

# Combining Evidence and Reasoning for Biomedical Fact-Checking

Mariano Barone  
mariano.barone@studenti.unina.it  
University of Naples Federico II  
Naples, Italy, Italy

Antonio Romano  
antonio.romano5@unina.it  
University of Naples Federico II  
Naples, Italy, Italy

Giuseppe Riccio  
giuseppe.riccio3@unina.it  
University of Naples Federico II  
Naples, Italy, Italy

Marco Postiglione  
marco.postiglione@northwestern.edu  
Northwestern University  
IL, United States, United States

Vincenzo Moscato  
vincenzo.moscato@unina.it  
University of Naples Federico II  
Naples, Italy, Italy

## Abstract

Misinformation in healthcare, from vaccine hesitancy to unproven treatments, poses risks to public health and trust in medical systems. While machine learning and natural language processing have advanced automated fact-checking, validating biomedical claims remains uniquely challenging due to complex terminology, the need for domain expertise, and the critical importance of grounding in scientific evidence. We introduce CER (Combining Evidence and Reasoning), a novel framework for biomedical fact-checking that integrates scientific evidence retrieval, reasoning via large language models, and supervised veracity prediction. By integrating the text-generation capabilities of large language models with advanced retrieval techniques for high-quality biomedical scientific evidence, CER effectively mitigates the risk of hallucinations, ensuring that generated outputs are grounded in verifiable, evidence-based sources. Evaluations on expert-annotated datasets (HealthFC, BioASQ-7b, SciFact) demonstrate state-of-the-art performance and promising cross-dataset generalization. Code and data are released for transparency and reproducibility: <https://github.com/PRAISELab-PicusLab/CER>.

## CCS Concepts

- **Computing methodologies** → **Natural language generation;**
- **Applied computing** → **Health care information systems.**

## Keywords

Fact-Checking, Healthcare, Generative AI, Large Language Models

### ACM Reference Format:

Mariano Barone, Antonio Romano, Giuseppe Riccio, Marco Postiglione, and Vincenzo Moscato. 2025. Combining Evidence and Reasoning for Biomedical Fact-Checking. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)*, July 13–18, 2025, Padua, Italy. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3726302.3729931>



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

*SIGIR '25, Padua, Italy*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1592-1/2025/07

<https://doi.org/10.1145/3726302.3729931>

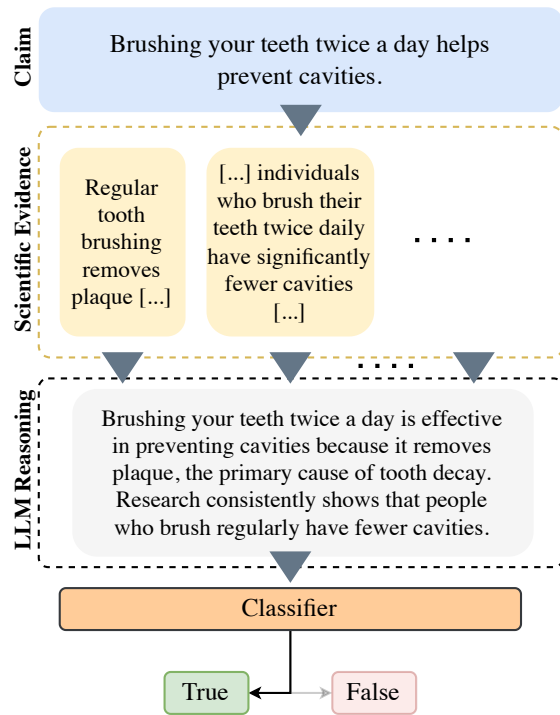
## 1 Introduction

Distinguishing reliable information from misinformation represents a *global challenge* that threatens the foundations of informed decision-making. Misinformation not only shapes public opinion and policy outcomes but also systematically erodes trust in authoritative information sources, thereby creating a recursive cycle that further complicates the pursuit of accurate knowledge [6]. The healthcare domain exemplifies the devastating impact of misinformation on public health outcomes. Misinformation surrounding evidence-based medical interventions — including cancer treatments [25], birth control<sup>1</sup>, and other forms of medical care — has led to devastating public health consequences, including treatment delays, increased mortality rates, and erosion of trust in healthcare systems [14]. For example, during the COVID-19 pandemic, misinformation about vaccines led to hesitancy, resulting in lower vaccination rates and increased spread of the virus. This misinformation was often propagated through social media platforms, where false claims about vaccine safety and efficacy spread rapidly [3]. Additionally, misinformation about unproven treatments, such as the use of hydroxychloroquine, caused confusion and sometimes harmful self-medication practices [38].

Machine learning (ML) and natural language processing (NLP) techniques have emerged as powerful tools for automated fact-checking, demonstrating promising results across various domains [40, 42]. While these computational approaches have been extensively studied and deployed in political fact-checking [35, 36, 47], their effectiveness in validating biomedical claims remains limited. Biomedical fact-checking presents distinct challenges that transcend the capabilities of current automated systems. Unlike many other domains, medical claims often incorporate domain-specific terminology, rely on complex causal relationships, and require sophisticated interpretation of empirical evidence [55]. *Furthermore, health-related misinformation demands particularly rigorous verification due to its potential impact on public health outcomes, necessitating both specialized domain knowledge and advanced reasoning capabilities that can contextualize claims within the broader framework of scientific evidence.*

Large language models (LLMs) have recently emerged as a promising approach to address these challenges in biomedical fact-checking.

<sup>1</sup><https://www.washingtonpost.com/health/2024/03/21/stopping-birth-control-misinformation/>



**Figure 1: Combining Evidence and Reasoning (CER) for biomedical fact-checking.** (1) Evidence supporting or opposing the biomedical claim is retrieved from scientific literature. (2) A large language model (LLM) is leveraged to justify the claim in light of the scientific evidence. (3) A classifier is trained to assess the veracity of the claim by leveraging the information collected in steps (1) and (2).

Their ability to process and reason about complex textual information, combined with their broad knowledge acquisition during pre-training, makes them potentially valuable tools for analyzing biomedical claims [31]. LLMs have demonstrated capabilities in understanding technical medical terminology, identifying relevant scientific evidence, and providing nuanced interpretations of research findings [43]. Retrieval-augmented generation (RAG) frameworks further enhance these capabilities by grounding LLM outputs in verified scientific literature, potentially mitigating the risk of hallucinations through direct access to peer-reviewed evidence [22, 28]. *By dynamically retrieving and incorporating relevant medical research during the analysis process, RAG systems could provide more reliable and verifiable fact-checking results.* However, these approaches still face important limitations that must be carefully considered: (1) LLMs, even when augmented with retrieval, can still generate plausible-sounding but factually incorrect information. (2) The quality of fact-checking remains heavily dependent on the relevance and reliability of the retrieved evidence, and challenges persist in efficiently identifying and incorporating the most pertinent scientific literature. (3) Additionally, their knowledge cutoff dates and potential biases in training data can lead to outdated or incomplete medical information.

In this work, we propose a novel approach to biomedical fact-checking that integrates three key components: systematic scientific evidence retrieval, LLM reasoning capabilities, and supervised veracity prediction. Our system, CER (Combining Evidence and Reasoning), addresses the challenges in biomedical claim verification through a multi-stage architecture (illustrated in Figure 1) that synthesizes (1) the input claim, (2) relevant scientific evidence from peer-reviewed literature, and (3) LLM-generated justifications. This architecture leverages the advanced reasoning capabilities of LLMs while maintaining factual grounding through explicit evidence linkage, thereby mitigating the risk of hallucinations. The proposed architecture follows an open-domain fact-checking paradigm, which eliminates reliance on a predefined set of evidence [1]. Rather than being restricted to a fixed corpus of golden evidence, this approach dynamically retrieves relevant information from an extensive external knowledge base, enabling a *content-driven* evaluation of claims. We evaluate CER on three benchmarks annotated by domain experts: HealthFC [45], BioASQ-7b [34], and SciFact [49]. The results show that CER consistently outperforms state-of-the-art methods, achieving up to a +3.43% improvement in F1 scores. To provide deeper insights, we analyze the performance impact of different retrieval strategies (sparse vs. dense) and examine the contributions of scientific evidence integration and LLM-based reasoning. Additionally, cross-dataset evaluations highlight the generalizability and robustness of our framework across diverse settings. While these results underscore its effectiveness in healthcare, the modular design of CER makes it adaptable to other domains requiring domain-specific evidence and advanced reasoning. To promote transparency and reproducibility, we release the code, evidence, and outputs generated during our experiments<sup>2</sup>.

