

Improving Health Question Answering with Reliable and Time-Aware Evidence Retrieval

Juraj Vladika and Florian Matthes

Department of Computer Science

Technical University of Munich

Garching, Germany

{juraj.vladika, matthes}@tum.de

Abstract

In today’s digital world, seeking answers to health questions on the Internet is a common practice. However, existing question answering (QA) systems often rely on using pre-selected and annotated evidence documents, thus making them inadequate for addressing novel questions. Our study focuses on the open-domain QA setting, where the key challenge is to first uncover relevant evidence in large knowledge bases. By utilizing the common retrieve-then-read QA pipeline and PubMed as a trustworthy collection of medical research documents, we answer health questions from three diverse datasets. We modify different retrieval settings to observe their influence on the QA pipeline’s performance, including the number of retrieved documents, sentence selection process, the publication year of articles, and their number of citations. Our results reveal that cutting down on the amount of retrieved documents and favoring more recent and highly cited documents can improve the final macro F1 score up to 10%. We discuss the results, highlight interesting examples, and outline challenges for future research, like managing evidence disagreement and crafting user-friendly explanations.

1 Introduction

In the digital era, using the Internet to search for health information has become a prevalent behavior (Jia et al., 2021). Users turn to seek health advice online due to its ease of access, wide coverage of information, convenience of searching, interactivity, and anonymity (Neely et al., 2021). Health information sought online includes anything regarding the symptoms and treatments of different diseases. In general, health information seeking can lead to enhanced patient involvement in medical decision-making, improved communication with care providers, and improved quality of life (Rutten et al., 2019). Nevertheless, finding trustworthy and relevant evidence in abundant online content

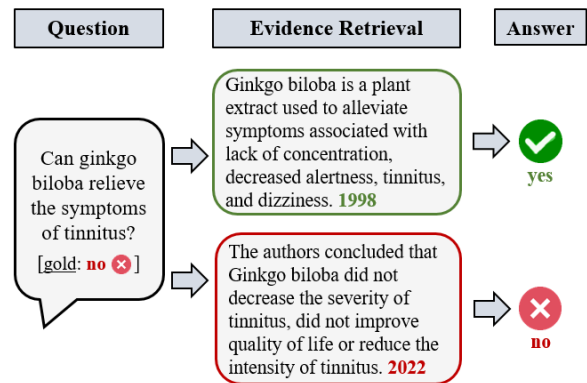


Figure 1: The question-answering system used in our study and an example question with two retrieved evidence documents and final predictions. This example shows how retrieving an outdated study caused an incorrect prediction (top), while a more recent study resulted with an accurate answer (bottom).

remains an open challenge (Battineni et al., 2020). Especially in the medical field, clinical recommendations can change with time, so finding the latest evidence is essential for reliable answers.

Interacting with online search engines and conversational systems is done with *question answering (QA)*. Technical solutions based on Machine Learning (ML) and Natural Language Processing (NLP) aim to automate this task and provide users with reliable answers to their inquiries. The purpose of QA systems is multi-fold: it helps scientists verify their research hypothesis by finding related studies, it allows lay users to find answers to their everyday health concerns, and enables factuality assessment of generative language models by fact-checking their responses over trusted evidence (Jin et al., 2021; Vladika and Matthes, 2023a).

While QA can work with answering questions over a provided document, we are focusing on the more realistic and challenging problem of *open-domain question answering*, where extensive collections of documents with diverse topics have to

be queried to find the relevant evidence (Chen and Yih, 2020a). The open-domain QA systems usually consist of two main components, a retriever and a reader (Zhu et al., 2021b). The retriever’s task is finding the relevant documents that will serve as the main source of evidence for answering the question. The reader (QA module) performs the reasoning process between the question and found evidence, and produces the final answer. While both components are essential for the system, we posit that the retrieval is a more challenging part, considering that the QA module receives the input from it and the quality of the final answer depends on the retrieved documents (Sauchuk et al., 2022).

Retrieving credible evidence documents relevant to the query ensures the final output’s quality. This is true for both the retrieval-powered text classification tasks and recently popular retrieval-augmented generation (RAG) approaches (Cuconasu et al., 2024). While much progress has been made in open-domain QA, addressing the challenges in retrieval settings still needs to be explored. These include assessing the quantity of documents needed to be retrieved for a reliable answer, the amount of evidence passages extracted from them, and the quality of the documents themselves, such as their recency and strength of findings. Figure 1 shows an example of a health question answered with two different retrieved documents – the more recent one has more up-to-date knowledge and findings.

To bridge this research gap, in this study, we perform an array of experiments to test the predictive performance of an open-domain QA system with different evidence retrieval strategies. We use three diverse datasets of biomedical and health questions, which contain discreet labels like *yes* and *no* as their final answers, and use their gold labels as ground truth. We use the large collection of 20 million biomedical research abstracts from PubMed as the knowledge base for evidence retrieval. By keeping the reader (answering module) fixed, we only vary the different retrieval aspects and measure the change in the QA performance as measured by classification metrics precision, recall, and F1. The settings we explore include the number of documents retrieved and sentences extracted from them, the articles’ publication year, and their number of citations. Our findings demonstrate that the QA performance is improved by accounting for the amount and quality of the retrieved documents.

To summarize, our research contributions are:

- We evaluate the performance of an open-domain QA pipeline for health questions, using biomedical research papers as evidence source, concerning the different number of documents retrieved and sentences extracted from them.
- Additionally, we evaluate the influence of the evidence quality parameters like year of publication and number of citations on the final predictive performance of the QA system, showing that time-aware evidence retrieval leads to improved performance.
- Finally, we take a deeper look into the results and provide insights with a qualitative analysis. We report on the problem of evidence disagreement and provide future directions on developing better health question answering systems to be deployed in the future.

We make our code and datasets publicly available in a GitHub repository.¹

2 Related work

In this section, we outline the work related to our study.

