe that these optimisation techniques, while important, may only lead to marginal improvements in performance. They may also prove inadequate when handling complex medical queries, which often require reasoning over the evidence. An example of this is answering multi-hop queries where the system must retrieve and reason over several pieces of supporting evidence [110]. Furthermore, we observed that even within RAG, LLMs may not always respond as expected. This can be attributed to several factors: 1) the information contained in the pre-training of the model may leak in the final model answer; 2) irrelevant information may be retrieved, leading to inaccurate responses; 3) LLM generation can be unpredictable, resulting in unexpected outputs. Also, our findings indicate a lack of standardisation in the evaluation metrics used across different studies (as seen in Table 5). This makes it difficult to compare the performance of the RAG systems across different studies, highlighting a need for more uniform evaluation metrics. Recently, a standard benchmark collection for evaluating clinical language understanding tasks was proposed [111]. Lastly, we find that the majority of the studies we reviewed do not address safety concerns. Ensuring safety is essential to ensure that the system causes no harm, protect patients' privacy and comply with regulations.

We acknowledge that our study is not exempt from limitations. First, we only included papers published in English, which may have excluded relevant studies in other languages. Second, key words used in the search may have excluded some RAG-based studies in healthcare because RAG terminologies continues to rapidly evolve. Third, we relied on the information reported by the authors of the original papers, which may have introduced some bias or errors in our analysis. Last, despite conducting a systematic review we might have missed some relevant studies because we limited our search to studies published between January 2020 - March 2024.

Despite RAG advancements, retrieval quality remains a challenge. Several issues can arise from poor retrieval quality from external knowledge sources. For example, when not all the relevant chunks are retrieved (low context recall), it is challenging for the LLM to produce complete and coherent text. Also, the retrieved chunks might not align with the user query, potentially lead to hallucinations. To identify and mitigate hallucinations, different methods have been proposed. The authors in [5] proposed Med-Halt, a domain specific benchmark to evaluate and reduce hallucination in LLMs. Other researchers have proposed overcoming hallucination by using human-in-the-loop, algorithmic corrections and fine-tuning [112]. However, the aforementioned studies do not specifically focus on RAG scenarios. Wu et al. [113] proposed a high quality manually annotated dataset called RAGTruth and achieved comparative performance on existing prompt based techniques using SOTA LLMs, for e.g, GPT-4. Further research should focus on retrieval issues, by developing novel methods to find the most relevant information for the query. In addition, curation of diverse benchmark datasets for hallucination detection in healthcare, going beyond multi-choice questions and more inline with clinical practise, should constitute a primary research endeavour.

Current LLMs are limited by their context window. The context window determines the number of tokens the model can process and generate at a given user session. A right balance should be sought by the user, who should provide a sufficient context to the model without exceeding its context length [114]. Inadequate context can lead to a lack of necessary information, while excessive irrelevant context can impair the ability to recall relevant context. Ongoing research is exploring the benefits of longer context in enabling models to access more information from external knowledge sources [115]. This is especially important in healthcare settings, where clinicians often rely on longitudinal data, such as clinical notes, lab values, and imaging data. To incorporate multi-modal data effectively, LLMs need to consider much longer contexts. Future work should prioritise evaluating the impact of longer context on LLMs for healthcare. Though a longer context window can enhance the properties of RAG, it is not clear how this can reduce hallucination. A recent study has demonstrated that even with RAG, out of all the medical queries, approximately 45% of the responses provided by GPT-4 were not completely backed by the URLs retrieved [116]. Future studies should focus on techniques to robustly handle noise, information integration and improving the validity of sources, e.g., using post-hoc citation-enhanced generation [117].

The emergence of multimodal LLMs that can understand and output text, images, and audio presents exciting prospects for RAG. In healthcare, the gap in visual semantics and language understanding has been addressed by vision-language models that correlate visual semantics from medical images and text from medical reports or EHRs. For instance, Liu et al. [118] proposed contrastive language-image pre-training using zero-shot prompting to provide additional knowledge to help the model making explainable and accurate diagnosis from medical images. Beyond images, multimodal LLMs grounded in an individual's specific data to estimate disease risk have been proposed [20]. Other researchers have proposed RAG driven frameworks to improve multimodal EHRs representation [119].

Human evaluations remains crucial in assessing the output generated by RAG-based models. While human evaluation remains important, automated solutions are being developed to improve evaluation of LLMs by assessing both the output generated and information used to generate the answers. For instance, the RAGAs framework [95] allows to evaluate both generator and retriever separately. Automated benchmarks have also been proposed, integrating the evaluation of the ability

of the system to retrieve and summarise relevant information [32]. Future research should explore combining automated metrics and human evaluation, ensuring that there is alignment. Finally, there is a need to shift from generic evaluation frameworks and benchmarks to contextual standardised evaluation metrics for RAG-based LLMs in healthcare.

While RAG holds immense promise in healthcare, its adoption must be guided by ethical principles, balancing innovation with patient privacy and safety. For instance, when external databases are accessed in RAG systems, there is a risk of inadvertently revealing sensitive information such as patient prescription information [58]. One way to overcome this challenge is to ensure that retrieval databases do not contain personally identifiable patient information. Additionally, composite structured prompting can be used to effectively extract retrieval data and evaluate privacy leakages by comparing LLM-generated outputs with the retrieved information [104]. Future studies, should explore novel measures to effectively overcome retrieval information leakages. Furthermore, it is crucial for LLM developers in healthcare to proactively address ethical issues throughout the AI development life cycle [120]. This proactive approach would foster trust and encourage the adoption of RAG-based LLMs in critical sectors such as healthcare [121,122]. Finally, if developed ethically, LLMs in medicine can potentially increase access and equity in healthcare, for example enabling clinical trials to be more inclusive or providing treatments that are tailored to patients from diverse demographics.