## 2 Combining Evidence and Reasoning (CER)

In this section, we present our proposed methodology to classify biomedical *claims*<sup>3</sup>. Our architecture, as illustrated in Figure 2, is divided into three main components: (1) Scientific Evidence Retrieval, (2) LLM Reasoning and (3) Veracity Prediction. The remainder of this section provides an in-depth description of each component.

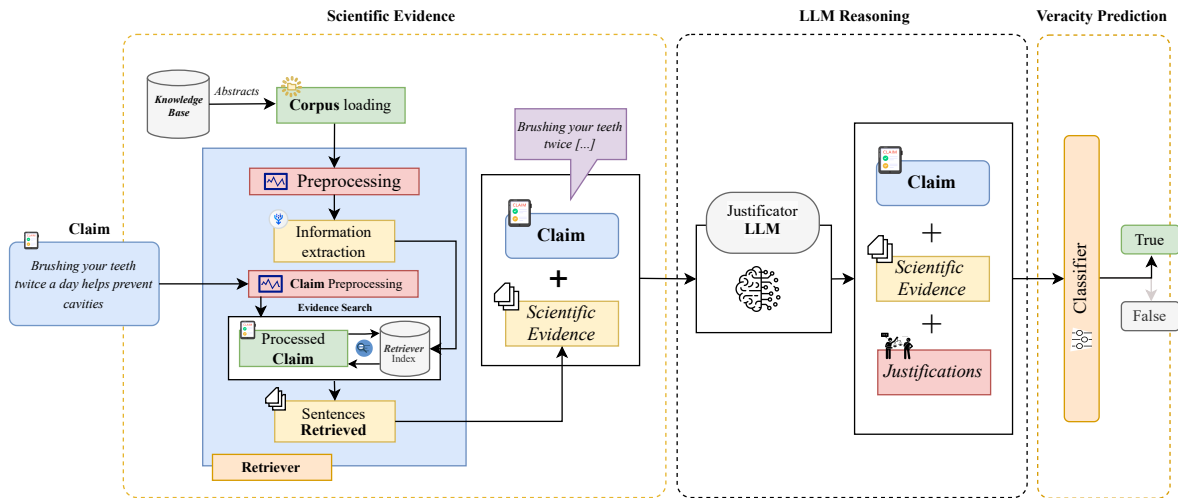
### 2.1 Scientific Evidence Retrieval

The workflow starts with the Scientific Evidence Retrieval module, which interfaces with scientific knowledge bases to extract domain-specific claims. Similarly to previous studies [24, 39, 44], we leverage PubMed<sup>4</sup> as our primary knowledge source, given its comprehensive repository of peer-reviewed biomedical literature. We specifically focus on article abstracts, which offer concise yet comprehensive summaries of research findings, methodologies, and conclusions. This choice is motivated by both the information density of abstracts and the common restrictions on access to full-text articles, while providing sufficient context for validating medical claims.

<sup>2</sup>Code is shared anonymously; data will be available upon publication.

<sup>3</sup>In this work, a *claim* refers to any statement or question whose truthfulness or factual accuracy can be evaluated based on external evidence. This broad definition encompasses declarative sentences as well as interrogative forms.

<sup>4</sup><https://zenodo.org/records/7849020>



**Figure 2: CER workflow.** Claims are preprocessed before being matched against a scientific corpus indexed using a retrieval system. Relevant evidence passages are retrieved and concatenated with the original claim as input to a large language model (LLM), which generates detailed justifications. A binary classifier then evaluates the claim’s veracity by analyzing the claim text, retrieved scientific evidence, and generated justifications to produce a final true/false assessment

*Retrieval Module.* To extract sentences from the PubMed knowledge base, we employed Sparse Retrieval and Dense Retrieval strategies, both widely used in open-domain question-answering systems. We utilized a BM25 inverted index [37] to rank documents based on term-matching relevance. The BM25 scoring function computes the relevance of a document  $D$  to a query  $Q$  as shown as follows:

$$\text{score}(Q, D) = \sum_{t \in Q} \text{IDF}(t) \cdot \frac{f(t, D) \cdot (k_1 + 1)}{f(t, D) + k_1 \cdot (1 - b + b \cdot \frac{|D|}{\text{avgdl}})},$$

where:

- $\text{IDF}(t)$  is the inverse document frequency, calculated as:

$$\text{IDF}(t) = \log \frac{N - n_t + 0.5}{n_t + 0.5},$$

with  $N$  as the total number of documents and  $n_t$  as the number of documents containing the term  $t$

- $f(t, D)$  is the term frequency of  $t$  in document  $D$
- $|D|$  is the length of the document  $D$ , and  $\text{avgdl}$  is the average document length in the collection
- $k_1$  and  $b$  are hyperparameters controlling term frequency saturation and document length normalization, respectively

Text preprocessing included normalization, tokenization, stemming, stopword removal, and punctuation standardization. The top  $k$  results ( $k = 20$ ) were ranked using cosine similarity over term-based features, offering efficiency and interpretability for term-centric queries. For Dense Retrieval, we employed the multi-qa-MiniLM-L6-cos-v1<sup>5</sup>, a pre-trained SBERT model optimized for semantic similarity tasks. This model generates dense vector embeddings for textual inputs, capturing semantic nuances such as synonyms and paraphrases. The similarity between a query  $Q$  and a document  $D$

<sup>5</sup><https://huggingface.co/sentence-transformers/multi-qa-MiniLM-L6-cos-v1>

is measured using cosine similarity:  $\text{sim}(Q, D) = \frac{\mathbf{q} \cdot \mathbf{d}}{\|\mathbf{q}\| \cdot \|\mathbf{d}\|}$ , where  $\mathbf{q}$  and  $\mathbf{d}$  are the embedding vectors of the query  $Q$  and document  $D$ , respectively, and  $\|\mathbf{q}\|$  and  $\|\mathbf{d}\|$  are their Euclidean norms.

*Combining Claims and Scientific Evidence.* Once the knowledge base is indexed, claims from the dataset are processed and used as search queries. The retrieval system identifies relevant sentences from the indexed database by evaluating term-level relevance between the query and the stored text. This step ensures that the retrieved evidence is contextually aligned with the given claim. For each claim, the system extracts up to three pieces of evidence from the retrieved results. The extracted evidence is structured in the following format: "claim [SEP] evidence\_1, evidence\_2, evidence\_3." Here, [SEP] acts as a separator token to distinguish the claim from the evidence pieces, while the evidence sentences are concatenated with commas to maintain clarity and consistency in the input format. Based on preliminary experiments, using three pieces of evidence usually provides sufficient support without redundancy or excessive processing costs. The final claim-evidence pairs are stored in a structured format and prepared for the next stage of processing, where they will be used as input for reasoning by a large language model (LLM). This integration of retrieval and reasoning ensures a streamlined pipeline for claim verification tasks.

