

HealthFC: Verifying Health Claims with Evidence-Based Medical Fact-Checking

Juraj Vladika, Phillip Schneider, Florian Matthes
Department of Computer Science, Technical University of Munich
Boltzmannstraße 3, Garching, Germany
{juraj.vladika, phillip.schneider, matthes}@tum.de

Abstract

In the digital age, seeking health advice on the Internet has become a common practice. At the same time, determining the trustworthiness of online medical content is increasingly challenging. Fact-checking has emerged as an approach to assess the veracity of factual claims using evidence from credible knowledge sources. To help advance automated Natural Language Processing (NLP) solutions for this task, in this paper we introduce a novel dataset HEALTHFC. It consists of 750 health-related claims in German and English, labeled for veracity by medical experts and backed with evidence from systematic reviews and clinical trials. We provide an analysis of the dataset, highlighting its characteristics and challenges. The dataset can be used for NLP tasks related to automated fact-checking, such as evidence retrieval, claim verification, or explanation generation. For testing purposes, we provide baseline systems based on different approaches, examine their performance, and discuss the findings. We show that the dataset is a challenging test bed with a high potential for future use.

Keywords: fact-checking, claim verification, question answering, misinformation, biomedical NLP

1. Introduction

Health can be defined as "a state of complete physical, mental, and social well-being" and is a popular point of discussion both in everyday life and in online spaces (Kühn and Rieger, 2017). The Internet has made seeking information about personal and public health easier than ever before. Many people have turned to online blogs and news portals as a source of evidence regarding health-related inquiries. According to recent reports, 72% of American adults (Fox and Duggan, 2013) and 55% of EU citizens (Eurostat, 2020) have searched online for medical conditions that they might have, and they first consult the Internet before deciding if they should visit a medical professional. They mostly try to find answers to their medical questions before going to a medical professional or make decisions about whether to consult a doctor or not.

With the increasing volume of new data generated daily and the rapid speed at which information is propagated in digital media, keeping track of trustworthy sources has become challenging. This has facilitated the spread of misinformation – content that is usually false, misleading, or not backed by any relevant knowledge sources. In the period of the COVID-19 pandemic, medical misinformation has led people to turn to unsafe drugs and unproven treatments (Pennycook et al., 2020; Zarcostas, 2020). The challenge of seeking credible health-related information is further amplified by the advent of digital health assistants and generative language models, which have the ability to generate eloquent responses for any input query yet are prone to "hallucinating" knowledge or omitting

Claim: Can regular intake of vitamin C prevent colds?

Evidence: The recommendation to take high-dose vitamin C at the first signs of a cold cannot be confirmed by studies. If cough, sniffing or sore throat are already present, vitamin C does not seem to have any detectable effect. The daily requirement for the vitamin is about 100 milligrams, with the recommendations slightly fluctuating [2,3]. This amount is contained in an apple, half a pepper or two tomatoes [4]. (...)

Verdict: **Refuted**

Claim: Does melatonin help against jet lag?

Evidence: This sounds plausible at first because melatonin plays an important role in sleep-wake rhythm [4]. We have found an overview of ten individual studies [1] and a newer individual study [2]. At random, the test subjects received melatonin or a dummy medication. Overall, the studies show that melatonin may help better against jet lag than a sham drug. (...)

Verdict: **Supported**

Table 1: Example of two claims from the dataset with a snippet of evidence documents and verdicts. Manually annotated evidence sentences are highlighted in violet.

important details (Ji et al., 2023).

The usual way for biomedical researchers to test their hypotheses related to human health is by conducting *clinical trials*. Clinical trials are carefully designed research studies that seek to investigate the efficacy and safety of biomedical or behavioral interventions in human subjects, which may include novel treatments such as vaccines, drugs,

dietary supplements, medical devices, or known interventions that require further examination (Piantadosi, 2017). When performed with high standards, clinical trials serve as a high-quality and trustworthy expert-curated source of evidence for health-related decisions. Multiple clinical trials related to the same topic are commonly combined into a *systematic review*. These reviews serve as a medical artifact providing guidelines concerning treatments and medical decisions with varying levels of evidence and strength of recommendation (Sekhon et al., 2017).

Fact-checking is the task of assessing factual claims that are contested, using relevant evidence from credible knowledge sources. It is a time-consuming task that is still usually performed manually by dedicated experts in journalism (Guo et al., 2022). Recently, solutions based on Machine Learning (ML) and Natural Language Processing (NLP) have been developed to automate parts of the fact-checking process. Considering the complexity of the task, current solutions are still far from achieving human-level performance. Still, they can be used to assist human fact-checkers in their work, such as discovering evidence (Nakov et al., 2021).