## 5. Conclusions

In this paper, we comprehensively outline RAG's advancement in grounding and improving the capabilities of LLMs in the medical domain. First, we discuss the available datasets used for grounding LLMs for healthcare tasks such as question-answer/dialogue and information retrieval. Second, we compare the models, and the retrieval and augmentation techniques employed by existing studies. Third, we assess evaluation frameworks proposed for RAG systems in the medical domain. Our results show that there is a growing interest in applying RAG to ground LLMs in healthcare, and proprietary LLMs are the most commonly used models. When it comes to evaluation of RAG-pipeline, our findings highlight the absence of a standardised framework for assessing RAG pipelines in the medical field. Despite these challenges, RAG has the potential to ground and customise the domain knowledge of LLMs for healthcare, by integrating dynamic data, standards, and a complete integration with individual scenarios. Therefore, revolutionise various areas in healthcare from drug development to disease prediction and personalised care management. Nevertheless for RAG to be effectively implemented in healthcare, it is essential to adequately address challenges such as integration of information, handling of noise, source factuality and ethical considerations. Finally, with continuous improvements LLMs will play an instrumental role in shaping the future of healthcare, driving innovation and enhancing patient care.

**Author Contributions:** LMA and PM were responsible for conducting the systematic review and analysis. All authors contributed to the study design, reviewed and contributed to the manuscript. All authors have read and agreed to the published version of the manuscript.

**Funding:** This research is supported by Boehringer Ingelheim Pharma GmbH & Co. KG.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** Not applicable.

**Acknowledgments:** We would like to express our deepest gratitude to reviewers, especially Lars Andersen and Dr Nicole Kraemer for reviewing our work. Their insightful comments and feedback has greatly improved our manuscript.

**Conflicts of Interest:** LMA, PM, SD and JS are employed by the by Boehringer Ingelheim Pharma GmbH & Co. KG, SB is an employee of Boehringer Ingelheim (China) Investment Co., Ltd.

## Abbreviations

The following abbreviations are used in this manuscript:

BERT	Bidirectional Encoder Representations from Transformers
GPT	Generative Pre-trained Transformers
HyKGE	Hypothesis Knowledge Graph Enhanced
KGs	Knowledge Graphs
LLM	Large Language Models
RAG	Retrieval Augmented Generation
QA	Question answering
SOTA	State-of-the-art (SOTA)

## References

1. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D.M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; Amodei, D. Language models are few-shot learners. Proceedings of the 34th International Conference on Neural Information Processing Systems; Curran Associates Inc.: Red Hook, NY, USA, 2020; NIPS'20.
2. OpenAI. GPT-4 Technical Report, 2023, [\[arXiv:cs.CL/2303.08774\]](https://arxiv.org/abs/cs/2303.08774).
3. Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; Bikel, D.; Blecher, L.; Ferrer, C.C.; Chen, M.; Cucurull, G.; Esiobu, D.; Fernandes, J.; Fu, J.; Fu, W.; Fuller, B.; Gao, C.; Goswami, V.; Goyal, N.; Hartshorn, A.; Hosseini, S.; Hou, R.; Inan, H.; Kardas, M.; Kerkez, V.; Khabsa, M.; Kloumann, I.; Korenev, A.; Koura, P.S.; Lachaux, M.A.; Lavril, T.; Lee, J.; Liskovich, D.; Lu, Y.; Mao, Y.; Martinet, X.; Mihaylov, T.; Mishra, P.; Molybog, I.; Nie, Y.; Poulton, A.; Reizenstein, J.; Rungta, R.; Saladi, K.; Schelten, A.; Silva, R.; Smith, E.M.; Subramanian, R.; Tan, X.E.; Tang, B.; Taylor, R.; Williams, A.; Kuan, J.X.; Xu, P.; Yan, Z.; Zarov, I.; Zhang, Y.; Fan, A.; Kambadur, M.; Narang, S.; Rodriguez, A.; Stojnic, R.; Edunov, S.; Scialom, T. Llama 2: Open Foundation and Fine-Tuned Chat Models, 2023, [\[arXiv:cs.CL/2307.09288\]](https://arxiv.org/abs/cs/2307.09288).
4. Gemini Team. Gemini: A Family of Highly Capable Multimodal Models, 2023, [\[arXiv:cs.CL/2312.11805\]](https://arxiv.org/abs/cs/2312.11805).
5. Pal, A.; Umapathi, L.K.; Sankarasubbu, M. Med-HALT: Medical Domain Hallucination Test for Large Language Models, 2023, [\[arXiv:cs.CL/2307.15343\]](https://arxiv.org/abs/cs/2307.15343).
6. Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y.J.; Madotto, A.; Fung, P. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.* **2023**, *55*. doi:10.1145/3571730.
7. Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y.T.; Li, Y.; Lundberg, S.; Nori, H.; Palangi, H.; Ribeiro, M.T.; Zhang, Y. Sparks of Artificial General Intelligence: Early experiments with GPT-4, 2023, [\[arXiv:cs.CL/2303.12712\]](https://arxiv.org/abs/cs/2303.12712).
8. Lee, J.; Yoon, W.; Kim, S.; Kim, D.; Kim, S.; So, C.H.; Kang, J. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics* **2019**, *36*, 1234–1240. doi:10.1093/bioinformatics/btz682.
9. Alsentzer, E.; Murphy, J.R.; Boag, W.; Weng, W.H.; Jin, D.; Naumann, T.; McDermott, M.B.A. Publicly Available Clinical BERT Embeddings, 2019, [\[arXiv:cs.CL/1904.03323\]](https://arxiv.org/abs/cs/1904.03323).
10. Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, 2019, [\[arXiv:cs.CL/1810.04805\]](https://arxiv.org/abs/cs/1810.04805).
11. Singhal, K.; Azizi, S.; Tu, T.; Mahdavi, S.S.; Wei, J.; Chung, H.W.; Scales, N.; Tanwani, A.; Cole-Lewis, H.; Pfohl, S.; Payne, P.; Seneviratne, M.; Gamble, P.; Kelly, C.; Babiker, A.; Schärli, N.; Chowdhery, A.; Mansfield, P.; Demner-Fushman, D.; Agüera y Arcas, B.; Webster, D.; Corrado, G.S.; Matias, Y.; Chou, K.; Gottweis, J.; Tomasev, N.; Liu, Y.; Rajkomar, A.; Barral, J.; Semturs, C.; Karthikesalingam, A.; Natarajan, V. Large language models encode clinical knowledge. *Nature* **2023**, *620*, 172–180. doi:10.1038/s41586-023-06291-2.
12. Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttrler, H.; Lewis, M.; tau Yih, W.; Rocktäschel, T.; Riedel, S.; Kiela, D. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks, 2021, [\[arXiv:cs.CL/2005.11401\]](https://arxiv.org/abs/cs/2005.11401).