## 2.2 LLM Reasoning

In this phase, CER leverages LLMs as reasoning assistants rather than autonomous decision-makers. This design choice explicitly mitigates the risk of hallucinations inherent in LLMs when used for standalone fact-checking. The LLM’s role is twofold: (1) to assess the claim’s *veracity* based on the provided scientific evidence and (2) to generate a detailed *justification* for this assessment. Both

<b>Dataset</b>	SciFact
<b>Claim</b>	A deficiency of folate increases blood levels of homocysteine.
<b>Justification</b>	<i>Folate deficiency has been shown to increase blood levels of homocysteine. This is due to the role of folate in the metabolism of homocysteine, where it acts as a coenzyme. In cases of folate deficiency, the remethylation pathway of homocysteine is inhibited, leading to its efflux into the circulation. This results in elevated homocysteine levels, which have been implicated in various adverse health conditions. Studies have reported a significant inverse correlation between plasma folate levels and circulating homocysteine levels, and a significant increase in plasma homocysteine concentrations in patients with marginal folate deficiency. Folate supplementation has been shown to reduce plasma homocysteine levels, even in patients with normal folate status and those with cobalamin deficiency.</i>

**Table 1: Example of LLM reasoning output for a scientific claim, showing the generated justification based on retrieved evidence. The claim originates from the SciFact dataset [49], a benchmark for scientific claim verification.**

the binary prediction and the accompanying reasoning serve as structured inputs for the downstream Veracity Prediction module, which will determine the final classification. The LLM’s analysis is guided by a prompt template that systematically combines the claim with the retrieved evidence (Table 1 shows an example; additional examples are provided in Appendix A.3), as shown below:

#### LLM Reasoning Prompt Structure

**Claim:** [claim].  
**Evidence:** [evidence 1, evidence 2, evidence 3].  
*Provide a justification for your assessment about the veracity of the claim.*

## 2.3 Veracity Prediction

While the LLM reasoning module generates structured analyses, relying solely on these outputs for final veracity decisions would be problematic due to potential reasoning gaps, inconsistencies between justifications and conclusions, or biases in the LLM’s decision process. To address these limitations, we implement a dedicated verification layer that evaluates both the LLM’s reasoning and the underlying evidence to produce more reliable classifications. Building upon established approaches in veracity prediction [48], we implement a classification framework that assigns one of three labels to each claim: “true”, “false”, or “insufficient evidence”. We explore and compare two distinct methodological approaches — i.e., zero-shot classification and fine-tuning — for the veracity prediction task.

*Zero-shot classification.* The first technique leverages a language model (LM) with zero-shot classification capabilities. In this approach, the LM processes claims alongside the extracted sentences and justifications, directly providing a classification output. Zero-shot classification refers to the ability of the model to generalize and perform tasks it has not been explicitly trained for. Formally, given a claim  $C$  and evidence  $E$ , the model predicts a class label  $y \in \{y_1, y_2, \dots, y_k\}$  using its pre-trained knowledge:  $P(y | C, E) = \text{LM}(C, E)$ , where  $P(y | C, E)$  represents the probability distribution over the possible labels.

This approach is particularly advantageous in scenarios where labeled data is unavailable, as it allows for rapid deployment without the need for additional task-specific training. The model utilizes its broad general knowledge acquired during pre-training to interpret the claim and evidence. However, the zero-shot approach has limitations in domain-specific tasks, where generalization may be inadequate for capturing nuanced or highly specialized information.

*Fine-tuning.* The second technique involves fine-tuning the LM, adapting it to the specific task by retraining it on a smaller, domain-specific dataset. Fine-tuning is especially effective for specialized tasks where domain-specific knowledge is critical. During this process, the LM learns to predict a target label  $y$  by minimizing the task-specific loss function  $\mathcal{L}$  over the fine-tuning dataset  $\mathcal{D}$ :

$$\mathcal{L} = - \sum_{(C, E, y) \in \mathcal{D}} \log P(y | C, E),$$

where  $\mathcal{D}$  consists of tuples  $(C, E, y)$ , with  $C$  as the claim,  $E$  as the evidence, and  $y$  as the correct label. In our case, we fine-tuned the model using claims, scientific evidence, and justifications generated in the preceding steps. This enabled the LM to learn directly from task-relevant examples, resulting in enhanced accuracy and improved fact-checking performance. Unlike the zero-shot approach, fine-tuning offers higher accuracy for domain-specific tasks but requires labeled data and additional computational resources, making it a more resource-intensive method.

## 3 Experiments


This section presents a comprehensive empirical evaluation of our proposed CER architecture across multiple datasets and experimental configurations. Our evaluation protocol encompasses four key objectives: (1) benchmarking CER against state-of-the-art closed- and open-source baselines, (2) conducting ablation studies to quantify the contribution of each pipeline component, (3) analyzing architectural variants through systematic comparison of retrieval strategies (dense vs. sparse), LLM reasoning approaches (diverse prompting techniques), and veracity prediction methods (zero-shot vs. fine-tuned classification), and (4) assessing the framework’s generalization capabilities through cross-dataset evaluation.

### 3.1 Experimental Setup

*3.1.1 Datasets.* We evaluate CER on three datasets: HealthFC [45], BioASQ-7b [34], and SciFact [49]. HealthFC [45] is a dataset for question answering and fact-checking, comprising 750 health-related claims labeled as true, false, or NEI (not enough information), verified by medical professionals using scientific evidence. BioASQ-7b

Dataset	Domain	✓	×	?
HealthFC	Health-related online searches	202	125	433
BioASQ-7b	Biomedical research	614	131	—
SciFact	Scientific claims	556	516	337

**Table 2: Overview of the datasets used in the experiments, including their domain and label distributions (✓ – true labels, × – false labels, ? – NEI labels).**



**Figure 3: Average sentence length distributions of HealthFC (a), BioASQ-7b (b), SciFact (c). Mean sentence length refers to the average number of characters per sentence in each document**

[34] contains 745 biomedical claims from the BioASQ challenge, labeled as true or false, supporting the development of question-answering systems for biomedical queries. SciFact [49] includes 1.4k expert-written scientific claims with evidence-based labels (“support”, “confute”, or NEI) and rationales, focusing on validating claims through specialized domain knowledge. These datasets were annotated by domain experts, with rigorous validation of the assigned labels to ensure accuracy and reliability. This careful curation guarantees that the datasets reflect high-quality, validated medical knowledge. Table 2 presents details of label distributions.

**3.1.2 Implementation and Training details.** To implement the Sparse and Dense Retriever, we utilized a system featuring an Intel(R) Core(TM) i7-9700K CPU @ 3.60GHz processor, 32GB of RAM, and a GT710 graphics card. All training experiments were conducted on a machine equipped with an Intel Xeon CPU with 2 vCPUs (virtual CPUs), 13GB of RAM, and an NVIDIA Tesla T4 GPU with 16GB of VRAM. The LLM used to generate justification in the “LLM Reasoning” module is *Mixtral-8x22B-Instruct-v0.1* [23]. Its Mixture-of-Experts (MoE) design excels at managing complex language structures and synthesizing information across diverse contexts, making it particularly effective for domain-specific tasks like medical text synthesis. Comparative experiments with other LLMs, including those pre-trained on biomedical corpora, are detailed in the Appendix A.2. DeBERTa (“deberta-v3-large-zeroshot-v1.1-all-33”) [20], BERT (“bert-base-uncased”) [12], and PubMedBERT [16] were tested for the implementation of the “Veracity Prediction” module. Figure 3 illustrates the sentence length distributions for HealthFC, BioASQ-7b, and SciFact, emphasizing the complexity and variability of the datasets used. To address input token limitations, we use 20 sentences per claim in the “Scientific Evidence” step and 4 sentences in the “Veracity Prediction” step, maximizing contextual information for justification creation and label prediction. Veracity Prediction models have been trained using original training and

Dataset	Train	Validation	Test
HealthFC	451	151	151
BioASQ-7b	447	150	150
SciFact	809	300	300

**Table 3: Training, validation, and test set distributions of HealthFC, BioASQ-7b, and SciFact.**

validation datasets, with hyperparameters optimized through grid and random search to ensure generalization and prevent overfitting (details in Appendix A.1). Performance was assessed on the independent test set to ensure unbiased metrics, with training and test splits detailed in Table 3. Evaluation relied on macro precision, macro recall, and macro F1 scores.