2.1 Biomedical Question Answering

Question Answering (QA) is a rapidly evolving knowledge-intensive NLP task, with over 80 QA datasets released in last two years (Rogers et al., 2023). Based on the availability of evidence for the question, it can be analyzed in a closed-domain or an open-domain setting. In closed-domain QA, the evidence comes from an already provided source document. This setting is also called Machine Reading Comprehension (MRC) since the goal is to build models that can comprehend from the given text how to answer the posed question (Baradaran et al., 2022). In open-domain QA, to which our work belongs, only the question and its final answer are known, and the QA system has to find appropriate evidence in a large document corpus or other type of collection (Chen and Yih, 2020b).

Based on the topic of questions, our work is related to the research on *science question answering*, aiming to answer questions related to natural sciences from resources like school curricula (Lu et al., 2022) or scholarly articles (Lee et al., 2023). More

¹<https://github.com/jvladika/Improving-Health-QA>

precisely, our work is part of *biomedical question answering* (Jin et al., 2022).

Biomedical QA can help biomedical researchers conduct their work by answering complex research questions (Jin et al., 2019), help clinical practitioners by answering clinical questions over health records (Vilares and Gómez-Rodríguez, 2019), or help consumers answer questions about their health concerns (Demner-Fushman et al., 2020). The last category, also called *health question answering*, is increasingly being adopted by consumers due to the rising popularity of conversational assistants (Budler et al., 2023). In our work, we cover both the datasets related to QA for biomedical research and consumer health.

2.2 Open-Domain Fact Verification and QA

Considering that the datasets we use only contain questions with discrete (*yes/no*) answers, our work is related to the task of automated fact verification (fact-checking). This task aims to verify the veracity of a factual claim based on credible evidence that supports it, refutes it, or does not provide enough information (Guo et al., 2022). Recent years have seen a rise in fact-checking datasets focusing on scientific knowledge, in particular health and medicine (Vladika and Matthes, 2023a).

While fact verification literature often works in a closed-domain setting with evidence documents provided, some recent work also explores the open-domain setting. Wadden et al. (2022) observe significant performance drops in F1 scores of final verdict predictions when increasing the evidence corpus from a few thousand to half a million documents. Pugachev et al. (2023) analyze how well consumer-health questions can be answered with built-in search engines of PubMed and Wikipedia. Expanding the scope even more, Vladika and Matthes (2024) compare the performance of semantic search and BM25 over PubMed and Wikipedia, as well as Google search, for verifying biomedical and health questions.

Some studies have analyzed the influence of time and quantity in evidence retrieval on downstream tasks. Allein et al. (2021) trained time-aware evidence ranking models for time-sensitive news claims and show performance improvement. Likewise, Schlichtkrull et al. (2024) constructed a dataset where evidence for given claims only appeared after the claim itself, thus eliminating temporal leaks. Regarding the document quantity,

Oh and Thorne (2023) analyze the influence of the number of retrieved evidence passages on the QA performance over two general QA datasets, showing that the performance often actually drops with the increasing number of retrieved snippets.

To the best of our knowledge, our paper uses the largest document collection so far for open-domain health QA by indexing the entire PubMed corpus of more than 20 million articles. Likewise, it is also the first study to test the influence of the number of documents retrieved on the final QA performance instead of fixing it to a certain number, like the commonly found 5 (Thorne et al., 2018) or 6 (Wadden et al., 2022); as well as testing the influence of the different number of sentences retrieved. While there has been existing research on time-aware evidence ranking in fact verification for news claims, our work is first to explore the temporal aspect for biomedical questions, as well as other evidence quality aspects like the number of citations of retrieved publications.

3 Foundations

In this section, we explain the foundations of the study, including the used datasets, the evidence corpus, and the structure of the used QA system.

3.1 Datasets

We chose three datasets of biomedical and health claims in English, built for different purposes. We only use the questions and final labels (answers) from the datasets in our experiments. While the datasets provide the gold evidence passages used to derive the answers, we do not utilize them since the idea of our open-domain QA setting is that the retriever component has to discover the relevant evidence in a large document corpus.

HealthFC (Vladika et al., 2023) is a question-answering and fact-checking dataset focusing on consumer health questions and common topics users search health advice online for. It includes diverse topics like dietary supplements, heart and lungs, reproductive health, cancer, and mental health. Medical practitioners manually answered and verified all the questions using the evidence from systematic reviews and clinical trials. There are 750 questions in total, out of which 205 are supported, 122 are refuted, and for 433 questions, there is not enough information (NEI) to answer. We use two variants of the dataset: **HealthFC-3**, which has all 750 claims and all three classes; and

HealthFC-2, which only has 327 supported and refuted claims with two classes.

BioASQ-7b (Nentidis et al., 2020) is a biomedical question-answering dataset constructed by biomedical researchers and designed to reflect the real information needs of biomedical experts. It is part of the ongoing challenge of the same name, focusing on biomedical semantic indexing and question answering. The evidence for answers comes from biomedical research publications, i.e., the same corpus of PubMed used in our study. Other than only exact answers, the BioASQ dataset also includes ideal answer summaries. The 7b version of the dataset we use has 745 claims, of which 614 are supported ("yes"), and 131 are refuted ("no").

TREC-Health (Pugachev et al., 2023) is a dataset of 117 popular health questions originating from two TREC shared challenges. TREC is an ongoing series of workshops centering on challenges in accurate information retrieval (Voorhees et al., 2005). The questions stem from two shared tasks: the TREC 2019 Decision Track (Abualsaud et al., 2020) and the TREC 2021 Health Misinformation Track (Clarke et al., 2021), both focusing on challenges with incorrect search engine results for health (mis)information. Questions cover common consumer health concerns, similar to HealthFC, but the two datasets do not overlap. The dataset consists of 113 questions, of which 61 are supported ("yes") and 52 refuted ("no").