13. Dong, Q.; Li, L.; Dai, D.; Zheng, C.; Wu, Z.; Chang, B.; Sun, X.; Xu, J.; Li, L.; Sui, Z. A Survey on In-context Learning, 2023, [\[arXiv:cs.CL/2301.00234\]](https://arxiv.org/abs/cs/2301.00234).
14. Das, S.; Saha, S.; Srihari, R.K. Diving Deep into Modes of Fact Hallucinations in Dialogue Systems, 2023, [\[arXiv:cs.CL/2301.04449\]](https://arxiv.org/abs/cs/2301.04449).
15. Wang, C.; Ong, J.; Wang, C.; Ong, H.; Cheng, R.; Ong, D. Potential for GPT Technology to Optimize Future Clinical Decision-Making Using Retrieval-Augmented Generation. *Annals of Biomedical Engineering* **2023**. doi:10.1007/s10439-023-03327-6.
16. Gao, Y.; Xiong, Y.; Gao, X.; Jia, K.; Pan, J.; Bi, Y.; Dai, Y.; Sun, J.; Guo, Q.; Wang, M.; Wang, H. Retrieval-Augmented Generation for Large Language Models: A Survey, 2024, [\[arXiv:cs.CL/2312.10997\]](https://arxiv.org/abs/cs/2312.10997).
17. Zhao, P.; Zhang, H.; Yu, Q.; Wang, Z.; Geng, Y.; Fu, F.; Yang, L.; Zhang, W.; Cui, B. Retrieval-Augmented Generation for AI-Generated Content: A Survey, 2024, [\[arXiv:cs.CV/2402.19473\]](https://arxiv.org/abs/cs/2402.19473).
18. Gao, L.; Ma, X.; Lin, J.; Callan, J. Precise Zero-Shot Dense Retrieval without Relevance Labels, 2022, [\[arXiv:cs.IR/2212.10496\]](https://arxiv.org/abs/cs/IR/2212.10496).
19. Al Ghadban, Y.; Lu, H.Y.; Adavi, U.; Sharma, A.; Gara, S.; Das, N.; Kumar, B.; John, R.; Devarasetty, P.; Hirst, J. Transforming Healthcare Education: Harnessing Large Language Models for Frontline Health Worker Capacity Building using Retrieval-Augmented Generation, 2023. doi:10.1101/2023.12.15.23300009.
20. Belyaeva, A.; Cosentino, J.; Hormozdiari, F.; Eswaran, K.; Shetty, S.; Corrado, G.; Carroll, A.; McLean, C.Y.; Furlotte, N.A. Multimodal LLMs for Health Grounded in Individual-Specific Data. *Machine Learning for Multimodal Healthcare Data*; Maier, A.K.; Schnabel, J.A.; Tiwari, P.; Stegle, O., Eds.; Springer Nature Switzerland: Cham, 2024; pp. 86–102.
21. Chen, X.; You, M.; Wang, L.; Liu, W.; Fu, Y.; Xu, J.; Zhang, S.; Chen, G.; Li, K.; Li, J. Evaluating and Enhancing Large Language Models Performance in Domain-specific Medicine: Osteoarthritis Management with DocOA, 2024, [\[arXiv:cs.CL/2401.12998\]](https://arxiv.org/abs/cs/2401.12998).
22. Chen, X.; Zhao, Z.; Zhang, W.; Xu, P.; Gao, L.; Xu, M.; Wu, Y.; Li, Y.; Shi, D.; He, M. EyeGPT: Ophthalmic Assistant with Large Language Models, 2024, [\[arXiv:cs.CL/2403.00840\]](https://arxiv.org/abs/cs/2403.00840).
23. Jeong, M.; Sohn, J.; Sung, M.; Kang, J. Improving Medical Reasoning through Retrieval and Self-Reflection with Retrieval-Augmented Large Language Models, 2024, [\[arXiv:cs.CL/2401.15269\]](https://arxiv.org/abs/cs/2401.15269).
24. Zakka, C.; Shad, R.; Chaurasia, A.; Dalal, A.R.; Kim, J.L.; Moor, M.; Fong, R.; Phillips, C.; Alexander, K.; Ashley, E.; Boyd, J.; Boyd, K.; Hirsch, K.; Langlotz, C.; Lee, R.; Melia, J.; Nelson, J.; Sallam, K.; Tullis, S.; Vogelsong, M.A.; Cunningham, J.P.; Hiesinger, W. Almanac — Retrieval-Augmented Language Models for Clinical Medicine. *NEJM AI* **2024**, 1, A1oa2300068, [\[https://ai.nejm.org/doi/pdf/10.1056/A1oa2300068\]](https://ai.nejm.org/doi/pdf/10.1056/A1oa2300068). doi:10.1056/A1oa2300068.
25. Liberati, A.; Altman, D.G.; Tetzlaff, J.; Mulrow, C.; Gøtzsche, P.C.; Ioannidis, J.P.A.; Clarke, M.; Devereaux, P.J.; Kleijnen, J.; Moher, D. The PRISMA Statement for Reporting Systematic Reviews and Meta-Analyses of Studies That Evaluate Health Care Interventions: Explanation and Elaboration. *PLOS Medicine* **2009**, 6, 1–28. doi:10.1371/journal.pmed.1000100.
26. Stefana, E.; Marciano, F.; Cocca, P.; Alberti, M. Predictive models to assess Oxygen Deficiency Hazard (ODH): A systematic review. *Safety Science* **2015**, 75, 1–14. doi:<https://doi.org/10.1016/j.ssci.2015.01.008>.
27. Hambarde, K.A.; Proen  a, H. Information Retrieval: Recent Advances and Beyond. *IEEE Access* **2023**, 11, 76581–76604. doi:10.1109/access.2023.3295776.
28. Chen, S.; Ju, Z.; Dong, X.; Fang, H.; Wang, S.; Yang, Y.; Zeng, J.; Zhang, R.; Zhang, R.; Zhou, M.; Zhu, P.; Xie, P. MedDialog: a large-scale medical dialogue dataset. *arXiv preprint arXiv:2004.03329* **2020**.
29. Liu, W.; Tang, J.; Cheng, Y.; Li, W.; Zheng, Y.; Liang, X. MedDG: An Entity-Centric Medical Consultation Dataset for Entity-Aware Medical Dialogue Generation, 2022, [\[arXiv:cs.CL/2010.07497\]](https://arxiv.org/abs/cs/2010.07497).
30. Tsatsaronis, G.; Balikas, G.; Malakasiotis, P.; Partalas, I.; Zschunke, M.; Alvers, M.R.; Weissenborn, D.; Krithara, A.; Petridis, S.; Polychronopoulos, D.; Almirantis, Y.; Pavlopoulos, J.; Baskiotis, N.; Gallinari, P.; Art  eres, T.; Ngomo, A.C.N.; Heino, N.; Gaussier, E.; Barrio-Alvers, L.; Schroeder, M.; Androutsopoulos, I.; Paliouras, G. An overview of the BIOASQ large-scale biomedical semantic indexing and question answering competition. *BMC Bioinformatics* **2015**, 16, 138. doi:10.1186/s12859-015-0564-6.
31. Abacha, A.B.; Agichtein, E.; Pinter, Y.; Demner-Fushman, D. Overview of the Medical Question Answering Task at TREC 2017 LiveQA. Text Retrieval Conference, 2017.