**3.1.3 Baselines.** We compared CER with open-source and closed-source fact-checkers, as shown in Table 4. Factiveverse.AI and Perplexity.AI are closed-source platforms that categorize evidence as “Supporting”, “Mixed”, or “Disputing” using sources like Google, Bing, Wikipedia, and FactSearch, with flexibility in managing the “Mixed” label. In particular, for Perplexity.AI, each claim was submitted as a direct query, and the system’s textual answer was manually mapped to three classes: responses expressing certainty and agreement were labeled as True, disagreement as False, and vague or inconclusive responses as Not Enough Information (NEI). Vladika and Matthes [44] proposes a voting-based veracity prediction architecture, leveraging sentences from knowledge bases, similar to our approach. Bekoulis et al. [8] focuses on document and sentence retrieval from Wikipedia for claim validation. Zaheer et al. [53] employs BigBird, a sparse-attention transformer, while Lan et al. [30] uses ALBERT for text classification, trained on datasets like FEVER, HOOVER, and FEVEROUS. Vladika et al. [45] applies pipeline and joint systems for medical claim fact-checking. Simple baselines (“All True”, “All False” and “All NEI”) were also included for comparison.

## 3.2 Results

**3.2.1 Comparison with state-of-the-art baselines.** In this work, we evaluated our CER framework in both zero-shot and fine-tuned configurations with state-of-the-art baselines. We used several pre-trained classification models at the foundation of CER, including general-domain models (BERT, DeBERTa-v3) and a domain-specific biomedical model (PubMedBERT). Results presented in Table 4 show that CER achieves consistent improvements across all datasets, with notable enhancements in F1 score and overall performance. Specifically, on the HealthFC dataset, fine-tuned CER achieves an F1 score of 69.90%, outperforming both its zero-shot configuration (54.61%) and the other models, emphasizing the benefits of fine-tuning on task-specific data. The results on BioASQ-7b further reinforce this trend: CER, fine-tuned, reaches a F1 score of 95.20%, surpassing all the other baselines. Notably, CER’s zero-shot performance on BioASQ-7b (F1 = 91.77%) remains competitive, indicating that it can leverage pre-existing knowledge from general-domain pretraining, which is crucial for zero-shot tasks where no additional labeled data is available. On the SciFact dataset, the fine-tuned CER model achieves an F1 score of 61.14%, while its zero-shot counterpart obtains an F1 score of 53.74%. Despite not

Category	Fact-Checker	HealthFC			BioASQ-7b			SciFact		
		Precision (%)	Recall (%)	F1 (%)	Precision (%)	Recall (%)	F1 (%)	Precision (%)	Recall (%)	F1 (%)
Baselines	All True (baseline)	8.97	33.33	14.14	82.39	1.0	90.34	13.71	33.33	19.43
	All False (baseline)	5.55	33.33	9.52	17.60	1.0	29.94	7.16	33.33	11.78
	All NEI (baseline)	18.79	33.33	24.04	–	–	–	12.47	33.33	18.15
Online platforms	Factiverse.AI [41]	53.47	30.19	38.59	91.31	84.01	87.51	50.25	46.80	48.46
	Perplexity.AI <sup>6</sup>	56.40	47.20	51.39	83.20	80.13	81.64	52.75	50.00	51.34
LLM Predictions	Mixtral LLM prediction <sup>7</sup>	40.10	41.24	40.61	72.43	70.69	71.54	45.89	40.60	43.02
	GPT4o mini [51]	40.52	39.88	39.06	53.85	52.71	52.79	39.01	37.64	38.31
Literature baselines	Vladika and Matthes [44]	52.60	44.50	40.60	60.90	64.40	61.70	46.00	42.30	44.07
	Bekoulis et al. [8]	41.72	49.36	45.21	47.42	52.36	49.76	54.12	51.20	52.62
	Zaheer et al. [53]	26.67	35.95	30.62	69.62	33.57	45.29	38.70	35.20	36.87
	Lan et al. [30]	21.40	39.36	27.72	74.60	50.37	60.13	34.45	31.85	33.10
	Vladika et al. [45]	68.24	66.84	67.53	–	–	–	–	–	–
CER ( <i>zero-shot</i> )	CER (DeBERTA v3-large)	58.26	51.48	54.61	93.10	90.48	91.77	55.72	51.90	53.74
CER ( <i>fine-tuned</i> )	CER (DeBERTA v3-large)	64.55	70.14	67.22	<b>93.70</b>	<b>96.75</b>	<b>95.20</b>	<b>61.44</b>	<b>60.72</b>	<b>61.14</b>
	CER (PubMedBERT)	67.55	<b>72.43</b>	<b>69.90</b>	89.76	90.02	89.88	57.70	56.13	56.91
	CER (BERT)	<b>69.92</b>	69.33	69.62	81.28	83.89	82.56	58.10	55.89	56.97

**Table 4: Comparison with state-of-the-art baselines. This table presents the macro precision, recall, and F1 scores for different baseline models evaluated on the experimented datasets. Best results are highlighted in bold. ‘–’ indicates instances where a technique could not be applied.**

Dataset	Scientific Evidence	P (%)	R (%)	F1 (%)
HealthFC	No	40.10	41.24	40.61
	Yes	<b>67.55</b>	<b>72.43</b>	<b>69.90</b>
BioASQ-7b	No	72.43	70.69	71.54
	Yes	<b>93.70</b>	<b>96.75</b>	<b>95.20</b>
SciFact	No	45.89	40.60	43.02
	Yes	<b>61.44</b>	<b>60.72</b>	<b>61.14</b>

**Table 5: Impact of Scientific Evidence Retrieval. This table shows the performance of CER with and without the scientific evidence retrieved from the PubMed knowledge base. P, R and F1 denote the macro precision, recall and F1 scores, respectively.**

being fine-tuned, the zero-shot version still outperforms all other baseline models. Interestingly, PubMedBERT, despite being pre-trained on biomedical data, performs slightly worse than CER on SciFact (F1 = 56.91%) and BioASQ-7b (F1 = 89.88%), suggesting that CER benefits from the combination of general-domain pretraining and fine-tuning, which may improve generalization across multiple biomedical domains. DeBERTa-v3’s strong performance on BioASQ-7b (F1 = 95.20%) and SciFact (F1 = 61.14%) is attributed to its advanced transformer architecture, which excels in handling more complex linguistic patterns, but fine-tuned CER consistently outperforms it on HealthFC and BioASQ-7b. Overall, the fine-tuned results for CER across HealthFC (F1 = 69.90%), BioASQ-7b (F1 = 95.20%), and SciFact (F1 = 61.14%) underscore its ability to generalize well across a range of datasets. The competitive zero-shot performance across all datasets further highlights the strength of CER in effectively leveraging both general-domain and biomedical pretraining for downstream tasks. These findings emphasize the importance of fine-tuning for specialized tasks while also demonstrating the utility of combining pre-trained models for robust performance in zero-shot settings.

Dataset	Retriever model	P (%)	R (%)	F1 (%)
HealthFC	Dense Retriever	64.09	71.10	67.36
	Sparse Retriever	<b>67.55</b>	<b>72.43</b>	<b>69.90</b>
BioASQ-7b	Dense Retriever	92.77	96.47	94.61
	Sparse Retriever	<b>93.70</b>	<b>96.75</b>	<b>95.20</b>
SciFact	Dense Retriever	<b>62.75</b>	<b>61.72</b>	<b>62.23</b>
	Sparse Retriever	61.44	60.72	61.14

**Table 6: Impact of Retriever Models. This table shows the performance of Dense and Sparse Retriever models on three datasets: HealthFC, BioASQ-7b, and SciFact. P, R, and F1 denote the macro precision, recall, and F1 scores, respectively.**

**3.2.2 Impact of Scientific Evidence Retrieval.** To quantify the contribution of the Scientific Evidence Retrieval module (Section 2.1) to the overall CER framework, we conducted an ablation study comparing the full pipeline against a variant without retrieved context for justification generation and veracity classification. Table 5 demonstrates that removing scientific evidence retrieval leads to substantial performance degradation across all evaluation metrics, with macro-F1 scores decreasing by up to 29.3%. *The performance decline can be attributed to the lack of grounded scientific evidence and the consequent increased occurrence of hallucinated content in the LLM’s reasoning process.*

**3.2.3 Evaluation of Dense vs Sparse Retrieval Approaches.** We evaluated both Dense and Sparse Retrieval approaches across the three experimental datasets. As shown in Table 6, the difference in performance between the two methods was marginal across all datasets. Sparse Retrieval demonstrated a slight advantage on HealthFC (F1: 69.90% vs. 67.36%), while Dense Retrieval achieved marginally better performance on both BioASQ-7b (F1: 95.20% vs. 95.08%) and SciFact (F1: 62.23% vs. 61.14%). Interestingly, the consistent performance across both retrieval paradigms suggests that our CER

<sup>6</sup>Perplexity.ai website

<sup>7</sup>Mixtral-8x22B model on HuggingFace

framework’s effectiveness is not heavily dependent on the choice of retrieval method. *This robustness stems from the framework’s justification mechanism, which appears to effectively leverage the retrieved scientific evidence regardless of the retrieval approach employed.* This finding has important implications for practical deployments, as it allows for flexible selection of retrieval methods based on computational constraints or specific domain requirements.