Table 1 gives an overview of the four datasets. With HealthFC and TREC-Health, we aim to target common health questions users would pose to a QA system, while BioASQ is selected in order to explore how do the QA results change for more complex, expert-gearred biomedical questions.

Dataset	Domain	✓	✗	?
HealthFC-3	everyday health	202	125	433
HealthFC-2	everyday health	202	125	—
BioASQ-7b	biomedical research	614	131	—
TREC-Health	consumer health	61	52	—

Table 1: The four datasets used in the experiments, including their domain and label distribution. (✓ – supported, ✗ – refuted, ? – not enough information)

3.2 Evidence Corpus

We approach the QA task in the open-domain setting, meaning that evidence is unknown when the question is posed and must first be discovered in a vast evidence collection. Given that we work with medical and health-related questions, we chose a collection of biomedical research publications as the source of evidence.

Our evidence corpus originates from PubMed, a large and trustworthy knowledge base of biomedical research publications (Canese and Weis, 2013). Considering that the full text of most of these publications is not freely accessible, we use only the abstracts, which are always available. This does not hinder the performance since medical abstracts often already include a verdict on their main research hypothesis. The US National Library of Medicine provides every year MEDLINE, a snapshot of currently available abstracts in PubMed that is updated once a year. We used the 2022 version found on the official website.² While this yields 33.4M abstracts, we pre-processed the data by removing any non-English papers, papers with no abstracts, and papers with unfinished abstracts, which yields 20.6 million abstracts.

3.3 QA System

The question-answering system used for our experiments is in the form of a pipeline, based on the pipeline system from Vladika and Matthes (2023b). This pipeline consists of two main parts: a *retriever* and a *reader*. The process of question answering is done sequentially, by first retrieving the evidence, performing reasoning over it, and finally producing a final answer. Our experiments focus on changing the different aspects of the retriever while keeping the reader completely fixed. This ensures that the experimental setup is consistent and that only one parameter is tested at a time.

In the retriever, given a question q and a corpus of n documents $D = \{d_1, d_2, \dots, d_n\}$, the goal is to select the top k most relevant documents g_1, \dots, g_k for the given query. The selection is done with a function $search(q, d)$, which compares the similarity of the question (query) and each document in the corpus. The documents in our corpus are abstracts of medical publications. While abstracts are shorter versions of full documents, they can still contain irrelevant sentences for producing the

²https://www.nlm.nih.gov/databases/download/pubmed_medline.html

Top k docs	HealthFC-3			HealthFC-2			TREC			BioASQ		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
1	49.2	44.1	40.1	63.3	62.7	62.9	61.7	60.9	60.7	71.8	57.4	58.8
5	53.9	44.7	38.8	67.6	57.8	55.2	63.8	58.7	55.8	65.5	65.0	65.2
10	55.1	45.2	39.5	66.4	56.5	53.1	67.6	63.0	57.2	67.7	63.9	65.4
15	58.2	44.4	39.1	67.5	55.6	51.1	65.6	58.4	54.5	68.2	63.4	65.1
20	49.8	43.7	38.5	63.8	54.9	50.5	66.9	58.3	53.7	67.5	63.0	64.5
50	48.4	42.9	38.1	62.2	53.3	47.6	67.5	57.1	51.4	67.5	60.4	62.1
100	47.5	44.3	37.6	56.3	44.6	44.6	64.5	56.3	50.8	66.3	59.6	61.2

Table 2: Results of final answer prediction over four datasets, with different values of *top k* documents retrieved during the process and used for majority voting. All scores are macro averaged.

final verdict.

In our first experiment, we use full documents g_1, \dots, g_k and the question q to predict the final answer. In the second experiment, we select only the top j most relevant sentences from the abstracts. From m candidate sentences s_1, s_2, \dots, s_m comprising the selected documents, top j sentences are selected as evidence sentences $\vec{e} = e_1, e_2, \dots, e_n$ with a function $select(q, s)$. These sentences are the most similar to the question q .

Finally, the answer is predicted from the given question q and evidence \vec{e} , where the evidence is either the complete documents (the first round of experiments) or a set of sentences (the second round of experiments). The reader model produces the final verdict and is one of three classes $y(q, \vec{e}) \in \{Refuted(0), Supported(1), Not\ Enough\ Information(2)\}$. While QA can be generative and elicit long answers, all datasets we use contain only the short *yes/no/unknown* answers. This makes using the standard classification metrics precision, recall, and F1 possible. We model answer prediction as the related task of recognizing entailment or natural language inference (NLI) and use an NLI model for the prediction.

In all experiments, majority voting is used to determine the final verdict. For the dataset HealthFC-3 with three classes, this can be one of the three classes (0, 1, 2). For other datasets, the majority is taken only from predictions of 0 and 1 (in case of a tie, 0 is predicted). Majority voting is chosen for convenience, but it is not optimal as the information on prediction disagreement. Future work should explore how to model the disagreement better.

For the $search(q, d)$ function that selects *top k* most relevant documents, we use BM25 as it was proven to be a strong baseline for retrieving documents for automated claim verification (Stamm-

bach et al., 2023). It also ensures higher precision at the cost of coverage, which aligns with our use case – we want the retrieved documents to be relevant before being passed to the reader. We use a sentence embedding model and semantic search to select the sentences most similar to our query from abstracts. For $select(q, s)$, we select the model SPICED (Wright et al., 2022), which is a recent sentence similarity model optimized for paraphrases of scientific claims. For the final answer prediction model (reader) $y(q, \vec{e})$, we choose the DeBERTav3 model (He et al., 2021) since it was shown to be a highly potent model for natural language understanding and reasoning tasks. We use the variant of the model optimized for NLI prediction (Laurer et al., 2024). We do not fine-tune the models on the datasets in our experiments because we want to simulate a realistic QA system that has to answer unseen questions.