While multiple datasets for automated fact-checking of health-related and biomedical claims have been constructed in recent years, none of them use clinical studies as their primary source of knowledge to determine a claim's veracity. This is a major drawback, considering the importance of clinical trials and systematic reviews in making health-related decisions in medicine. Furthermore, most datasets provide only top-level labels like "true" and "false" with no information regarding the level of evidence and certainty in the label. Finally, almost all datasets contain claims solely in English.

To address these research gaps, we introduce a novel dataset MEDIZIN-TRANSPARENT, constructed from the portal² run by the University for Continuing Education Krems (Danube University Krems).

Our study presents the following contributions:

1. We introduce a bilingual German and English dataset, featuring 750 health-related claims and richly annotated metadata: veracity labels from a team of medical experts, level of evidence, and explanatory documents written in lay language describing clinical studies used for assessment. We additionally provide manually annotated evidence sentences from documents. The dataset enables testing of various NLP tasks related to automated fact-checking.
2. We develop diverse baseline systems to benchmark the performance of evidence selection and verdict prediction on the dataset and describe the findings and challenges.
3. We provide additional insight and experiments

related to different evidence sources and levels of evidence in the verification process.

We provide the dataset and code in a public GitHub repository.¹

2. Related Work

2.1. Medical NLP Tasks

Healthcare is a popular application domain in artificial intelligence and natural language processing. The complexity of language found in sources like biomedical publications and clinical trial reports makes it a challenging domain to work with. To overcome these obstacles, general-purpose NLP models are pre-trained and fine-tuned on domain-specific biomedical and scientific texts. This includes models like SciBERT (Beltagy et al., 2019) and BioBERT (Lee et al., 2020). Biomedical NLP tasks include a wide array of common NLP tasks, such as natural language inference (Jullien et al., 2023), named entity recognition (Zhao et al., 2019), dialogue systems (Zeng et al., 2020), evidence inference (DeYoung et al., 2020b), or text summarization (Abacha et al., 2021).

A knowledge-intensive NLP task related to fact-checking is *question answering* (QA). In particular, *biomedical question answering* can be divided into four groups: scientific, clinical, examination, and consumer health (Jin et al., 2022). The first three groups target questions that help medical professionals and researchers conduct their work. Examples include PubMedQA (Jin et al., 2019) and BioASQ-QA (Krithara et al., 2023). Our work is mostly similar to consumer-health QA, where the goal is to help the general population seek medical advice, and the produced answers should be user-friendly (Demner-Fushman et al., 2019). Recent QA datasets focus on producing long answers for medical inquiries by generative language models, as seen in ExpertQA (Malaviya et al., 2023). The dataset in our study intuitively belongs to automated fact-checking since it assesses claim veracity.

2.2. Medical Fact-Checking

Numerous datasets for automated fact-checking have been released in recent years (Guo et al., 2022). Most of these datasets are related to society, politics, and general online rumors. Examples include MultiFC (Augenstein et al., 2019) or the Snopes dataset (Hanselowski et al., 2019), where the authors leveraged existing claims and explanations from professional fact-checking platforms to construct the dataset. We also followed such an approach, focusing on health claims. Most fact-checking datasets contain claims and evidence

¹<https://github.com/jvladika/HealthFC/>

solely written in English. The only other dataset we found with some claims in German is the multilingual dataset X-Fact (Gupta and Srikumar, 2021), which focuses on challenges in cross-lingual transfer for automated fact-checking.

Datasets with biomedical and health-related claims started emerging since 2020 due to online content related to the pandemic of COVID-19 (Vladika and Matthes, 2023a). These datasets differ with respect to their primary source of claims and evidence. Datasets like SciFact (Wadden et al., 2020) and HealthVer (Sarrouti et al., 2021) feature expert-written claims stemming from biomedical research publications and user search queries, respectively, and pair them with abstracts of scientific publications that provide evidence for assessing the claim. On the other hand, datasets COVID-Fact (Saakyan et al., 2021) and CoVERT (Mohr et al., 2022) take social media posts to gather the claims, the former pairing claims from Reddit with accompanying evidence articles, and the latter taking causative biomedical claims from Twitter posts paired with manually annotated evidence documents from Google search results. Similar to our study in terms of dataset construction is PubHealth (Kotonya and Toni, 2020), which uses news titles of articles on public health from dedicated fact-checking websites as claims and accompanying article text as the evidence source. Still, the dataset is relatively noisy since the news titles often do not make a factual and atomic claim. Also related is RedHOT (Wadhwa et al., 2023), a dataset of medical questions from reddit that matches them with relevant clinical trial reports. Recent work by Wüthrich et al. (2024) focuses on determining the verifiability of biomedical claims based on claim properties.

To the best of our knowledge, we introduce the first dataset for medical fact-checking to use clinical trials and systematic reviews as its main source of evidence. More precisely, it utilizes knowledge from clinical studies as its initial source and presents it understandably for everyday users in a form of articles. Furthermore, this is the only dataset of health claims to feature the strength of found evidence in its labels and a short explanation paragraph for every verdict decision. It also covers a wider variety of topics concerning all segments of human health, when compared to other datasets focusing only on COVID-19-related claims (as seen in HealthVer, COVID-Fact, and CoVERT).