**3.2.4 Evaluation of Dense vs Sparse Retrieval Approaches.** We evaluated both Dense and Sparse Retrieval approaches across the three experimental datasets. As shown in Table 6, the difference in performance between the two methods was marginal across all datasets. Sparse Retrieval demonstrated a slight advantage on HealthFC (F1: 69.90% vs. 67.36%), while Dense Retrieval achieved marginally better performance on both BioASQ-7b (F1: 95.20% vs. 95.08%) and SciFact (F1: 62.23% vs. 61.14%). Interestingly, the consistent performance across both retrieval paradigms suggests that our CER framework is relatively robust to the choice of retrieval method. *This may be attributed to the justification mechanism, which is designed to synthesize relevant information from the retrieved scientific evidence, even when retrieval quality varies slightly.* This finding has important implications for practical deployments, as it allows for flexible selection of retrieval methods based on computational constraints or specific domain requirements.

**3.2.5 Impact of LLM Reasoning.** To evaluate the influence of prompt construction on the performance of LLM, we conducted experiments with various configurations to analyze how different reasoning factors affect prediction accuracy. This prompt – designed using a Chain-of-Thought (CoT) [50] – consists in three sections: (1) the «SYS» section describes the LLM’s role and the tasks it needs to perform, (2) the «Context» section provides sentences extracted from the knowledge base, equipping the LLM with the necessary information to classify the claim and justify its prediction, and (3) the «Question» section presents the claim that the LLM is asked to assess.

In Table 7, we present the results of an ablation study conducted using CER to evaluate the impact of different prompt components on LLM reasoning performance. Specifically, we investigate the performance when the prompt omits (1) the assignment of the *Doctor* role, (2) the inclusion of scientific evidence, and (3) the requirement to provide a justification. Results show that the full CER prompt, incorporating the *Doctor* role assignment, scientific evidence, and justification requirement, achieves the highest performance in terms of Precision (P), Recall (R), and F1-score. In particular, on the HealthFC dataset, the CER prompt achieves an F1 score of 69.90%, outperforming the ablated versions. Removing the *Doctor* role assignment results in a substantial drop to 33.26% F1, highlighting the importance of framing the LLM as a medical expert. Similarly, omitting scientific evidence or the justification requirement leads to F1-scores of 40.13% and 39.11%, respectively. The BioASQ-7b dataset further reinforces this observation, but removing scientific evidence has a less drastic, though still noticeable, impact, reducing the F1-score to 91.63%. On the SciFact dataset, the CER prompt achieves an F1-score of 61.14%. Similar to the other datasets, ablating the *Doctor* role assignment, scientific evidence, or

Dataset	Prompt Construction	P (%)	R (%)	F1 (%)
HealthFC	CER	<b>67.55</b>	<b>72.43</b>	<b>69.90</b>
	w/o <i>Doctor</i> role assignment	37.10	30.15	33.26
	w/o scientific evidence	41.81	38.59	40.13
	w/o justification	40.38	37.93	39.11
BioASQ-7b	CER	<b>93.70</b>	<b>96.75</b>	<b>95.20</b>
	w/o <i>Doctor</i> role assignment	92.12	51.01	65.69
	w/o scientific evidence	92.80	90.48	91.63
	w/o justification	80.13	62.14	69.99
SciFact	CER	<b>61.44</b>	<b>60.72</b>	<b>61.14</b>
	w/o <i>Doctor</i> role assignment	30.87	27.08	28.85
	w/o scientific evidence	39.65	34.49	36.89
	w/o justification	42.01	40.64	41.31

**Table 7: Impact of LLM Reasoning.** This table shows the performance of CER on the experimented datasets by varying the prompt with a combination of sentences and justifications. The results are macro-averaged across experiments.

justification requirement leads to considerable performance degradation. These results support our hypothesis that *combining scientific evidence with LLM reasoning, facilitated by explicit justifications and an appropriate role assignment, improves the accuracy of claim verification.*

The template used to generate justifications is shown below:


#### LLM Reasoning Prompt Template

«SYS»: You are a helpful, respectful, and honest Doctor. Always answer as helpfully as possible using the context text provided. Try to give an explanation based on the scientific evidence as follows: <Sentences returned by the Scientific Evidence module>. The claim is as follows: claim. Elaborate the scientific evidence to generate new information. Use only the knowledge in the results to answer. Be formal and use the third person. Provide a maximum 200-word response without directly mentioning the Context. Use it but don’t refer to it in your answer. Begin with: “Label:” followed by the justification label (“Yes”, “No”, or “NEI”).

«Context»: Scientific evidence retrieved.

«Question»: Claim typed by the user.

**3.2.6 Cross-dataset evaluations.** This experiment evaluated the CER architecture’s generalization across biomedical datasets in a cross-dataset training and testing scenario. For the purpose of proper cross-dataset performance evaluation, given that BioASQ-7b contains only two classes (“true” or “false”), it was decided to introduce HealthFC-2 and SciFact-2, which are the original datasets previously encountered but without the “NEI” (Not Enough Information) class. Training on the diverse BioASQ-7b dataset improved generalization to specialized datasets (HealthFC-2, SciFact-2), demonstrating the importance of larger and more diverse datasets for cross-domain robustness. In contrast, training on specialized datasets was less effective for generalizing to BioASQ-7b. These results suggest that while the CER architecture shows some generalization



**Figure 4: Cross-dataset evaluation results. Three matrix of macro precision, recall, and F1 scores for cross-dataset evaluations across the three experimental datasets.**

ability, fine-tuning or transfer learning may further enhance performance, particularly when transitioning from specialized to broader datasets.

## 4 Related work

Fact-checking has evolved through various approaches, including natural language analysis for claim verification [11], evidence retrieval and matching from knowledge bases [13], and machine learning models for classifying claims as true or false [2, 4, 46]. The advent of Artificial Intelligence (AI) has introduced innovative methodologies to enhance these processes, making fact-checking more efficient and scalable. Comprehensive surveys, such as those by Guo et al. [17] and Kotonya et al. [29], analyze these techniques, identifying key research challenges and proposing frameworks for claim detection, evidence retrieval, and verification. These frameworks have catalyzed the development of online fact-checking tools like ClaimBuster [19], Full Fact<sup>8</sup>, The Factual (IsThisCredible?)<sup>9</sup>, and Google’s Fact Check Explorer<sup>10</sup>. Hartley et al. [18] evaluated these tools, comparing their accuracy with independent fact-checking organizations and emphasizing their significance in combating misinformation. Initially developed to verify political claims and online content, fact-checking has become essential for maintaining information accuracy, particularly given its influence on public opinion and decision-making. Research into identifying linguistic markers of fake news, such as the work by Rashkin et al. [35], has expanded to include automated fact-checking in social media contexts. Bhatnagar et al. [9] demonstrated the efficacy of both “out-of-the-box” and fine-tuned summarization models for these tasks. The importance of fact-checking has become particularly pronounced in the healthcare domain, where misinformation about treatments, diagnoses, and disease prevention can pose significant public health risks. During the COVID-19 pandemic, Mohr et al. [33] highlighted the detrimental effects of such misinformation. To address these challenges, researchers have adapted political fact-checking methodologies to healthcare, creating new datasets and techniques. For instance, Barve et al. [5] developed an architecture employing a Content Similarity Vector to match URLs with reliable sources. Similarly, Gong et al. [15], Kim and Choi [27] proposed a