4 Experiments

We conduct three groups of experiments to test the influence of different retrieval parameters on the performance of our QA system.

Number of retrieved documents. The first group of experiments consisted of keeping the QA pipeline consistent but increasing the number of retrieved documents (*top k*) that are forwarded to the final QA module. The motivation behind this was to find the fine balance between covering enough different studies but not saturating the module with noise and irrelevant articles. Considering the increasing popularity of retrieval-enhanced systems such as retrieval-augmented generation (RAG) pipelines (Lewis et al., 2020), retrieving only the relevant amount of documents or chunks is a significant challenge. We use BM25 as the retrieval technique because of its efficiency and its

Top j sents	HealthFC-3			HealthFC-2			TREC			BioASQ		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
1	40.7	40.0	35.9	57.4	57.4	57.4	59.8	59.3	59.1	56.9	62.0	52.6
3	39.0	41.7	36.6	60.2	56.9	56.0	59.3	57.6	56.4	58.8	64.1	58.2
5	52.2	44.6	39.8	59.8	56.3	54.9	59.7	57.3	59.3	58.5	62.8	58.6
10	46.9	44.4	40.0	61.4	56.9	55.3	63.1	58.8	56.4	50.5	62.5	60.2
15	45.4	44.2	39.8	61.9	57.3	55.8	60.9	56.9	53.8	60.7	64.2	61.5
20	52.6	44.5	40.6	64.3	58.2	56.8	61.5	58.0	55.7	60.9	64.4	61.7

Table 3: Results of final answer prediction over four datasets, with different values of $top\ j$ sentences retrieved during the process and used for majority voting. All scores are macro averaged.

focus on enhancing precision instead of recall.

Number of retrieved sentences. Instead of taking the whole documents and sending them with the question to the reader, the second group of experiments selected only the $top\ j$ most relevant sentences within all abstracts and used those as evidence. In this setup, we first retrieve the top 20 most similar abstracts with BM25. Afterward, all abstracts are split into sentences, which are embedded with the sentence-embedding model SPICED. After that, the $top\ j$ most similar sentences to the question q , according to cosine similarity, are chosen from the pool of sentences (so multiple sentences can come from the same abstract). The QA module calculates the entailment probability between the question and each selected sentence, and finally, majority voting is performed.

Year of publication and number of citations. The third and fourth group of experiments focused less on the technical parameters of the retrieval but more on the *quality* of the discovered evidence. So far, not many studies have leveraged the metadata of retrieved evidence documents for enhancing medical and health-related question answering or fact verification. We use two metadata parameters that should intuitively have an influence on the quality of the performance – *year of publication* of the retrieved research publication and the *number of citations* it has. The year was provided among the metadata that comes with PubMed, but getting the number of citations was more challenging, considering it is not found in the MEDLINE dump. Therefore, we utilized the *Semantic Scholar API* (Ammar et al., 2018) by querying it with the PubMed ID (PMID) of the retrieved article and then calling the API to get the number of citations. Once we had both numbers, experiments consisted of filtering the pool of possible evidence documents

by posing a restriction on the minimum year of publication and the minimum number of citations. Out of the top k documents we retrieved, only those published after a certain year, or those with at least a certain number of citations, were selected as the final evidence documents. The rest of the workflow is the same as in the first experiment: the documents are passed to the QA module and the answer is predicted.

Given that we work with discrete answers and labels, the evaluation metrics we used are macro-averaged versions of classification metrics precision, recall, and F1 score. Macro averaging implies that the arithmetic mean of the metric for each individual class is taken (e.g., $F1_{macro} = \sum_i F1_{class\ i}/n$). For HealthFC-3, this is an average of three classes, while for the other datasets, it is an average of two classes. The motivation behind using the macro version is that it considers all classes equally important. We posit that in a deployed health QA system, users would be interested not only in detecting positive answers, but the system effectively discerning between both negative and positive answers to questions. This also follows the literature on automated fact-checking, which commonly uses macro-averaged scores (Bekoulis et al., 2021).

All the experiments were run on a single Nvidia V100 GPU with 16 GB of VRAM. The process of retrieving the top 100 most relevant documents for each dataset used one computation hour. The process of predicting the final answer with the top 100 most relevant documents also used one computation hour.

5 Results

In this section, we present and describe the results of the conducted experiments. Table 2 shows the final classification scores when changing the number

Year	Precision	Recall	F1 Score
≥ 2020	59.7	60.3	58.7
≥ 2018	59.6	58.0	57.9
≥ 2015	61.1	56.0	53.9
≥ 2010	63.4	55.6	52.8
≥ 2005	68.1	56.5	52.0
≥ 2000	66.1	56.8	51.8
≥ 1990	65.6	55.4	51.3
≥ 1980	64.2	54.7	50.0

Table 4: Results of final answer prediction over HealthFC-2, with different limitations on the earliest year of the retrieved evidence documents. All scores are macro averaged.

of documents retrieved. Similarly, Table 3 shows the final performance for different numbers of top sentences extracted. Tables 4 and 5 show the influence of filtering evidence documents based on their year of publication and number of citations.

5.1 Retrieved Documents and Sentences

An interesting trend is observed when looking at Table 2. For all four datasets, the worst performance was when retrieving the highest amount of documents (50 and 100) and slowly increased towards the lower values. In the case of HealthFC-2 and TREC-Health, taking just the top document gave the best value of F1 (although the best macro precision and recall were with 5 and 10 documents). For BioASQ, taking into account just the top document is considerably worse than all other settings (because of poor recall), but the overall trend also holds for this dataset. As expected, the most challenging is the ternary version of HealthFC, but even there, the F1 performance increase of +2.5% is observed from 100 to 1 document. For binary HealthFC and TREC-Health, jumps from 100 to 1 of +18% and +10% are seen.