32. Lozano, A.; Fleming, S.L.; Chiang, C.C.; Shah, N. Clinfo.ai: An Open-Source Retrieval-Augmented Large Language Model System for Answering Medical Questions using Scientific Literature, 2023. [\[arXiv:cs.IR/2310.16146\]](https://arxiv.org/abs/2310.16146).

33. Jin, Q.; Dhingra, B.; Liu, Z.; Cohen, W.W.; Lu, X. PubMedQA: A Dataset for Biomedical Research Question Answering, 2019, [\[arXiv:cs.CL/1909.06146\]](https://arxiv.org/abs/1909.06146).

34. Jin, D.; Pan, E.; Oufattolle, N.; Weng, W.H.; Fang, H.; Szolovits, P. What Disease Does This Patient Have? A Large-Scale Open Domain Question Answering Dataset from Medical Exams. *Applied Sciences* **2021**, *11*. doi:10.3390/app11146421.

35. Ma, L.; Han, J.; Wang, Z.; Zhang, D. CephGPT-4: An Interactive Multimodal Cephalometric Measurement and Diagnostic System with Visual Large Language Model, 2023, [\[arXiv:cs.AI/2307.07518\]](https://arxiv.org/abs/2307.07518).

36. Chen, W.; Li, Z.; Fang, H.; Yao, Q.; Zhong, C.; Hao, J.; Zhang, Q.; Huang, X.; Peng, J.; Wei, Z. A Benchmark for Automatic Medical Consultation System: Frameworks, Tasks and Datasets. *Bioinformatics* **2022**, [/bioinformatics/btac817/48290490/btac817.pdf btac817](https://doi.org/10.1093/bioinformatics/btac817), doi:10.1093/bioinformatics/btac817.

37. Zeng, H. Measuring Massive Multitask Chinese Understanding, 2023, [\[arXiv:cs.CL/2304.12986\]](https://arxiv.org/abs/2304.12986).

38. Boteva, V.; Gholipour, D.; Sokolov, A.; Riezler, S. A Full-Text Learning to Rank Dataset for Medical Information Retrieval. *Advances in Information Retrieval*; Ferro, N.; Crestani, F.; Moens, M.F.; Mothe, J.; Silvestri, F.; Di Nunzio, G.M.; Hauff, C.; Silvello, G., Eds.; Springer International Publishing: Cham, 2016; pp. 716–722.

39. Roberts, K.; Alam, T.; Bedrick, S.; Demner-Fushman, D.; Lo, K.; Soboroff, I.; Voorhees, E.; Wang, L.L.; Hersh, W.R. TREC-COVID: rationale and structure of an information retrieval shared task for COVID-19. *Journal of the American Medical Informatics Association* **2020**, *27*, 1431–1436, [\[https://academic.oup.com/jamia/article-pdf/27/9/1431/34153771/ocaa091.pdf\]](https://academic.oup.com/jamia/article-pdf/27/9/1431/34153771/ocaa091.pdf). doi:10.1093/jamia/ocaa091.

40. Johnson, A.E.W.; Pollard, T.J.; Berkowitz, S.J.; Greenbaum, N.R.; Lungren, M.P.; Deng, C.y.; Mark, R.G.; Horng, S. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific Data* **2019**, *6*, 317. doi:10.1038/s41597-019-0322-0.

41. Ramesh, V.; Chi, N.A.; Rajpurkar, P. Improving Radiology Report Generation Systems by Removing Hallucinated References to Non-existent Priors, 2022, [\[arXiv:cs.CL/2210.06340\]](https://arxiv.org/abs/2210.06340).

42. Guo, Y.; Qiu, W.; Leroy, G.; Wang, S.; Cohen, T. Retrieval augmentation of large language models for lay language generation. *Journal of Biomedical Informatics* **2024**, *149*, 104580. doi:10.1016/j.jbi.2023.104580.

43. Pal, A.; Umapathi, L.K.; Sankarasubbu, M. MedMCQA: A Large-scale Multi-Subject Multi-Choice Dataset for Medical domain Question Answering. *Proceedings of the Conference on Health, Inference, and Learning*; Flores, G.; Chen, G.H.; Pollard, T.; Ho, J.C.; Naumann, T., Eds. PMLR, 2022, Vol. 174, *Proceedings of Machine Learning Research*, pp. 248–260.

44. Ben Abacha, A.; Mrabet, Y.; Sharp, M.; Goodwin, T.; Shooshan, S.E.; Demner-Fushman, D. Bridging the Gap between Consumers' Medication Questions and Trusted Answers. MEDINFO 2019, 2019.

45. Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; Steinhardt, J. Measuring Massive Multitask Language Understanding, 2021, [\[arXiv:cs.CY/2009.03300\]](https://arxiv.org/abs/2009.03300).

46. Xiong, G.; Jin, Q.; Lu, Z.; Zhang, A. Benchmarking Retrieval-Augmented Generation for Medicine, 2024, [\[arXiv:cs.CL/2402.13178\]](https://arxiv.org/abs/2402.13178).