word-level heterogeneous graph to better capture semantic relationships between claims and evidence. Additionally, datasets like BioASQ-7b [34], HealthFC [45], and SciFact [49] have been essential for testing automated fact-checking systems in health domain, emphasizing the importance of using scientific evidence to make such systems more accurate as pointed out by Vladika et al. [44]. Fine-tuned Transformer models have shown superior performance over large language models (LLMs) in tasks such as claim detection, while LLMs excel in generative tasks like question decomposition for evidence retrieval as described by Zhang et al. [54]. However, LLMs face challenges such as hallucination, which can degrade their performance. To address this, Bèchard et al. [7] recommended employing techniques like retrieval-augmented generation (RAG) to integrate external knowledge and mitigate these issues. Building on this foundation, our study presents an innovative framework that addresses the limitations of current LLM-based approaches in medical claim verification. By integrating structured knowledge bases with LLMs, our method grounds generated explanations in verified information, reducing issues like hallucination and enhancing reliability. Simultaneously, we leverage fine-tuned transformer models optimized for claim verification tasks to improve predictive accuracy. This hybrid approach effectively balances interpretability and precision, providing transparent, evidence-based insights crucial for safeguarding public health.

## 5 Conclusion & Future Work

In this work, we presented a novel approach to biomedical fact-checking, named CER (Combining Evidence and Reasoning). Our system leverages the strengths of both scientific evidence retrieval and LLM reasoning to achieve robust and accurate claim verification. We evaluated CER on three publicly available datasets and demonstrated its effectiveness through consistent improvements over prior state-of-the-art methods. Furthermore, we conducted an ablation study to analyze the impact of different design choices and shed light on the interplay between evidence integration and LLM reasoning. Cross-dataset evaluations confirm CER’s strong generalization ability, crucial for applications across diverse biomedical domains. Future work could extend CER’s evidence retrieval to additional biomedical databases for richer context and improved accuracy. Efforts should also focus on enhancing domain generalization through adaptive training or creating diverse datasets encompassing broader biomedical topics and claim types.

*Limitations.* The architecture relies on pre-trained language models for the generation of justifications. These models have known limitations, including a tendency to generate hallucinated content (incorrect but plausible-sounding information) [52], particularly when dealing with zero-shot tasks. While our pipeline mitigates this issue with the supervised *Veracity Prediction* step, generating justifications with models that are free from the hallucinations problem could potentially improve prediction accuracy.

## Acknowledgements

This work was conducted with the financial support of (1) the PNRR MUR project PE0000013-FAIR and (2) the Italian ministry of economic development, via the ICARUS (Intelligent Contract Automation for Rethinking User Services) project (CUP: B69J23000270005).

<sup>8</sup>Full Fact’s website

<sup>9</sup>The Factual’s website

<sup>10</sup>Google’s Fact Check Explorer website

LLM	Justification
Mixtral [23]	Recent studies have established a link between bites from the Lone Star Tick ( <i>Amblyomma americanum</i> ) and the development of red meat allergy. This allergy is mediated by the production of IgE antibodies against the carbohydrate galactose- $\alpha$ -1,3-galactose ( $\alpha$ -gal), which is present in the tick’s saliva. The tick bite triggers the production of these antibodies, leading to a hypersensitivity reaction upon subsequent consumption of red meat. This phenomenon has been observed globally, with the Lone Star Tick identified as a common vector. Therefore, it is accurate to state that a bite from the Lone Star Tick can cause the victim to become allergic to red meat.
M42-health/Llama3-Med42-8B [10]	Lone Star Tick - Amblyomma Americanum (pdf) (cdc.gov) Lone Star Tick (Amblyomma americanum) - Vector of Choice for Rickettsial Diseases (vets4pets.com) Lone Star Tick (Amblyomma americanum) - Animal Bites and M.D. Facts (paws.org) CDC - Tick Size Comparison (cdc.gov) Lone Star Tick - Amblyomma americanum
BioGPT [32]	Repeat the prompt without generating a justification
Healix-1.1B-V1-Chat-dDPO <sup>11</sup>	Continue repeating 'Answer: Yes' without providing any context or justification, even if the prompt does not include information about the label or a predicted answer.
Writer/palmyra-med-70b-32k [26]	The context mentions that the rat pancreatitis-associated protein (pap) is a 16 kDa secretory protein. This indicates that pap is a protein that is secreted from cells, as it is found in the pancreatic juice after induction of pancreatic inflammation. The context also states that pap is synthesized as a preprotein with a predicted molecular weight of 16,630, further supporting the idea that it is a secretory protein. Therefore, based on the information provided, it can be concluded that pap is a secreted protein.

Table 8: Example of different justifications generated for the same claim and evidences using different LLMs.

Fact-Checker	HealthFC			BioASQ-7b		
	P (%)	R (%)	F1 (%)	P (%)	R (%)	F1 (%)
CER (DeBERTa-v3)	46.59	48.66	47.26	75.25	76.26	75.64
CER (PubMedBERT)	55.27	56.66	55.32	68.41	68.66	68.53
CER (BERT)	56.11	58.00	53.31	53.33	40.00	45.71

Table 9: Results obtained using the Palmyra LLM, which was implemented specifically for justification creation.

Dataset	HealthFC
Claim	Does pine wood improve sleep?
Justification	The context provided does not contain specific information regarding the direct impact of pine wood on sleep improvement. While it mentions the long-lasting odor of certain pine species that has been reported to relieve stress in humans, it does not establish a direct link between pine wood and improved sleep. Therefore, based on the given context, it is not possible to definitively state that pine wood improves sleep. More specific research would be needed to confirm this claim.

Table 10: Evaluation of the claim from the HealthFC dataset.

Dataset	BioASQ-7b
Claim	Can capsules or tablets with lactase improve the symptoms of lactose intolerance and make dairy products more tolerable?
Justification	The administration of lactase supplements, such as capsules or tablets, has been shown to improve lactose digestion in individuals with impaired or absent intestinal lactase activity. This is supported by a randomized, placebo-controlled, double-blind, 5-arm crossover study, which demonstrated that different lactase preparations and freeze-dried yogurt culture can affect gastrointestinal lactose digestion and reduce symptoms of lactose intolerance. Furthermore, microbial-derived beta-galactosidase enzyme preparations have been shown to improve in vivo lactose digestion and tolerance through enhanced gastrointestinal digestion of lactose. Therefore, it can be concluded that capsules or tablets with lactase can improve the symptoms of lactose intolerance and make dairy products more tolerable.

Table 11: Evaluation of the claim from the BioASQ-7b dataset.

## A Appendix

### A.1 Fine-tuning parameters

To determine the optimal parameter values, we designed a structured experimental plan based on best practices from the literature

on similar prediction tasks and tested which parameter combinations yielded the best results [21]. We tested learning rate values ranging from  $2 \cdot e^{-5}$  to  $4 \cdot e^{-5}$ , finding that  $2 \cdot e^{-5}$  offered the best balance between convergence and performance. Batch size for training and evaluation were set to 8 to optimize memory usage and model performance (per device). Several epoch values were tested, with 5 epochs proving sufficient for convergence without overfitting. A weight decay of 0.20 helped prevent overfitting while maintaining model generalization. Gradient accumulation steps have been set to 2, allowing for larger effective batch sizes without exceeding memory constraints. The metric used to select the best model was the macro F1-score, and the save strategy was based on epochs.

### A.2 Healthcare LLM

We also generated justifications using other LLMs trained in the medical field to determine if their specialization could improve the generated justifications. The LLMs selected for this analysis were 'bioGPT' [32], 'm42-health/Llama3-Med42-8B' [10], 'Healix-1.1B-V1-Chat-dDPO'<sup>11</sup>, and 'writer/palmyra-med-70b-32k' [26].

These LLMs are smaller than Mixtral, which is used in the CER architecture. They are less powerful, and most failed to accurately grasp the context and information retrieved from PubMed, as shown in the justifications in Table 8. The only LLM capable of generating justifications comparable to Mixtral was "writer/palmyra-med-70b-32k" [26]. Consequently, we proceeded to generate results following the CER architecture using the justifications generated by this model, as shown in Table 9.