3. Dataset Construction

3.1. Data Source

The dataset was constructed from the publicly available data on the web portal *Medizin Transparent*.²

² <https://medizin-transparent.at/>

It is a project connected with Cochrane Austria. Cochrane is an international charitable organization formed to organize medical research findings to facilitate evidence-based decisions about health interventions involving health professionals, patients, and policymakers (Kleinstäuber et al., 1996).

The team of Medizin Transparent uses a systematic approach to perform fact-checking (Kerschner et al., 2015), described in their guidelines.³ The process usually starts with a user inquiry regarding a health-related issue. In addition, health claims that are currently trending on popular news portals are considered as well. Then, this inquiry is formed to a precisely defined question that is used to search through several research databases dealing with biomedical research, where relevant studies are manually filtered down. The preference as a primary source is given to systematic reviews since they present a comprehensive synthesis of results on a research topic in previously published studies. If no systematic reviews are available, the conclusions are drawn from as many informative individual studies as possible to make the best-informed decision. The narrowed-down studies are assessed with regard to quality and significance using previously defined criteria, ensuring the trustworthiness and consistency of the sources. The quality of studies is checked by at least two people from the project's scientific team. The results are summarized by an author, checked by a medical professional, and described in a comprehensible and easily understandable way for a wide audience.

We constructed a scraping project with the Python library Scrapy⁴ and collected all the text from the articles. Because the crawled articles from the portal are exclusively written in German, we translated them into English to provide a wider reach and alignment with similar datasets in English. Claims and explanations of the verdicts were translated with the DeepL API.⁵ For the article text, DeepL could not be used due to the character limits of the free API version. Instead, we translated the longer document texts with the Opus-MT library (Tiedemann and Thottingal, 2020), an open-source tool with a proven record of generating translations of high quality. All the translated articles were read by the authors during the evidence annotation process and any spotted mistakes in translation were manually corrected by the authors who are native German and fluent English speakers, to ensure high quality of the provided text.

³ Guidelines can be found here.

⁴ <https://scrapy.org/>

⁵ <https://www.deepl.com/pro-api>

3.2. Claims and Labels

The two main components of the dataset are claims and evidence documents. Each of the 750 claims is paired with a single evidence document. These evidence documents can be found directly on the fact-checking portal [Medizin-transparent.at](https://www.medizin-transparent.at) and were written and proofread by its medical team. The claim veracity labels also come directly from the medical experts. One specific aspect of the veracity labels in the dataset is that, on top of providing a positive ("true") or negative ("false") label, there is additionally a three-point scale denoting the *level of evidence*. This refers to how strong the findings from clinical trials were and how certain the veracity label of the claim is based on available evidence. The medical team follows internal guidelines on determining which of the three scores of level of evidence to assign to each claim.³ It is based on the GRADE framework, used for giving clinical practice recommendations ([Gopalakrishna et al., 2014](#)).

Following other common fact-checking datasets, we map all the veracity labels to three final high-level verdicts: *supported*, *refuted*, and *not enough information* (NEI). On top of these three labels, there is also a label for the aforementioned level of evidence, in case of the *supported* and *refuted* claims. In some other datasets, like SciFact, the NEI label signifies that no relevant evidence documents related to the claim are present in the dataset. On the other hand, claims labeled with NEI in the dataset of our study will always be paired with an evidence document. This evidence document usually reports how no relevant clinical studies were found in academic databases or those found are lacking in quality, and therefore, a reliable verdict on the claim's veracity cannot be made.

3.3. Evidence Annotation

Even though the evidence documents for each claim contain enough information to make a final verdict, not everything in the documents is relevant – they often contain background information that is interesting for readers, but not necessary to make a final decision. Hence, we decided to annotate individual sentences that provide evidence (rationale) for making a final verdict on claim's veracity.

Two authors served as annotators. They followed a systematic annotation process by first reading the claim along with its stated verdict and then the full article. All sentences in the article were split automatically using a sentence tokenizer from NLTK. The task was to select only those sentences that make a statement on a claim's veracity. The maximum number of sentences to be selected as rationales was capped at 5. It was empirically determined that rarely are more than 5 sentences needed to make a verdict and this number follows the convention of

other fact verification datasets like FEVER ([Thorne et al., 2018](#)) and COVID-Fact ([Saakyan et al., 2021](#)). The annotators held regular meetings to discuss and resolve any uncertainties in labeling.