47. Ge, J.; Sun, S.; Owens, J.; Galvez, V.; Gologorskaya, O.; Lai, J.C.; Pletcher, M.J.; Lai, K. Development of a Liver Disease-Specific Large Language Model Chat Interface using Retrieval Augmented Generation. *medRxiv* **2023**, [\[https://www.medrxiv.org/content/early/2023/11/11/2023.11.10.23298364.full.pdf\]](https://www.medrxiv.org/content/early/2023/11/11/2023.11.10.23298364.full.pdf). doi:10.1101/2023.11.10.23298364.

48. Long, C.; Subburam, D.; Lowe, K.; dos Santos, A.; Zhang, J.; Saduka, N.; Horev, Y.; Su, T.; Cote, D.; Wright, E. ChatENT: Augmented Large Language Models for Expert Knowledge Retrieval in Otolaryngology - Head and Neck Surgery. *medRxiv* **2023**, [\[https://www.medrxiv.org/content/early/2023/08/21/2023.08.18.23294283.full.pdf\]](https://www.medrxiv.org/content/early/2023/08/21/2023.08.18.23294283.full.pdf). doi:10.1101/2023.08.18.23294283.

49. Soong, D.; Sridhar, S.; Si, H.; Wagner, J.S.; Sá, A.C.C.; Yu, C.Y.; Karagoz, K.; Guan, M.; Hamadeh, H.; Higgs, B.W. Improving accuracy of GPT-3/4 results on biomedical data using a retrieval-augmented language model, 2023, [\[arXiv:cs.CL/2305.17116\]](https://arxiv.org/abs/2305.17116).

50. Thompson, W.E.; Vidmar, D.M.; Freitas, J.K.D.; Pfeifer, J.M.; Fornwalt, B.K.; Chen, R.; Altay, G.; Manghnani, K.; Nelsen, A.C.; Morland, K.; Stumpe, M.C.; Miotto, R. Large Language Models with Retrieval-Augmented Generation for Zero-Shot Disease Phenotyping, 2023, [\[arXiv:cs.AI/2312.06457\]](https://arxiv.org/abs/2312.06457).

51. Gao, Y.; Li, R.; Caskey, J.; Dligach, D.; Miller, T.; Churpek, M.M.; Afshar, M. Leveraging A Medical Knowledge Graph into Large Language Models for Diagnosis Prediction, 2023, [\[arXiv:cs.CL/2308.14321\]](https://arxiv.org/abs/cs.CL/2308.14321).
52. Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P.J. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer, 2023, [\[arXiv:cs.LG/1910.10683\]](https://arxiv.org/abs/cs.LG/1910.10683).
53. Jiang, X.; Zhang, R.; Xu, Y.; Qiu, R.; Fang, Y.; Wang, Z.; Tang, J.; Ding, H.; Chu, X.; Zhao, J.; Wang, Y. Think and Retrieval: A Hypothesis Knowledge Graph Enhanced Medical Large Language Models, 2023, [\[arXiv:cs.CL/2312.15883\]](https://arxiv.org/abs/cs.CL/2312.15883).
54. Xiao, S.; Liu, Z.; Zhang, P.; Muennighoff, N. C-Pack: Packaged Resources To Advance General Chinese Embedding, 2023, [\[arXiv:cs.CL/2309.07597\]](https://arxiv.org/abs/cs.CL/2309.07597).
55. Jin, M.; Yu, Q.; Zhang, C.; Shu, D.; Zhu, S.; Du, M.; Zhang, Y.; Meng, Y. Health-LLM: Personalized Retrieval-Augmented Disease Prediction Model, 2024, [\[arXiv:cs.CL/2402.00746\]](https://arxiv.org/abs/cs.CL/2402.00746).
56. Kang, B.; Kim, J.; Yun, T.R.; Kim, C.E. Prompt-RAG: Pioneering Vector Embedding-Free Retrieval-Augmented Generation in Niche Domains, Exemplified by Korean Medicine, 2024, [\[arXiv:cs.CL/2401.11246\]](https://arxiv.org/abs/cs.CL/2401.11246).
57. Ke, Y.; Jin, L.; Elangovan, K.; Abdullah, H.R.; Liu, N.; Sia, A.T.H.; Soh, C.R.; Tung, J.Y.M.; Ong, J.C.L.; Ting, D.S.W. Development and Testing of Retrieval Augmented Generation in Large Language Models – A Case Study Report, 2024, [\[arXiv:cs.CL/2402.01733\]](https://arxiv.org/abs/cs.CL/2402.01733).
58. Li, Y.; Li, Z.; Zhang, K.; Dan, R.; Jiang, S.; Zhang, Y. ChatDoctor: A Medical Chat Model Fine-Tuned on a Large Language Model Meta-AI (LLaMA) Using Medical Domain Knowledge, 2023, [\[arXiv:cs.CL/2303.14070\]](https://arxiv.org/abs/cs.CL/2303.14070).
59. Manathunga, S.S.; Illangasekara, Y.A. Retrieval Augmented Generation and Representative Vector Summarization for large unstructured textual data in Medical Education, 2023, [\[arXiv:cs.CL/2308.00479\]](https://arxiv.org/abs/cs.CL/2308.00479).
60. Markey, N.; El-Mansouri, I.; Rensonnet, G.; van Langen, C.; Meier, C. From RAGs to riches: Using large language models to write documents for clinical trials, 2024, [\[arXiv:cs.CL/2402.16406\]](https://arxiv.org/abs/cs.CL/2402.16406).
61. Miao, J.; Thongprayoon, C.; Suppadungsuk, S.; Garcia Valencia, O.A.; Cheungpasitporn, W. Integrating Retrieval-Augmented Generation with Large Language Models in Nephrology: Advancing Practical Applications. *Medicina* 2024, 60. doi:10.3390/medicina60030445.
62. Murugan, M.; Yuan, B.; Venner, E.; Ballantyne, C.M.; Robinson, K.M.; Coons, J.C.; Wang, L.; Empey, P.E.; Gibbs, R.A. Empowering personalized pharmacogenomics with generative AI solutions. *Journal of the American Medical Informatics Association* 2024, p. ocae039, [\[https://academic.oup.com/jamia/advance-article-pdf/doi/10.1093/jamia/ocae039/56880048/ocae039.pdf\]](https://academic.oup.com/jamia/advance-article-pdf/doi/10.1093/jamia/ocae039/56880048/ocae039.pdf). doi:10.1093/jamia/ocae039.
63. Neupane, S.; Mitra, S.; Mittal, S.; Golilarz, N.A.; Rahimi, S.; Amirlatifi, A. MedInsight: A Multi-Source Context Augmentation Framework for Generating Patient-Centric Medical Responses using Large Language Models, 2024, [\[arXiv:cs.CL/2403.08607\]](https://arxiv.org/abs/cs.CL/2403.08607).
64. Ong, J.C.L.; Jin, L.; Elangovan, K.; Lim, G.Y.S.; Lim, D.Y.Z.; Sng, G.G.R.; Ke, Y.; Tung, J.Y.M.; Zhong, R.J.; Koh, C.M.Y.; Lee, K.Z.H.; Chen, X.; Chng, J.K.; Than, A.; Goh, K.J.; Ting, D.S.W. Development and Testing of a Novel Large Language Model-Based Clinical Decision Support Systems for Medication Safety in 12 Clinical Specialties, 2024, [\[arXiv:cs.CL/2402.01741\]](https://arxiv.org/abs/cs.CL/2402.01741).
65. Chen, J.; Xiao, S.; Zhang, P.; Luo, K.; Lian, D.; Liu, Z. BGE M3-Embedding: Multi-Lingual, Multi-Functionality, Multi-Granularity Text Embeddings Through Self-Knowledge Distillation, 2024, [\[arXiv:cs.CL/2402.03216\]](https://arxiv.org/abs/cs.CL/2402.03216).
66. Parmanto, B.; Aryoyudanta, B.; Soekinto, W.; Setiawan, I.M.A.; Wang, Y.; Hu, H.; Saptono, A.; Choi, Y.K. Development of a Reliable and Accessible Caregiving Language Model (CaLM), 2024, [\[arXiv:cs.CL/2403.06857\]](https://arxiv.org/abs/cs.CL/2403.06857).
67. Quidwai, M.A.; Lagana, A. A RAG Chatbot for Precision Medicine of Multiple Myeloma. *medRxiv* 2024, [\[https://www.medrxiv.org/content/early/2024/03/18/2024.03.14.24304293.full.pdf\]](https://www.medrxiv.org/content/early/2024/03/18/2024.03.14.24304293.full.pdf). doi:10.1101/2024.03.14.24304293.
68. Ranjit, M.; Ganapathy, G.; Manuel, R.; Ganu, T. Retrieval Augmented Chest X-Ray Report Generation using OpenAI GPT models, 2023, [\[arXiv:cs.CL/2305.03660\]](https://arxiv.org/abs/cs.CL/2305.03660).
69. Li, J.; Selvaraju, R.R.; Gotmare, A.D.; Joty, S.; Xiong, C.; Hoi, S. Align before Fuse: Vision and Language Representation Learning with Momentum Distillation, 2021, [\[arXiv:cs.CV/2107.07651\]](https://arxiv.org/abs/cs.CV/2107.07651).
70. Rau, A.; Rau, S.; Zöller, D.; Fink, A.; Tran, H.; Wilpert, C.; Nattenmüller, J.; Neubauer, J.; Bamberg, F.; Reisert, M.; Russe, M.F. A Context-based Chatbot Surpasses Radiologists and Generic ChatGPT in Following the ACR Appropriateness Guidelines. *Radiology* 2023, 308, e230970, [\[https://doi.org/10.1148/radiol.230970\]](https://doi.org/10.1148/radiol.230970). doi:10.1148/radiol.230970.