<sup>11</sup>Healix-1.1B model on HuggingFace

### A.3 LLM Reasoning output examples

Table 10 and 11 show additional examples of LLM Reasoning outputs on claims extracted from HealthFC and BioASQ-7b, respectively.

## References

- [1] Sahar Abdelnabi, Rakibul Hasan, and Mario Fritz. 2022. Open-Domain, Content-based, Multi-modal Fact-checking of Out-of-Context Images via Online Resources. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*. IEEE, New Orleans, LA, USA, 14920–14929. doi:10.1109/CVPR52688.2022.01452
- [2] Tariq Alhindi, Savvas Petridis, and Samaranda Muresan. 2018. Where is Your Evidence: Improving Fact-checking by Justification Modeling. In *Proceedings of the First Workshop on Fact Extraction and VERification, FEVER@EMNLP 2018, Brussels, Belgium, November 1, 2018*, James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal (Eds.). Association for Computational Linguistics, Brussels, Belgium, 85–90. doi:10.18653/V1/W18-5513
- [3] Oberiri Destiny Apuke and Bahiyah Omar. 2021. Fake news and COVID-19: modelling the predictors of fake news sharing among social media users. *Telematics Informatics* 56 (2021), 101475. doi:10.1016/J.TELE.2020.101475
- [4] Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. Generating Fact Checking Explanations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel R. Tetreault (Eds.). Association for Computational Linguistics, Online, 7352–7364. doi:10.18653/V1/2020.ACL-MAIN.656
- [5] Yashoda Barve and Jatinderkumar R. Saini. 2023. Detecting and classifying online health misinformation with 'Content Similarity Measure (CSM)' algorithm: an automated fact-checking-based approach. *J. Supercomput.* 79, 8 (2023), 9127–9156. doi:10.1007/S11227-022-05032-Y
- [6] 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 (2020), 2055207620948996. doi:10.1177/2055207620948996 arXiv:https://doi.org/10.1177/2055207620948996 PMID: 32944269.
- [7] Patrice Béchar and Orlando Marquez Ayala. 2024. Reducing hallucination in structured outputs via Retrieval-Augmented Generation. *CoRR* abs/2404.08189 (2024), -. doi:10.48550/ARXIV.2404.08189 arXiv:2404.08189
- [8] Giannis Bekoulis, Christina Papagiannopoulou, and Nikos Deligiannis. 2021. A Review on Fact Extraction and Verification. *ACM Comput. Surv.* 55, 1, Article 12 (Nov. 2021), 35 pages. doi:10.1145/3485127
- [9] Varad Bhatnagar, Diptesh Kanojia, and Kameswari Chebrolu. 2022. Harnessing Abstractive Summarization for Fact-Checked Claim Detection. In *Proceedings of the 29th International Conference on Computational Linguistics, COLING 2022, Gyeongju, Republic of Korea, October 12-17, 2022*, Nicoletta Calzolari, Chu-Ren Huang, Hansaem Kim, James Pustejovsky, Leo Wanner, Key-Sun Choi, Pum-Mo Ryu, Hsin-Hsi Chen, Lucia Donatelli, Heng Ji, Sadao Kurohashi, Patrizia Paggio, Nianwen Xue, Seokhwan Kim, Youngyun Hahm, Zhong He, Tony Kyungil Lee, Enrico Santus, Francis Bond, and Seung-Hoon Na (Eds.). International Committee on Computational Linguistics, Gyeongju, Republic of Korea, 2934–2945. https://aclanthology.org/2022.coling-1.259
- [10] Clément Christophe, Praveen K Kanithi, Thatagata Raha, Shadab Khan, and Marco AF Pimentel. 2024. Med42-v2: A Suite of Clinical LLMs. arXiv:arXiv:2408.06142 https://arxiv.org/abs/2408.06142
- [11] Anubrata Das, Houjiang Liu, Venelin Kovatchev, and Matthew Lease. 2023. The state of human-centered NLP technology for fact-checking. *Information Processing & Management* 60 (03 2023), 103219. doi:10.1016/j.ipm.2022.103219
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR* abs/1810.04805 (2018), -. arXiv:1810.04805 http://arxiv.org/abs/1810.04805
- [13] Sushant Gautam. 2024. FactGenius: Combining Zero-Shot Prompting and Fuzzy Relation Mining to Improve Fact Verification with Knowledge Graphs. arXiv:2406.01311 [cs.CL] https://arxiv.org/abs/2406.01311
- [14] Michael A Gisondi, Rachel Barber, Jemery Samuel Faust, Ali Raja, Matthew C Strehlow, Lauren M Westafer, and Michael Gottlieb. 2022. A deadly infodemic: social media and the power of COVID-19 misinformation. e35552 pages.
- [15] Haisong Gong, Weizhi Xu, Shu Wu, Qiang Liu, and Liang Wang. 2024. Heterogeneous Graph Reasoning for Fact Checking over Texts and Tables. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, Michael J. Wooldridge, Jennifer G. Dy, and Sriraam Natarajan (Eds.). AAAI Press, Vancouver, Canada, 100–108. doi:10.1609/AAAI.V38I1.27760
- [16] Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2022. Domain-Specific Language Model Pretraining for Biomedical Natural Language Processing. *ACM Trans. Comput. Heal.* 3, 1 (2022), 2:1–2:23. doi:10.1145/3458754
- [17] Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2021. A Survey on Automated Fact-Checking. *CoRR* abs/2108.11896 (2021), -. arXiv:2108.11896 https://arxiv.org/abs/2108.11896
- [18] Russell Hartley. 2024. Efficacy Analysis of Online Artificial Intelligence Fact-Checking Tools. *The International Review of Information Ethics* 33, 1 (April 2024), -. doi:10.29173/irrie502
- [19] Naeemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. Toward Automated Fact-Checking: Detecting Check-worthy Factual Claims by ClaimBuster. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13 - 17, 2017*. ACM, Halifax, NS, Canada, 1803–1812. doi:10.1145/3097983.3098131
- [20] Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. DeBERTaV3: Improving DeBERTa using ELECTRA-Style Pre-Training with Gradient-Disentangled Embedding Sharing. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, Kigali, Rwanda, -. https://openreview.net/forum?id=sE7-XhLxHA
- [21] Mathav Raj J, Kushala VM, Harikrishna Warriar, and Yogesh Gupta. 2024. Fine Tuning LLM for Enterprise: Practical Guidelines and Recommendations. arXiv:2404.10779 [cs.SE] https://arxiv.org/abs/2404.10779
- [22] Minbyul Jeong, Jiwoong Sohn, Mujeen Sung, and Jaewoo Kang. 2024. Improving medical reasoning through retrieval and self-reflection with retrieval-augmented large language models. *Bioinformatics* 40, Supplement\_1 (2024), i119–i129.
- [23] Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, Léo Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2024. Mixtral of Experts. arXiv:2401.04088 [cs.LG] https://arxiv.org/abs/2401.04088
- [24] Qiao Jin, Won Kim, Qingyu Chen, Donald C Comeau, Lana Yeganova, W John Wilbur, and Zhiyong Lu. 2023. MedCPT: Contrastive Pre-trained Transformers with large-scale PubMed search logs for zero-shot biomedical information retrieval. *Bioinformatics* 39, 11 (Nov. 2023), -. doi:10.1093/bioinformatics/btad651
- [25] Skyler B Johnson, Matthew Parsons, Tanya Dorff, Meena S Moran, John H Ward, Stacey A Cohen, Wallace Akerley, Jessica Bauman, Joleen Hubbard, Daniel E Spratt, et al. 2022. Cancer misinformation and harmful information on Facebook and other social media: a brief report. *JNCI: Journal of the National Cancer Institute* 114, 7 (2022), 1036–1039.
- [26] Kiran Kamble and Waseem Alshikh. 2023. Palmyra-Med: Instruction-Based Fine-Tuning of LLMs Enhancing Medical Domain Performance. doi:10.13140/RG.2.2.30939.75046
- [27] Jiseong Kim and Key-Sun Choi. 2020. Unsupervised Fact Checking by Counter-Weighted Positive and Negative Evidential Paths in A Knowledge Graph. In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, Donia Scott, Núria Bel, and Chengqing Zong (Eds.). International Committee on Computational Linguistics, Online, 1677–1686. doi:10.18653/V1/2020.COLING-MAIN.147
- [28] Seonho Kim and Juntae Yoon. 2024. VAIv bio-discovery service using transformer model and retrieval augmented generation. *BMC bioinformatics* 25, 1 (2024), 273.
- [29] Neema Kotonya and Francesca Toni. 2020. Explainable Automated Fact-Checking for Public Health Claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 7740–7754. doi:10.18653/V1/2020.EMNLP-MAIN.623
- [30] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. *CoRR* abs/1909.11942 (2019), -. arXiv:1909.11942 http://arxiv.org/abs/1909.11942
- [31] Sheng Jie Lui, Cheng Xiang, and Shonali Krishnaswamy. 2024. Knowledge Aware Automated Health Claims Processing with Medical Ontologies and Large Language Models. In *Proceedings of the AAAI Symposium Series, Vol. 4. AAAI Press, Online*, 199–209.
- [32] Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng Zhang, Hoifung Poon, and Tie-Yan Liu. 2022. BioGPT: generative pre-trained transformer for biomedical text generation and mining. *Briefings in Bioinformatics* 23, 6 (09 2022), -. doi:10.1093/bib/bbac409 arXiv:https://academic.oup.com/bib/article-pdf/23/6/bbac409/47144271/bbac409.pdf bbac409.
- [33] Isabelle Mohr, Amelie Wüthrich, and Roman Klinger. 2022. CoVERT: A Corpus of Fact-checked Biomedical COVID-19 Tweets. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference, LREC 2022, Marseille, France, 20-25 June 2022*, Nicoletta Calzolari, Frédéric Béchet, Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Jan Odijk, and Stelios Piperidis (Eds.). European Language Resources Association, Marseille, France, 244–257. https://aclanthology.org/2022.lrec-1.26
- [34] Anastasios Nentidis, Konstantinos Bougiatiotis, Anastasia Krithara, and Georgios Paliouras. 2020. Results of the seventh edition of the BioASQ Challenge. *CoRR* abs/2006.09174 (2020), -. arXiv:2006.09174 https://arxiv.org/abs/2006.09174
- [35] Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of Varying Shades: Analyzing Language in Fake News and Political Fact-Checking. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, Martha Palmer, Rebecca Hwa, and Sebastian Riedel (Eds.). Association for