When looking at Table 3, the results are less consistent than in the previous case. In fact, for datasets HealthFC-3 and BioASQ, the effect is rather opposite to the one in the previous experiment. The performance kept dropping as the number of selected sentences was lowered. The increased amount of knowledge in the bigger corpus of sentences helped the performance. For the binary HealthFC and TREC-Health, the numbers generally increased towards the lower number of sentences retrieved, but there wasn't a consistent pattern. Overall, the experiment showed that, in

general, adding sentences that are relevant and semantically similar to the question increases the performance as opposed to adding more full documents to the QA system.

# Cits.	Precision	Recall	F1 Score
≥ 100	59.7	58.9	59.1
≥ 75	59.6	57.9	58.0
≥ 50	65.2	58.5	56.9
≥ 25	65.0	57.8	55.8
≥ 10	57.6	55.6	54.8
≥ 5	67.0	56.9	53.6
≥ 0	66.4	56.5	53.1

Table 5: Results of final answer prediction over HealthFC-2, with different limitations on the minimum number of citations of the retrieved evidence documents. All scores are macro averaged.

5.2 Evidence Quality

Table 4 tested the influence of the recency of the paper (year of publication) on the predictive performance. Once again, an interesting trend can be noted. The more recent the selected documents were, the more accurate the answers to health questions in our system. A similar phenomenon can be observed in Table 5 for the number of citations of the papers. When limiting the selection to more and more cited papers, the final score kept increasing. Nevertheless, the highest scores in Tables 3 and 4 are still lower than the top-1-document performance from Table 4.

It intuitively makes sense that the more recent papers will provide the latest knowledge and insights into a research hypothesis, which then also better aligns with the gold labels of our datasets. This is also slightly biased by the recent creation date of our datasets, where annotators had access to the most recent knowledge. Likewise, the better the reputation of a paper (more citations), it could

		average	median	std.dev.
Year	✓	2012.6	2014	8.1
	×	2009.3	2011	9.8
# Cits.	✓	73.7	30	171.9
	×	63.6	26	122.2

Table 6: Average value, median, and standard deviation of years and citations of top 10 evidence documents retrieved for questions from HealthFC-3. They are grouped by those documents that yielded correct predictions (✓) and those with incorrect predictions (×).

Question	Document #1	Document #2	Document #3
Is intense physical activity associated with longevity? (Supported)	Evidence, mainly from cross-sectional studies, suggests that physical activity is a potentially important modifiable factor associated with physical performance and strength in older age. It is unclear whether the benefits of physical activity accumulate across life or whether there are sensitive periods when physical activity is more influential. (Cooper et al., 2011) (Not Enough Info)	Physical activity plays an important role for achieving healthy aging by promoting independence and increasing the quality of life. (...) Indeed, there is evidence to suggest that increasing exercise intensity in older adults may be associated with greater reductions in the risk of cardiovascular disease and mortality. (El Hajj Boutros et al., 2019) (Supported)	Exercise training above the public health recommendations provides additional benefits regarding disease protection and longevity. Endurance exercise, including high-intensity training to improve cardiorespiratory fitness promotes longevity and slows down aging. Strength training should be added to slow down loss of muscle mass, associated with aging and disease. (Pedersen, 2019) (Supported)
Is dexamethasone recommended for treatment of intracerebral hemorrhage? (Refuted)	Dexamethasone contributed to many serious adverse events. Conclusions: Given the small sample size, these preliminary results have not shown a clear beneficial effect of dexamethasone against placebo in our patients. (Prud'homme et al., 2016) (Refuted)	Overall, there is no evidence of a beneficial or adverse effect of corticosteroids in patients with either SAH or PICH. Confidence intervals are wide and include clinically significant effects in both directions. (Feigin et al., 2005) (Not Enough Info)	Dexamethasone has been used to treat cerebral edema associated with brain abscess. (...) Conclusions: In patients with a brain abscess treated with antibiotics, the use of dexamethasone was not associated with increased mortality. (Simjian et al., 2018) (Supported)

Table 7: Example questions and retrieved evidence from the top three documents. In the first case, retrieving the top-3 performs better than retrieving just the top-1 document. In the second case, retrieving just the single best document gives a more precise verdict.

be assumed it will be a more reliable indicator of a correct answer. This wasn't the case for all of the questions, and there were many examples where older or less cited papers aligned better with the gold labels. Still, Table 6 provides some statistics that show that the trend generally holds. On average, those documents that produced a correct verdict were around three years more recent (in both the mean and median) while having an average of 10 citations more (median four citations). On the other hand, all categories have considerable standard deviations, showing many outliers and exceptions to this rule.

Another challenging factor is that there could seemingly be an inverse correlation between the age of a paper and its number of citations – older publications have had more time to gather a bigger number of citations. After deeper inspection, we observed that being cited a lot over time is only true for high-quality studies. Overall, striking a balance by finding both those studies that are both recent and already highly cited is an optimal strategy.

6 Discussion

In this section, we discuss and provide deeper insights uncovered in the results.

6.1 Qualitative Analysis

Specific effects of retrieving multiple documents for open-domain health question answering are shown in Table 7. For the first question, just retrieving the best document would have led to a study that does not provide a definitive conclusion to the question. However, the second and third most similar documents that were retrieved support

the given research hypothesis. This held true even for further documents that cannot be shown in the table because of space constraints. On the other hand, the second question would have been more appropriately assessed by just looking at the best document instead of the top 3 documents. In fact, the second and third documents do not explicitly talk about the given question but rather a variation.