71. Russe, M.F.; Fink, A.; Ngo, H.; Tran, H.; Bamberg, F.; Reisert, M.; Rau, A. Performance of ChatGPT, human radiologists, and context-aware ChatGPT in identifying AO codes from radiology reports. *Scientific Reports* **2023**, *13*, 14215. doi:10.1038/s41598-023-41512-8.

72. Shi, W.; Zhuang, Y.; Zhu, Y.; Iwinski, H.; Wattenbarger, M.; Wang, M.D. Retrieval-Augmented Large Language Models for Adolescent Idiopathic Scoliosis Patients in Shared Decision-Making. Proceedings of the 14th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics; Association for Computing Machinery: New York, NY, USA, 2023; BCB '23. doi:10.1145/3584371.3612956.

73. Soman, K.; Rose, P.W.; Morris, J.H.; Akbas, R.E.; Smith, B.; Peetoom, B.; Villouta-Reyes, C.; Ceroni, G.; Shi, Y.; Rizk-Jackson, A.; Israni, S.; Nelson, C.A.; Huang, S.; Baranzini, S.E. Biomedical knowledge graph-enhanced prompt generation for large language models, 2023, [arXiv:cs.CL/2311.17330].

74. Unlu, O.; Shin, J.; Mailly, C.J.; Oates, M.F.; Tucci, M.R.; Varughese, M.; Wagholicar, K.; Wang, F.; Scirica, B.M.; Blood, A.J.; Aronson, S.J. Retrieval Augmented Generation Enabled Generative Pre-Trained Transformer 4 (GPT-4) Performance for Clinical Trial Screening. *medRxiv* **2024**, [https://www.medrxiv.org/content/early/2024/02/08/2024.02.08.24302376.full.pdf]. doi:10.1101/2024.02.08.24302376.

75. Vaid, A.; Lampert, J.; Lee, J.; Sawant, A.; Apakama, D.; Sakhija, A.; Soroush, A.; Lee, D.; Landi, I.; Bussola, N.; Nabeel, I.; Freeman, R.; Kovatch, P.; Carr, B.; Glicksberg, B.; Argulian, E.; Lerakis, S.; Kraft, M.; Charney, A.; Nadkarni, G. Generative Large Language Models are autonomous practitioners of evidence-based medicine, 2024, [arXiv:cs.AI/2401.02851].

76. Wang, Y.; Ma, X.; Chen, W. Augmenting Black-box LLMs with Medical Textbooks for Clinical Question Answering, 2023, [arXiv:cs.CL/2309.02233].