In order to assess the inter-annotator agreement in the labeling process, 50 evidence documents (6.7%) were selected for mutual annotation by both authors. Cohen's κ coefficient ([Cohen, 1960](#)) was determined to be 0.72. Cohen suggested interpreting the values of κ between 0.61 and 0.80 as substantial agreement. This is comparable to Cohen's κ of 0.70 in [Hanselowski et al. \(2019\)](#) and 0.71 in [Wadden et al. \(2020\)](#), as well as Fleiss' κ coefficient ([Fleiss, 1971](#)) of 0.68 in [Thorne et al. \(2018\)](#) and 0.74 in [Hu et al. \(2022\)](#).

4. Dataset Description

In this chapter, we will provide descriptive analytics and statistics of the dataset, and outline some specific characteristics and challenges.

4.1. General Overview of Dataset

The dataset consists of 750 scientifically fact-checked health claims and evidence articles. The dataset is available in both English and German. [Figure 1](#) shows the number of yearly published articles over the project's time span. The plotted distribution reveals a significant increase in the number of articles per year, peaking in 2016 with 105 articles. After that, around 80 to 90 claims were fact-checked annually until the number dropped again in 2022. Since articles can get outdated with time as new clinical studies are published, the portal's team checked some selected claims and updated the verdicts as appropriate – in this case, the date of the article reflects the latest revision. After the release of this dataset, it is possible some verdicts get outdated with time, especially those before 2018, when the team updated their internal guidelines. Therefore, the dataset reflects the latest status at the end of 2022 and should be interpreted as such.

The dataset covers a diverse range of health topics, encompassing many subdomains. Listing all covered topics would go beyond the scope of this brief dataset description. To gain insight into the most frequently covered subdomains, a subset of the top ten topics is visualized in [Figure 2](#). The chart depicts the relative share of fact-checks among the top ten topics. It is evident that inquiries about eating habits are most popular, since the topics dietary supplements and nutrition account for respectively 18% and 15%. Dietary topics are a complex topic for which an abundance of health advice can be found on the Internet. The third most popular topic is the immune system. It plays a vital role in bodily defense and is thus responsible




Figure 1: Number of collected health fact-check articles by year of publication.

for many health conditions. Other prominent sub-domains focus on specific body systems, such as the respiratory, musculoskeletal, or cardiovascular systems. Alternative and complementary medicine covers non-traditional forms of healing.

Apart from general health domains that remain consistent over time, the dataset also contains topics that depend on current trends and events. One such topic is COVID-19, which has dominated the news and public health discussions since its outbreak in 2019. The COVID-19 pandemic has impacted people’s health and well-being worldwide. It has highlighted the importance of online resources as a primary information medium for people seeking health and medical advice.




Figure 2: Distribution of the top ten most popular health topics in the collected dataset.

4.2. Descriptive Statistics of Dataset

The dataset comprises health claims, evidence articles, verdicts, and manually annotated evidence sentences that support the final verdict. Each evidence article also contains an explanation paragraph that serves as a short summary of the article and a justification of the verdict. Table 2 summarizes descriptive statistics of the English texts in the novel dataset. It includes the mean, standard deviation, minimum, and maximum values for various aspects of the dataset. An interesting observation

is that the word count of explanations has a mean of 40.0 and ranges from only 7 up to 103 words. On the other hand, the word counts of the manually selected evidence sentences are almost twice as high. This demonstrates that, despite limiting the evidence sentences to a maximum of five, the original explanations (summaries) of the verdicts are more concise. These short explanatory summaries could be used for the task of explanation generation in future work.

Aspect	μ	σ	Min	Max
No. of evid. sent.	3.4	1.2	1	5
No. of all sentences	59.0	25.3	16	168
Words: articles	857	369	244	2677
Words: explanations	40.0	18.3	7	103
Words: evid. sent.	76.6	32.2	15	189

Table 2: Quantitative statistics of the dataset.

The absolute frequency of three verdict labels (refuted, supported, NEI) for all 750 articles is presented in Figure 3. The distribution in the chart is highly skewed towards the NEI verdict, being the majority class with 423 articles. This is to be expected due to the complex nature of research in the health field and due to the strict guidelines followed by Cochrane for giving a conclusive verdict. Health claims are often subject to ongoing research and clinical trials, and there may not be enough evidence available at the time of the assessment to determine whether a claim is true or false.




Figure 3: Evidence level count by verdict label. NEI denotes "not enough information".

The number of articles where claims are supported (202) is less than half of those in the NEI class, while the fewest belong to the refuted (125) category. Supported and refuted claims additionally have levels of evidence provided on a three-point scale, which refer to the frequency and strength of discovered evidence in clinical studies related to the claim. The distribution of the level of evidence

is also skewed towards lower levels. This demonstrates once again the challenge of making decisive health assessments. A similar distribution was found by [Howick et al. \(2022\)](#), who manually analyzed 2,428 Cochrane reviews and found that only 6% of interventions had high-quality evidence to support the effects, with an additional 20% having evidence of moderate quality for supported claims.