Examples of positive influences of taking into account the qualitative properties of the evidence into account, namely the year of publication and number of citations, are shown in Table 8. For the first question, a publication with above 100 citations directly answers the question and matches the gold annotation for this claim from the dataset. On the other hand, the second most similar retrieved document was a study with only three citations, which seems to be inconclusive about or even refuting the research hypothesis. Similarly, the answer to the second question on the effects of ginkgo is skewed by the top document from the 1990s that talks about the presumed positive effects of ginkgo. Two decades later, another meta-review analysis of systematic reviews on ginkgo biloba showed that it was not proven to help with tinnitus. This nicely demonstrates the changing nature of scientific knowledge and scientific consensus throughout time.

6.2 Future Directions

Based on our findings and discussion, we see that future work could focus on these directions:

- **Strength of evidence.** Taking into account the year of publication and number of citations has proven to be an effective strategy for en-

Question	Document #1	Document #2
Can stress promote dementia? (Supported)	To test the hypothesis that high job stress during working life might lead to an increased risk of dementia and Alzheimer’s disease (AD) in late life. (...) Lifelong work-related psychosocial stress, characterized by low job control and high job strain, was associated with increased risk of dementia and AD in late life, independent of other known risk factors. (Wang et al., 2020) (Supported) [121 citations]	Patients with Alzheimer’s disease (AD) or dementia are increasing in numbers as the population worldwide ages. Mid-life psychological stress, psychosocial stress and post-traumatic stress disorder have been shown to cause cognitive dysfunction. The mechanisms behind stress-induced AD or dementia are not known. (Zhu et al., 2021a) (Refuted) [3 citations]
Can ginkgo biloba relieve the symptoms of tinnitus? (Refuted)	Ginkgo biloba is a plant extract used to alleviate symptoms associated with cognitive deficits, e.g., decreased memory performance, lack of concentration, decreased alertness, tinnitus, and dizziness. Pharmacologic studies have shown that the therapeutic effect of ginkgo(...) (Søholm, 1998) (Supported) [Year: 1998]	We identified three systematic reviews including four primary studies, all corresponding to randomized trials. We concluded the use of Ginkgo biloba probably does not decrease the severity of tinnitus. In addition, it does not reduce the intensity of tinnitus or improve the quality of life of patients. (Kramer and Ortigoza, 2018) (Refuted) [Year: 2018]

Table 8: Example questions and retrieved evidence from two different documents, where only one of them provided a correct final answer. The more recent and the more cited papers provided better performance.

hancing the health question-answering performance. Similarly, further metadata could be taken into account to augment the process. In medical research, the strength of evidence is an important factor, and systems like GRADE are used to assess it (Balshem et al., 2011). Different types of studies, such as a single study, a randomized clinical trial, and a systematic review, all have different strengths. Including this could be an important factor in improving the reliability of answers.

- **Evidence disagreement and variation.** We observed how different studies and sources can come to differing conclusions regarding a claim. In this paper, we chose the majority vote among the *top k* documents as the final decision, but this diminishes the information about the prediction uncertainty. This is part of the broader ML problem of learning with disagreements (Leonardelli et al., 2023) and modeling human label variation (Baan et al., 2024). While usually applied to uncertainty in data annotation, it could also be applied in the future to uncertainty in answering questions with diverse evidence documents.
- **Interpretability and user-centric results.** Other than just predicting the final answer, the end users posing biomedical and health questions would appreciate making the results more interpretable. This includes aspects such as displaying the different evidence documents, highlighting important sections,

and showing prediction probabilities. Modern large language models (LLMs) could be used to enhance the reasoning process and to generate user-friendly explanations of model predictions and decisions (Lamm et al., 2021).

7 Conclusion

In this paper, we conducted a number of experiments assessing the performance of a health question-answering system in an open-domain setting. Moving away from the standard setup of working with a small evidence corpus, we expand the knowledge sources to a large corpus of more than 20 million biomedical research abstracts. We measured the answer prediction performance over three diverse datasets of health questions while varying four aspects: number of retrieved documents, number of extracted sentences, and evidence quality in the form of year of publication and number of citations. Our results show that a lower number of documents retrieved leads to better performance, with the ideal spot in the 1–5 range. We also show that the performance is improved and made more reliable by using a time-aware evidence retrieval process, i.e., retrieving only the highly cited and more recent papers. Our research leaves room for exploration of disagreeing and conflicting evidence, generating explanations for end users, and including further metadata. We hope our research will encourage more exploration of the open-domain health question answering setting and addressing real-world user needs.

Limitations

The question-answering pipeline used in this paper is a complex system with multiple factors – the choice of the retrieval method, the sentence embedding model, the QA model, and the prediction threshold. It is possible that some incorrect predictions weren't due to the change in parameters but due to the faulty answer prediction. To account for this, we kept the reader part of the pipeline constant and fixed, to ensure a comparable setup. We focused on reporting only those findings and patterns that we observed were commonly occurring after a thorough analysis of retrieved evidence for each claim. Additionally, the final answer prediction for a given question was done by taking the label from the majority vote. This discards the information about evidence disagreement and label variation, which are also crucial in this domain. Future work should focus on this challenge.

Ethical Considerations

The study focuses on the medical domain and answering health-related questions. This is a sensitive field where problems like misinformation, model hallucination, and incorrect evidence retrieval can lead to harmful consequences, disinformation spread, and societal effects. The question-answering system in this study shows promising performance but is still not ready for deployment and widespread use by end users, considering incorrect predictions, hallucinations, and lack of model interpretability.

Acknowledgements

This research has been supported by the German Federal Ministry of Education and Research (BMBF) grant 01IS17049 Software Campus 2.0 (TU München).