77. Wang, J.; Yang, Z.; Yao, Z.; Yu, H. JMLR: Joint Medical LLM and Retrieval Training for Enhancing Reasoning and Professional Question Answering Capability, 2024, [arXiv:cs.CL/2402.17887].

78. Wornow, M.; Lozano, A.; Dash, D.; Jindal, J.; Mahaffey, K.W.; Shah, N.H. Zero-Shot Clinical Trial Patient Matching with LLMs, 2024, [arXiv:cs.CL/2402.05125].

79. Yu, H.; Guo, P.; Sano, A. Zero-Shot ECG Diagnosis with Large Language Models and Retrieval-Augmented Generation. Proceedings of the 3rd Machine Learning for Health Symposium; Hegselmann, S.; Parziale, A.; Shanmugam, D.; Tang, S.; Asiedu, M.N.; Chang, S.; Hartvigsen, T.; Singh, H., Eds. PMLR, 2023, Vol. 225, *Proceedings of Machine Learning Research*, pp. 650–663.

80. Ziletti, A.; D'Ambrosi, L. Retrieval augmented text-to-SQL generation for epidemiological question answering using electronic health records, 2024, [arXiv:cs.CL/2403.09226].

81. Team, L.D. LlamaIndex: A Framework for Context-Augmented LLM Applications, 2024.

82. Team, L.D. Langchain: A framework for developing applications powered by large language models, 2024.

83. Cuconasu, F.; Trappolini, G.; Siciliano, F.; Filice, S.; Campagnano, C.; Maarek, Y.; Tonello, N.; Silvestri, F. The Power of Noise: Redefining Retrieval for RAG Systems, 2024, [arXiv:cs.IR/2401.14887].

84. Yang, R.; Marrese-Taylor, E.; Ke, Y.; Cheng, L.; Chen, Q.; Li, I. Integrating UMLS Knowledge into Large Language Models for Medical Question Answering, 2023, [arXiv:cs.CL/2310.02778].

85. Varshney, D.; Zafar, A.; Behera, N.K.; Ekbal, A. Knowledge graph assisted end-to-end medical dialog generation. *Artificial Intelligence in Medicine* **2023**, *139*, 102535. doi:https://doi.org/10.1016/j.artmed.2023.102535.

86. Varshney, D.; Zafar, A.; Behera, N.K.; Ekbal, A. Knowledge grounded medical dialogue generation using augmented graphs. *Scientific Reports* **2023**, *13*, 3310. doi:10.1038/s41598-023-29213-8.

87. Morris, J.H.; Soman, K.; Akbas, R.E.; Zhou, X.; Smith, B.; Meng, E.C.; Huang, C.C.; Ceroni, G.; Schenk, G.; Rizk-Jackson, A.; Harroud, A.; Sanders, L.; Costes, S.V.; Bharat, K.; Chakraborty, A.; Pico, A.R.; Mardirossian, T.; Keiser, M.; Tang, A.; Hardi, J.; Shi, Y.; Musen, M.; Israni, S.; Huang, S.; Rose, P.W.; Nelson, C.A.; Baranzini, S.E. The scalable precision medicine open knowledge engine (SPOKE): a massive knowledge graph of biomedical information. *Bioinformatics* **2023**, *39*, btad080, [https://academic.oup.com/bioinformatics/article-pdf/39/2/btad080/49278346/btad080.pdf]. doi:10.1093/bioinformatics/btad080.

88. Jin, Q.; Wang, Z.; Floudas, C.S.; Sun, J.; Lu, Z. Matching Patients to Clinical Trials with Large Language Models, 2023, [arXiv:cs.CL/2307.15051].

89. Douze, M.; Guzhva, A.; Deng, C.; Johnson, J.; Szilvassy, G.; Mazaré, P.E.; Lomeli, M.; Hosseini, L.; Jégou, H. The Faiss library, 2024, [arXiv:cs.LG/2401.08281].

90. Lin, X.V.; Chen, X.; Chen, M.; Shi, W.; Lomeli, M.; James, R.; Rodriguez, P.; Kahn, J.; Szilvassy, G.; Lewis, M.; Zettlemoyer, L.; Yih, S. RA-DIT: Retrieval-Augmented Dual Instruction Tuning, 2023, [arXiv:cs.CL/2310.01352].