5. Baselines

In this section, we introduce the problem and describe baseline systems used to benchmark the performance on evidence selection and veracity prediction on the dataset. We experimented with two types of systems (pipeline and joint) and four different base language models. Pipeline and joint systems are used from [Vladika and Matthes \(2023b\)](#).

5.1. Problem Statement

The process of automated fact-checking in the dataset consists of two major components: evidence selection and veracity prediction.

Evidence selection. A binary classification task, where given a claim c and an evidence document consisting of n sentences s_1, s_2, \dots, s_n , the task is to train a model that predicts $z_i = 1[s_i \text{ is an evidence sentence}]$.

Veracity prediction. A ternary classification task, where for a given claim c and k previously selected evidence sentences $\vec{e} = e_1, e_2, \dots, e_k$, the goal is to predict one of the three classes of the final verdicts: $y(c, \vec{e}) \in \{\text{SUPPORTED}, \text{REFUTED}, \text{NOT ENOUGH INFO (NEI)}\}$.

5.2. Pipeline Systems

The intuitive approach is to develop two separate models – one for evidence selection and another for veracity prediction. The evidence sentences selected by the first model are used as input for veracity prediction in the next step. It is also common to use the same underlying base model in both steps and fine-tune it for these two different tasks ([DeYoung et al., 2020a](#)).

Each candidate sentence s_i from the document is concatenated with the claim c to obtain candidate sequences in the form of $a_i = [s_i; SEP; c]$. Each sequence is encoded with a base language model to obtain their dense representation: $h_i = BERT(a_i)$. This representation is then fed to the classifier model Multi-Layer Perceptron (MLP) that assigns probabilities on the candidate sentence being evidence: $p_i, \bar{p}_i = softmax(MLP(h_i))$. Finally, a selection function labels sentences with a probability over a threshold (fixed at 0.5) as evidence sentences: $z_i = p_i > 0.5$. In the end, the model selected k final evidence sentences e_1, e_2, \dots, e_k as

input for the next step. We set k to 5 sentences after experimenting with different values.

The task of veracity prediction is commonly modeled in automated fact-checking as the established task Natural Language Inference (NLI), or more specifically, Recognizing Textual Entailment (RTE), which aims to infer the logical relation (entailment/contradiction/neutral) between a hypothesis and a premise. In our case, the hypothesis is the claim c , and the premise is a concatenation of evidence sentences $e = [e_1; e_2; \dots; e_k]$. These two are joined as $x = [c; SEP; e]$ and embedded as $w = BERT(x)$. The final model for sequence classification has to learn the function $\hat{y}(c; e) = softmax(MLP(w))$, which is the probability of a veracity label for the claim c given evidence e . The class with the highest probability score is selected as the final verdict $v(c; e) = argmax(y)$.

5.3. Joint Systems

Another approach is a system that jointly learns both the tasks of evidence retrieval and veracity prediction. This type of training leverages multi-task learning (MTL) and is beneficial because of data efficiency, reduced overfitting, and faster learning with auxiliary information ([Crawshaw, 2020](#)). This is achieved by modeling a unified representation of the claim and the document used for both tasks and a joint loss function that combines the evidence selection loss and veracity prediction loss.

The claim c is concatenated together with all of the sentences s_1, s_2, \dots, s_n in the document to obtain a claim+document sequence $seq = [c; SEP; s_1; SEP; s_2; \dots; SEP; s_n]$.⁶ This sequence is embedded as $h = BERT(seq) = [h_c; SEP; h_{s_1}; \dots; SEP; h_{s_n}]$. The representation of each candidate sentence $h_{s_i} = [h_{w_1}, h_{w_2}, \dots, h_{w_m}]$ is singled out from the initial representation and passed to a binary linear classifier that calculates the probabilities of the sentence being evidence: $p_i, \bar{p}_i = softmax(MLP(h_{s_i}))$. Those sentences that are above the 0.5 threshold are selected and used to form the final claim+evidence representation $h_f = [h_c; h_{e_1}, h_{e_2}, \dots, h_{e_k}]$. This representation is given to a ternary classifier that predicts the verdict $v = argmax(softmax(MLP(h_f)))$.

5.4. Encoding Models

To encode the text, we experimented with a number of underlying base models that we found representative of different aspects we wanted to test. BERT ([Devlin et al., 2019](#)) is used as the representative vanilla pre-trained language model (PLM), which

⁶This sequence had to be truncated to 512 or 1024 tokens with respect to input limitations, which is good enough for the vast majority of documents in the dataset.