- Computational Linguistics, ACL, 2931–2937. doi:10.18653/V1/D17-1317
- [36] Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of Varying Shades: On Political Fact-Checking and Fake News. *Conference on Empirical Methods in Natural Language Processing 1* (2017), -. <https://par.nsf.gov/biblio/10074110>
- [37] Stephen E. Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Found. Trends Inf. Retr.* 3, 4 (2009), 333–389. doi:10.1561/15000000019
- [38] Michael S Saag. 2020. Misguided use of hydroxychloroquine for COVID-19: the infusion of politics into science. *Jama* 324, 21 (2020), 2161–2162.
- [39] Mourad Sarrouti, Asma Ben Abacha, Yassine Mrabet, and Dina Demner-Fushman. 2021. Evidence-based Fact-Checking of Health-related Claims. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Punta Cana, Dominican Republic, 3499–3512. doi:10.18653/v1/2021.findings-emnlp.297
- [40] Sebastian Walter Schuetz, Tracy Ann Sykes, and Viswanath Venkatesh. 2021. Combating COVID-19 fake news on social media through fact checking: antecedents and consequences. *Eur. J. Inf. Syst.* 30, 4 (2021), 376–388. doi:10.1080/0960085X.2021.1895682
- [41] Vinay Setty. 2024. Surprising Efficacy of Fine-Tuned Transformers for Fact-Checking over Larger Language Models. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024, Washington DC, USA, July 14–18, 2024*, Grace Hui Yang, Hongning Wang, Sam Han, Claudia Hauff, Guido Zuccon, and Yi Zhang (Eds.). ACM, Washington DC, USA, 2842–2846. doi:10.1145/3626772.3661361
- [42] Koosha Sharifani, Mahyar Amini, Yaser Akbari, and Javad Godarzi. 2022. Operating Machine Learning across Natural Language Processing Techniques for Improvement of Fabricated News Model. 20 - 44 pages.
- [43] Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. 2023. Large language models encode clinical knowledge. *Nature* 620, 7972 (2023), 172–180.
- [44] Juraj Vladika and Florian Matthes. 2024. Improving Health Question Answering with Reliable and Time-Aware Evidence Retrieval. *CoRR* abs/2404.08359 (2024), -. doi:10.48550/ARXIV.2404.08359 arXiv:2404.08359
- [45] Juraj Vladika, Phillip Schneider, and Florian Matthes. 2024. HealthFC: Verifying Health Claims with Evidence-Based Medical Fact-Checking. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20–25 May, 2024, Torino, Italy*, Nicoletta Calzolari, Min-Yen Kan, Véronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (Eds.). ELRA and ICCL, Torino, Italy, 8095–8107. <https://aclanthology.org/2024.lrec-main.709>
- [46] Juraj Vladika, Phillip Schneider, and Florian Matthes. 2024. MedREQAL: Examining Medical Knowledge Recall of Large Language Models via Question Answering. *CoRR* abs/2406.05845 (2024), -. doi:10.48550/ARXIV.2406.05845 arXiv:2406.05845
- [47] Nguyen Vo and Kyumin Lee. 2019. Learning from Fact-checkers: Analysis and Generation of Fact-checking Language. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21–25, 2019*, Benjamin Piwowarski, Max Chevalier, Éric Gaussier, Yoelle Maarek, Jian-Yun Nie, and Falk Scholer (Eds.). ACM, Paris, France, 335–344. doi:10.1145/3331184.3331248
- [48] Soroush Vosoughi, Mostafa 'Neo' Mohsenvand, and Deb Roy. 2017. Rumor Gauge: Predicting the Veracity of Rumors on Twitter. *ACM Trans. Knowl. Discov. Data* 11, 4 (2017), 50:1–50:36. doi:10.1145/3070644
- [49] David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or Fiction: Verifying Scientific Claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16–20, 2020*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 7534–7550. doi:10.18653/V1/2020.EMNLP-MAIN.609
- [50] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (Eds.). NeurIPS, New Orleans, LA, USA, -. [http://papers.nips.cc/paper\\_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html)
- [51] Zhifei Xie and Changqiao Wu. 2024. Mini-Omni2: Towards Open-source GPT-4o with Vision, Speech and Duplex Capabilities. arXiv:2410.11190 [eess.AS] <https://arxiv.org/abs/2410.11190>
- [52] Jia-Yu Yao, Kun-Peng Ning, Zhen-Hui Liu, Munan Ning, and Li Yuan. 2023. LLM Lies: Hallucinations are not Bugs, but Features as Adversarial Examples. *CoRR* abs/2310.01469 (2023), -. doi:10.48550/ARXIV.2310.01469 arXiv:2310.01469
- [53] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontañón, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, and Amr Ahmed. 2020. Big Bird: Transformers for Longer Sequences. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). NeurIPS, Online, -. <https://proceedings.neurips.cc/paper/2020/hash/c8512d142a2d849725f31a9a7a361ab9-Abstract.html>
- [54] Xuan Zhang and Wei Gao. 2024. Reinforcement Retrieval Leveraging Fine-grained Feedback for Fact Checking News Claims with Black-Box LLM. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20–25 May, 2024, Torino, Italy*, Nicoletta Calzolari, Min-Yen Kan, Véronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (Eds.). ELRA and ICCL, Torino, Italy, 13861–13873. <https://aclanthology.org/2024.lrec-main.1209>
- [55] Bu Zhong. 2023. Going beyond fact-checking to fight health misinformation: A multi-level analysis of the Twitter response to health news stories. *Int. J. Inf. Manag.* 70 (2023), 102626. doi:10.1016/J.IJINFOMGT.2023.102626