91. Chen, T.; Guestrin, C. Xgboost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 785–794.
92. Luo, L.; Li, Y.F.; Haffari, G.; Pan, S. Reasoning on Graphs: Faithful and Interpretable Large Language Model Reasoning, 2024, [\[arXiv:cs/2310.01061\]](https://arxiv.org/abs/cs/2310.01061).
93. Hong, S.; Zhuge, M.; Chen, J.; Zheng, X.; Cheng, Y.; Zhang, C.; Wang, J.; Wang, Z.; Yau, S.K.S.; Lin, Z.; Zhou, L.; Ran, C.; Xiao, L.; Wu, C.; Schmidhuber, J. MetaGPT: Meta Programming for A Multi-Agent Collaborative Framework, 2023, [\[arXiv:cs.AI/2308.00352\]](https://arxiv.org/abs/cs/2308.00352).
94. Weng, L. LLM-powered Autonomous Agents. *lilianweng.github.io* 2023.
95. Es, S.; James, J.; Espinosa-Anke, L.; Schockaert, S. RAGAS: Automated Evaluation of Retrieval-Augmented Generation, 2023, [\[arXiv:cs/2309.15217\]](https://arxiv.org/abs/cs/2309.15217).
96. Team, C.E.D. Continuous Eval: an open-source package created for granular and holistic evaluation of GenAI application pipelines, 2024.
97. Team, T.D. TruLens: Evaluate and Track LLM Applications, 2024.
98. Team, D.D. DeepEval: the open-source LLM evaluation framework, 2024.
99. Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; Stoyanov, V. RoBERTa: A Robustly Optimized BERT Pretraining Approach, 2019, [\[arXiv:cs/1907.11692\]](https://arxiv.org/abs/cs/1907.11692).
100. Steinbock, B. *The Oxford Handbook of Bioethics*; Oxford University Press, 2009. doi:10.1093/oxfordhb/9780199562411.001.0001.
101. Carlini, N.; Ippolito, D.; Jagielski, M.; Lee, K.; Tramer, F.; Zhang, C. Quantifying Memorization Across Neural Language Models, 2023, [\[arXiv:cs.LG/2202.07646\]](https://arxiv.org/abs/cs/2202.07646).
102. Zhang, C.; Ippolito, D.; Lee, K.; Jagielski, M.; Tramèr, F.; Carlini, N. Counterfactual Memorization in Neural Language Models, 2023, [\[arXiv:cs/2112.12938\]](https://arxiv.org/abs/cs/2112.12938).
103. Lee, J.; Le, T.; Chen, J.; Lee, D. Do Language Models Plagiarize? Proceedings of the ACM Web Conference 2023; Association for Computing Machinery: New York, NY, USA, 2023; WWW '23, p. 3637–3647. doi:10.1145/3543507.3583199.
104. Zeng, S.; Zhang, J.; He, P.; Xing, Y.; Liu, Y.; Xu, H.; Ren, J.; Wang, S.; Yin, D.; Chang, Y.; Tang, J. The Good and The Bad: Exploring Privacy Issues in Retrieval-Augmented Generation (RAG), 2024, [\[arXiv:cs/2402.16893\]](https://arxiv.org/abs/cs/2402.16893).
105. Vassilev, A.; Oprea, A.; Fordyce, A.; Anderson, H. Adversarial machine learning: A taxonomy and terminology of attacks and mitigations. Technical report, National Institute of Standards and Technology, 2024.
106. Wu, S.; Koo, M.; Blum, L.; Black, A.; Kao, L.; Fei, Z.; Scalzo, F.; Kurtz, I. Benchmarking Open-Source Large Language Models, GPT-4 and Claude 2 on Multiple-Choice Questions in Nephrology. *NEJM AI* 2024, 1, AIdbp2300092, [\[https://ai.nejm.org/doi/pdf/10.1056/AIdbp2300092\]](https://ai.nejm.org/doi/pdf/10.1056/AIdbp2300092). doi:10.1056/AIdbp2300092.
107. Greene, R.; Sanders, T.; Weng, L.; Neelakantan, A. New and improved embedding model, Dec. 15, 2022.
108. Amugongo, L.M.; Bidwell, N.J.; Corrigan, C.C. Invigorating Ubuntu Ethics in AI for healthcare: Enabling equitable care. Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency; Association for Computing Machinery: New York, NY, USA, 2023; FAccT '23, p. 583–592. doi:10.1145/3593013.3594024.
109. Bergman, A.S.; Hendricks, L.A.; Rauh, M.; Wu, B.; Agnew, W.; Kunesch, M.; Duan, I.; Gabriel, I.; Isaac, W. Representation in AI Evaluations. Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency; Association for Computing Machinery: New York, NY, USA, 2023; FAccT '23, p. 519–533. doi:10.1145/3593013.3594019.
110. Tang, Y.; Yang, Y. MultiHop-RAG: Benchmarking Retrieval-Augmented Generation for Multi-Hop Queries, 2024, [\[arXiv:cs/2401.15391\]](https://arxiv.org/abs/cs/2401.15391).
111. Goodwin, T.R.; Demner-Fushman, D. Clinical Language Understanding Evaluation (CLUE), 2022, [\[arXiv:cs/2209.14377\]](https://arxiv.org/abs/cs/2209.14377).
112. Ahmad, M.A.; Yaramis, I.; Roy, T.D. Creating Trustworthy LLMs: Dealing with Hallucinations in Healthcare AI, 2023, [\[arXiv:cs/2311.01463\]](https://arxiv.org/abs/cs/2311.01463).
113. Wu, Y.; Zhu, J.; Xu, S.; Shum, K.; Niu, C.; Zhong, R.; Song, J.; Zhang, T. RAGTruth: A Hallucination Corpus for Developing Trustworthy Retrieval-Augmented Language Models, 2023, [\[arXiv:cs/2401.00396\]](https://arxiv.org/abs/cs/2401.00396).
114. Gu, A.; Dao, T. Mamba: Linear-Time Sequence Modeling with Selective State Spaces, 2023, [\[arXiv:cs.LG/2312.00752\]](https://arxiv.org/abs/cs/2312.00752).

115. Gemini Team, G. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. [https://storage.googleapis.com/deepmind-media/gemini/gemini\\_v1\\_5\\_report.pdf](https://storage.googleapis.com/deepmind-media/gemini/gemini_v1_5_report.pdf), 2024. [Accessed 26-02-2024].
116. Wu, K.; Wu, E.; Cassasola, A.; Zhang, A.; Wei, K.; Nguyen, T.; Riantawan, S.; Riantawan, P.S.; Ho, D.E.; Zou, J. How well do LLMs cite relevant medical references? An evaluation framework and analyses, 2024, [arXiv:cs.CL/2402.02008].
117. Li, W.; Li, J.; Ma, W.; Liu, Y. Citation-Enhanced Generation for LLM-based Chatbots, 2024, [arXiv:cs.CL/2402.16063].
118. Liu, J.; Hu, T.; Zhang, Y.; Gai, X.; Feng, Y.; Liu, Z. A ChatGPT Aided Explainable Framework for Zero-Shot Medical Image Diagnosis, 2023, [arXiv:eess.IV/2307.01981].
119. Zhu, Y.; Ren, C.; Xie, S.; Liu, S.; Ji, H.; Wang, Z.; Sun, T.; He, L.; Li, Z.; Zhu, X.; Pan, C. REALM: RAG-Driven Enhancement of Multimodal Electronic Health Records Analysis via Large Language Models, 2024, [arXiv:cs.AI/2402.07016].
120. Amugongo, L.M.; Kriebitz, A.; Boch, A.; Lütge, C. Operationalising AI ethics through the agile software development lifecycle: a case study of AI-enabled mobile health applications. *AI and Ethics* 2023. doi:10.1007/s43681-023-00331-3.
121. Tucci, V.; Saary, J.; Doyle, T.E. Factors influencing trust in medical artificial intelligence for healthcare professionals: a narrative review. *Journal of Medical Artificial Intelligence* 2021, 5.
122. Amugongo, L.M.; Kriebitz, A.; Boch, A.; Lütge, C. Mobile Computer Vision-Based Applications for Food Recognition and Volume and Calorific Estimation: A Systematic Review. *Healthcare* 2023, 11.

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