System	Base Model	Evidence Selection			Veracity Prediction			Oracle Ver. Pred.
		Precision	Recall	F1 Macro	Precision	Recall	F1 Macro	F1 Macro
pipeline	XLM-R (German)	48.0 _{1.5}	48.3 _{3.6}	48.1 _{2.1}	59.0 _{9.2}	59.1 _{8.4}	58.4 _{9.1}	73.6 _{6.5}
	XLM-R (English)	51.9 _{2.1}	52.9 _{4.0}	52.3 _{1.4}	64.8 _{7.4}	60.0 _{5.7}	60.4 _{6.1}	74.7 _{4.4}
	BERT	51.4 _{3.4}	51.0 _{1.9}	51.2 _{1.6}	50.0 _{4.8}	50.5 _{4.2}	50.1 _{4.6}	69.9 _{5.6}
	BioBERT	52.2 _{2.0}	54.4 _{2.5}	53.2 _{1.0}	62.6 _{5.3}	56.3 _{4.0}	57.2 _{4.2}	78.1 _{3.8}
	DeBERTa	54.8 _{3.7}	56.6 _{3.3}	55.5 _{1.0}	67.6 _{4.6}	64.5 _{4.0}	65.1 _{3.2}	81.9 _{1.4}
joint	BERT	79.1 _{3.8}	70.0 _{2.1}	73.2 _{1.7}	69.4 _{4.0}	65.5 _{3.5}	66.9 _{4.4}	—
	BioBERT	64.2 _{2.2}	74.2 _{2.8}	67.4 _{1.2}	65.1 _{3.9}	63.0 _{4.2}	63.1 _{3.7}	—
	DeBERTa	71.8 _{2.8}	75.2 _{3.5}	73.4 _{1.4}	68.2 _{4.6}	66.8 _{4.0}	67.5 _{3.2}	—

Table 3: Results of all baseline systems and models in the form of the mean and standard deviation (subscript) of a 5-fold cross-validation over the dataset.

gives a good initial insight into the performance of PLMs on the dataset. BioBERT (Lee et al., 2020), an extension of BERT that was fine-tuned to abstracts of biomedical scientific publications, is used to check whether the medical terminology and relations it learned will help assess the claims in this dataset. Additionally, DeBERTa-v3 (He et al., 2021), an improvement of BERT with enhanced training procedure based on disentangled attention, was chosen because it has proven to be powerful for natural language understanding (NLU) tasks, in particular natural language inference and entailment recognition. Finally, XLM-RoBERTa (Conneau et al., 2020) is chosen to contrast the performance between the English and German versions of the dataset because it is a powerful multilingual model that was also shown to work well on NLP tasks involving German text (Vladika et al., 2022).

6. Experiments

In this section, we report on the experimental setup and results achieved by different systems.

6.1. Setup

We performed an array of experiments to test the performance of baseline systems on the two common fact-checking tasks. Considering the dataset’s relatively small size, we opted out of declaring a small subset of the dataset to be a test set, but instead split it into 5 folds of equal size and equal label distribution and then performed a 5-fold cross-validation procedure with the final scores being shown with a mean and standard deviation. These five splits are released together with the dataset for easier reproducibility. The hyperparameters were mostly the same for all models and setups: learning rate 10^{-5} , warmup rate 0.06, weight decay 0.01, batch size 4, epochs 7. For all of the models, their *Large* version was used, imported from the HuggingFace repository. The experiments were run on a single Nvidia V100 GPU card with 16 GB of VRAM, for one computation hour per experiment.

6.2. Results

The final results of main experiments are shown in Table 3. The results show the mean and standard deviation over the 5-fold cross-validation of precision, recall, and F1 score, which are useful classification metrics for a dataset with an imbalanced label distribution. All the metrics are macro-averaged scores over the three classes. All experiments were run on the English version of the dataset, except *XLM-R (German)*, which was run on the German version of the dataset. The task of *Evidence Selection* consisted of predicting for each candidate sentence whether it belongs or not to evidence sentences, with only about 6% of all candidate sentences had a positive label in this task. The selected sentences in this task are passed over to models in the next task, the *Veracity Prediction*. It is a three-class classification problem with the goal of predicting one of the three classes. The models used for Veracity Prediction were fine-tuned to predict the label with gold sentences, but during inference time, they used model-selected sentences from the previous step. In the last column, *Oracle Verdict Prediction*, we show the scenario where manually annotated ("gold") evidence sentences were used as input to the label-prediction model.

Other than these main experiments, some additional experiments with different settings were performed, in order to test particular challenges in open-domain verification of the dataset. Their results are shown in Table 5. The experiment *Claims only* predicts the veracity by only taking into account the claim text, with no evidence at all. For the experiment *Google snippets*, we ran a search over the Google Search API (on September 1, 2023) with our claims in English and collected the snippets from the first 10 results. These snippets were concatenated as evidence, fine-tuned, and claim veracity was predicted. This is to test the open-domain claim verification, as Google snippets were used as evidence in other fact-checking datasets (Augenstein et al., 2019; Hu et al., 2022). The experiment *Gold explanations* utilizes explanatory summaries written by authors at the beginning of every fact-check article to check how useful these

summaries are for veracity prediction. All of these experiments used the same 5 folds as the previous table. Finally, the experiment *Level of Evidence* aimed to predict one of the three categories of the level of evidence (*low*, *medium*, or *high*), for which the distribution is shown in Figure 3 – it is worth noting that the gold labels are also uncertain because there is a variation in labels of human annotators and a level of subjectivity in the GRADE system.