References

- Mustafa Abualsaud, Christina Lioma, Maria Maistro, Mark D. Smucker, Guido, and Zuccon. 2020. [Overview of the trec 2019 decision track](#).
- Liesbeth Allein, Isabelle Augenstein, and Marie-Francine Moens. 2021. Time-aware evidence ranking for fact-checking. *Journal of Web Semantics*, 71:100663.
- Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Miles Crawford, Doug Downey, Jason Dunkelberger, Ahmed Elgohary, Sergey Feldman, Vu Ha, Rodney Kinney, Sebastian Kohlmeier,

Kyle Lo, Tyler Murray, Hsu-Han Ooi, Matthew Peters, Joanna Power, Sam Skjonsberg, Lucy Lu Wang, Chris Wilhelm, Zheng Yuan, Madeleine van Zuylen, and Oren Etzioni. 2018. [Construction of the literature graph in semantic scholar](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers)*, pages 84–91, New Orleans - Louisiana. Association for Computational Linguistics.

Joris Baan, Raquel Fernández, Barbara Plank, and Wilker Aziz. 2024. [Interpreting predictive probabilities: Model confidence or human label variation?](#) In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 268–277, St. Julian's, Malta. Association for Computational Linguistics.

Howard Balshem, Mark Helfand, Holger J Schünemann, Andrew D Oxman, Regina Kunz, Jan Brozek, Gunn E Vist, Yngve Falck-Ytter, Joerg Meerpohl, Susan Norris, et al. 2011. Grade guidelines: 3. rating the quality of evidence. *Journal of clinical epidemiology*, 64(4):401–406.

Razieh Baradaran, Razieh Ghiasi, and Hossein Amirkhani. 2022. A survey on machine reading comprehension systems. *Natural Language Engineering*, 28(6):683–732.

Gopi Battineni, Simone Baldoni, Nalini Chintalapudi, Getu Gamo Sagaro, Graziano Pallotta, Giulio Nittari, and Francesco Amenta. 2020. Factors affecting the quality and reliability of online health information. *Digital health*, 6:2055207620948996.

Giannis Bekoulis, Christina Papagiannopoulou, and Nikos Deligiannis. 2021. A review on fact extraction and verification. *ACM Computing Surveys (CSUR)*, 55(1):1–35.

Leona Cilar Budler, Lucija Gosak, and Gregor Stiglic. 2023. Review of artificial intelligence-based question-answering systems in healthcare. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 13(2):e1487.

Kathi Canese and Sarah Weis. 2013. Pubmed: the bibliographic database. *The NCBI handbook*, 2(1).

Danqi Chen and Wen-tau Yih. 2020a. Open-domain question answering. In *Proceedings of the 58th annual meeting of the association for computational linguistics: tutorial abstracts*, pages 34–37.

Danqi Chen and Wen-tau Yih. 2020b. [Open-domain question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, pages 34–37, Online. Association for Computational Linguistics.

Charles L. A. Clarke, Saira Rizvi, Mark D. Smucker, Maria Maistro, and G. Zuccon. 2021. [Overview of the trec 2021 health misinformation track](#). In *Text Retrieval Conference*.

- Rachel Cooper, Gita D Mishra, and Diana Kuh. 2011. Physical activity across adulthood and physical performance in midlife: findings from a british birth cohort. *Am. J. Prev. Med.*, 41(4):376–384.
- Florin Cuconasu, Giovanni Trappolini, Federico Siciliano, Simone Filice, Cesare Campagnano, Yoelle Maarek, Nicola Tonellotto, and Fabrizio Silvestri. 2024. The power of noise: Redefining retrieval for rag systems. *arXiv preprint arXiv:2401.14887*.
- Dina Demner-Fushman, Yassine Mrabet, and Asma Ben Abacha. 2020. Consumer health information and question answering: helping consumers find answers to their health-related information needs. *Journal of the American Medical Informatics Association*, 27(2):194–201.
- Guy El Hajj Boutros, José A Morais, and Antony D Karelis. 2019. Current concepts in healthy aging and physical activity: A viewpoint. *J. Aging Phys. Act.*, 27(5):755–761.
- V L Feigin, N Anderson, G J E Rinkel, A Algra, J van Gijn, and D A Bennett. 2005. Corticosteroids for aneurysmal subarachnoid haemorrhage and primary intracerebral haemorrhage. *Cochrane Database Syst. Rev.*, (3):CD004583.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. **Deberta: Decoding-enhanced bert with disentangled attention**. In *International Conference on Learning Representations*.
- Xiaoyun Jia, Yan Pang, and Liangni Sally Liu. 2021. **Online health information seeking behavior: A systematic review**. *Healthcare*, 9(12).
- Di Jin, Eileen Pan, Nassim Oufattole, Wei-Hung Weng, Hanyi Fang, and Peter Szolovits. 2021. What disease does this patient have? a large-scale open domain question answering dataset from medical exams. *Applied Sciences*, 11(14):6421.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. **PubMedQA: A dataset for biomedical research question answering**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577, Hong Kong, China. Association for Computational Linguistics.
- Qiao Jin, Zheng Yuan, Guangzhi Xiong, Qianlan Yu, Huaiyuan Ying, Chuanqi Tan, Mosha Chen, Songfang Huang, Xiaozhong Liu, and Sheng Yu. 2022. Biomedical question answering: a survey of approaches and challenges. *ACM Computing Surveys (CSUR)*, 55(2):1–36.
- Felipe Kramer and Ángela Ortigoza. 2018. Ginkgo biloba for the treatment of tinnitus. *Medwave*, 18(6):e7295.
- Matthew Lamm, Jennimaria Palomaki, Chris Alberti, Daniel Andor, Eunsol Choi, Livio Baldini Soares, and Michael Collins. 2021. **QED: A Framework and Dataset for Explanations in Question Answering**. *Transactions of the Association for Computational Linguistics*, 9:790–806.
- Moritz Laurer, Wouter Van Atteveldt, Andreu Casas, and Kasper Welbers. 2024. Less annotating, more classifying: Addressing the data scarcity issue of supervised machine learning with deep transfer learning and bert-nli. *Political Analysis*, 32(1):84–100.
- Yoonjoo Lee, Kyungjae Lee, Sunghyun Park, Dasol Hwang, Jaehyeon Kim, Hong-in Lee, and Moontae Lee. 2023. Qasa: advanced question answering on scientific articles. In *International Conference on Machine Learning*, pages 19036–19052. PMLR.
- Elisa Leonardelli, Gavin Abercrombie, Dina Almanea, Valerio Basile, Tommaso Fornaciari, Barbara Plank, Verena Rieser, Alexandra Uma, and Massimo Poesio. 2023. **SemEval-2023 task 11: Learning with disagreements (LeWiDi)**. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2304–2318, Toronto, Canada. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521.
- Stephen Neely, Christina Eldredge, and Ron Sanders. 2021. Health information seeking behaviors on social media during the covid-19 pandemic among american social networking site users: survey study. *Journal of medical Internet research*, 23(6):e29802.
- Anastasios Nentidis, Konstantinos Bougiatiotis, Anastasia Krithara, and Georgios Paliouras. 2020. Results of the seventh edition of the bioasq challenge. In *Machine Learning and Knowledge Discovery in Databases*, pages 553–568, Cham. Springer International Publishing.
- Philhoon Oh and James Thorne. 2023. **Detrimental contexts in open-domain question answering**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11589–11605, Singapore. Association for Computational Linguistics.

- Bente Klarlund Pedersen. 2019. Which type of exercise keeps you young? *Curr. Opin. Clin. Nutr. Metab. Care*, 22(2):167–173.
- Michel Prud’homme, François Mathieu, Nicolas Marcotte, and Sylvine Cottin. 2016. A pilot placebo controlled randomized trial of dexamethasone for chronic subdural hematoma. *Can. J. Neurol. Sci.*, 43(2):284–290.
- Alexander Pugachev, Ekaterina Artemova, Alexander Bondarenko, and Pavel Braslavski. 2023. Consumer health question answering using off-the-shelf components. In *European Conference on Information Retrieval*, pages 571–579. Springer.
- Anna Rogers, Matt Gardner, and Isabelle Augenstein. 2023. Qa dataset explosion: A taxonomy of nlp resources for question answering and reading comprehension. *ACM Computing Surveys*, 55(10):1–45.
- Lila J. Finney Rutten, Kelly D. Blake, Alexandra J. Greenberg-Worisek, Summer V. Allen, Richard P. Moser, and Bradford W. Hesse. 2019. Online health information seeking among us adults: Measuring progress toward a healthy people 2020 objective. *Public Health Reports*, 134(6):617–625. PMID: 31513756.
- Artsiom Sauchuk, James Thorne, Alon Halevy, Nicola Tonello, and Fabrizio Silvestri. 2022. On the role of relevance in natural language processing tasks. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1785–1789.
- Michael Schlichtkrull, Zhijiang Guo, and Andreas Vlachos. 2024. Averitec: A dataset for real-world claim verification with evidence from the web. *Advances in Neural Information Processing Systems*, 36.
- Thomas Simjian, Ivo S Muskens, Nayan Lamba, Ismael Yunusa, Kristine Wong, Raymond Veronneau, Annick Kronenburg, H Bart Brouwers, Timothy R Smith, Rania A Mekary, and Marika L D Broekman. 2018. Dexamethasone administration and mortality in patients with brain abscess: A systematic review and meta-analysis. *World Neurosurg.*, 115:257–263.
- B Sjøholm. 1998. Clinical improvement of memory and other cognitive functions by ginkgo biloba: review of relevant literature. *Adv. Ther.*, 15(1):54–65.
- Dominik Stambach, Boya Zhang, and Elliott Ash. 2023. The choice of textual knowledge base in automated claim checking. *ACM Journal of Data and Information Quality*, 15(1):1–22.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and VERification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- David Vilares and Carlos Gómez-Rodríguez. 2019. HEAD-QA: A healthcare dataset for complex reasoning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 960–966, Florence, Italy. Association for Computational Linguistics.
- Juraj Vladika and Florian Matthes. 2023a. Scientific fact-checking: A survey of resources and approaches. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6215–6230, Toronto, Canada. Association for Computational Linguistics.
- Juraj Vladika and Florian Matthes. 2023b. Sebis at SemEval-2023 task 7: A joint system for natural language inference and evidence retrieval from clinical trial reports. In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 1863–1870, Toronto, Canada. Association for Computational Linguistics.
- Juraj Vladika and Florian Matthes. 2024. Comparing knowledge sources for open-domain scientific claim verification. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2103–2114, St. Julian’s, Malta. Association for Computational Linguistics.
- Juraj Vladika, Phillip Schneider, and Florian Matthes. 2023. Healthfc: A dataset of health claims for evidence-based medical fact-checking.
- Ellen M Voorhees, Donna K Harman, et al. 2005. *TREC: Experiment and evaluation in information retrieval*, volume 63. Citeseer.
- David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Iz Beltagy, Lucy Lu Wang, and Hannaneh Hajishirzi. 2022. SciFact-open: Towards open-domain scientific claim verification. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4719–4734, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Hui-Xin Wang, Maria Wahlberg, Anita Karp, Bengt Winblad, and Laura Fratiglioni. 2020. Psychosocial stress at work is associated with increased dementia risk in late life. *Alzheimers. Dement.*, 8(2):114–120.
- Dustin Wright, Jiaxin Pei, David Jurgens, and Isabelle Augenstein. 2022. Modeling information change in science communication with semantically matched paraphrases. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1783–1807, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Donghui Zhu, Axel Montagne, and Zhen Zhao. 2021a. Alzheimer’s pathogenic mechanisms and underlying sex difference. *Cell. Mol. Life Sci.*, 78(11):4907–4920.
- Fengbin Zhu, Wenqiang Lei, Chao Wang, Jianming Zheng, Soujanya Poria, and Tat-Seng Chua. 2021b. Retrieving and reading: A comprehensive survey on open-domain question answering.