7. Discussion

In this section, we discuss the main findings of the experiments, provide deeper insights, perform a qualitative error analysis, and outline the open challenges for future work.

7.1. Main Results Analysis

As can be seen in Table 3, the basic BERT model provides results slightly above 50.0 (F1 macro) for both tasks, which is solid considering the dataset is imbalanced. Still, the biomedical model BioBERT outperforms BERT in both types of systems, which show the benefit of using a domain-specific model for a dataset that includes biomedical terminology. Even though the dataset consists of text written to be understandable to everyday users, it still features a wide array of medical terms and nuances that were probably better captured by BioBERT. Nevertheless, DeBERTa outperformed both BERT and BioBERT in the pipeline system, by a massive margin on veracity prediction. This shows the power of this model for the task of natural language inference in general. This shows that being optimized for good performance in a specific NLP task like entailment recognition can beat a simpler model that is optimized for a specific domain.

When looking at the performance of XLM-RoBERTa for the parallel German and English corpus, it is evident that it worked better for English, especially for evidence sentence selection. This likely stems from the fact that even though the model is multilingual, English was still the most prevalent dataset during pre-training. Still, the results for German are decent while leaving room for improvement. Developing language-specific NLP solutions for a task like this is useful because the speakers of the said language will often seek health advice on the Internet in their native language, so improving the performance on the German version of the dataset remains an open challenge.

The veracity prediction performance with oracle sentences is by far superior to the setup where the model has to select evidence sentences on its own. This shows that detecting appropriate evidence spans and arguments in unstructured text, for a given claim or query, is a challenging prob-

lem. Furthermore, the joint systems show a clear dominance for evidence selection and veracity prediction. Especially for evidence selection, the significant improvement stems from the fact that the learned representation takes the whole document into context and contextualizes sentences to their surroundings. The task of veracity prediction is also improved, which shows the clear benefit of multi-task learning and joint task modeling.

7.2. Qualitative Error Analysis

To get a deeper insight into the baseline model performance, we do a qualitative error analysis.

Claim: Are vegetables prepared in a microwave oven less healthy than those prepared in other ways?
Gold evidence: <i>If there are health problems related to the microwave, then this is not because the microwave ingredients are destroyed or changed, but because it is simply too unhealthy to eat altogether. The week-long feeding with always several times warmed up in the microwave has not led to any signs of poisoning.</i>
Selected evidence: Recent studies see the matter more differentiated: Depending on the variety of vegetables and the exact type of preparation, the microwave can sometimes be better and sometimes worse suited than, for example, cooking [2-12]. <i>If there are health problems related to the microwave, then this is not because the microwave ingredients are destroyed or changed, but because it is simply too unhealthy to eat altogether.</i>
Gold label: Refuted Predicted: Supported
Claim: Does cat's claw improve joint disease symptoms?
Gold evidence: Whether cat claw helps better in rheumatoid arthritis or osteoarthritis than a placebo cannot be reliably estimated on the basis of the available studies. We can't make any statements about effectiveness.
Selected evidence: People with joint disease often have a high level of suffering, especially if the medically prescribed products do not help enough. In the case of joint wear (arthrosis) and rheumatoid arthritis, preparations that promise a rapid remedy and are available free of prescription have an economic impact. In laboratory and animal studies there was evidence of a possible anti-inflammatory effect of cat claw extracts. As a whole, only a few patients were involved, and very different preparations were tested. This also includes various joint diseases.
Gold label: NEI Predicted: Supported
Claim: Does brain training boost intelligence?
Gold evidence: Those who train their memory only get better in these exercises, i.e. only in working memory and not in other aspects of intelligence. A review from 2013 casts doubt on how much cognitive training can be helpful for children and adolescents with various mental developmental disorders.
Selected evidence: [None]
Gold label: Refuted Predicted: NEI

Table 4: Examples of claims, with gold and selected evidence snippets, where the DeBERTa baseline made incorrect predictions.

Table 4 shows examples where the best-performing DeBERTa model made incorrect predictions. In the first example, even though the baseline model retrieved one of the gold evidence sentences with a refuting conclusion, it also retrieved mentions of earlier studies that supported the claim, which confused the final NLI prediction model. In the second example, only evidence snippets talking

about a possible supporting effect were retrieved, which made it hard to predict the gold NEI verdict. In the third example, none of the sentences from the article were above the required threshold to be selected, which led to an incorrect NEI verdict. Improvements to the verification process would include changing the prediction thresholds.

7.3. Challenges in Open-Domain Verification

Table 5 shows the results of the additional experiments. For claims only, the classifier is slightly better than random, which shows DeBERTa model utilized its internal world knowledge for predictions but is still considerably worse than any other setup from Table 3. This shows there are no linguistic patterns spoiling the results, and evidence is needed for a genuine verdict prediction. Fact-checking with Google snippets performs poorly and indicates that these snippets are not informative enough to reach a conclusion, on top of possibly coming from untrustworthy sources. Future work could explore the open-domain claim verification performance using sources such as Wikipedia or PubMed. The performance with explanation summaries is decent but still lacking when compared to using evidence sentences. Future work could utilize these human-written summaries to produce natural language explanations to justify predicted claim verdicts. Finally, the prediction of the level of evidence is also considerably poor and indicates how utilizing this aspect in fact-checking is yet to be explored and refined.

Experiment	P	R	F1
Claims only	37.4 _{2.2}	43.8 _{1.0}	39.3 _{0.5}
Google snippets	44.9 _{6.0}	48.0 _{2.4}	44.4 _{1.6}
Gold explanations	46.8 _{1.8}	57.4 _{2.5}	51.2 _{2.1}
Level of evidence	40.3 _{0.6}	44.4 _{0.8}	42.2 _{0.8}

Table 5: Results of additional experiments with DeBERTa on veracity prediction with different evidence sources and predicting the level of evidence.

We also refer to the work of [Vladika and Matthes \(2024\)](#) on comparing knowledge sources for open-domain scientific claim verification. In this work, authors test the verification performance on this study’s dataset and three other datasets using only evidence retrieved from knowledge bases Wikipedia, PubMed, and Google Search. They use only a subset of the dataset with 327 supported and refuted claims. The results show that evidence retrieved from Wikipedia achieved a binary F1 prediction score 76.5, PubMed 72.0, and Google snippets 74.5. For datasets with complex biomedical claims, PubMed outperformed Wikipedia. This shows that the dataset from our study is better suited for verifi-

cation over Wikipedia, owing to its focus on common health concerns asked by a wide audience.

As future work, it would be interesting to test the performance of recent Large Language Models (LLMs) on this dataset, especially those specialized for the biomedical domain like PMC-LLaMA ([Wu et al., 2023](#)), Med-PaLM ([Singhal et al., 2023](#)), or BioMistral ([Labrak et al., 2024](#)).

8. Conclusion

We introduce a novel fact-checking dataset for verifying claims related to everyday health-related concerns. It comprises 750 claims based on users’ online inquiries, rich metadata including final verdict labels, explanation paragraphs, full evidence documents, and manually annotated rationale sentences. We describe the dataset creation and collection process in detail and present descriptive statistics. Finally, we provide results of extensive experiments with two types of baseline systems with multiple base models and show that joint systems with full-document representation outperform the more common pipeline systems. We anticipate that the dataset can help advance the state of automated medical fact-checking and be used for NLP tasks not covered in this paper, such as open-domain verification, explanation generation, and conversational health assistants.

Limitations

Our study employed a rigorous research design that involved the collection, annotation, as well as analysis of a data corpus about fact-checks of health claims. However, it is crucial to acknowledge certain constraints. For one thing, the crawled text data was automatically translated from German to English, which may have resulted in translation errors, especially in view of particular layman’s terms or idioms that are difficult to translate. Still, upon manual annotation of evidence sentences, we corrected any spotted errors and inconsistencies.

Moreover, owing to the German-speaking readership of the Austria-based portal Medizin Transparent, a few articles might focus on topics related to healthcare practices or products popular in Germany, Austria, and Switzerland (e.g., *Does taking LaVita promote health or physical and mental performance?*). Nevertheless, a vast majority of health claims in the dataset talk about common health concerns and can be applied globally without being restricted to a specific country or region. As a last point, it is possible that errors have occurred during the dataset annotation, leading to a not ideal selection of evidence sentences. To mitigate this, the annotators discussed their labels in regular meetings to establish a clear understanding of the task.

Acknowledgements

This research has been supported by the German Federal Ministry of Education and Research (BMBF) grant 01IS17049 Software Campus 2.0 (TU München). We would like to thank the anonymous reviewers for insightful and helpful feedback.

We would like to thank the authors of articles from the Medizin-Transparent team of the University for Continuing Education Krems (Danube University Krems) for their valuable work over the years, which has enabled the construction of the dataset. This includes Bernd Kerschner, Jana Meixner, Teresa König, Iris Hinneburg, Julia Harlfinger, Claudia Christof, Jörg Wipplinger, Iris Mair, Verena Ahne, Tanja Wolf, Björn Bernitt, Lilith Teusch, and others.

Bibliographical References

- Asma Ben Abacha, Yassine M'rabet, Yuhao Zhang, Chaitanya Shivade, Curtis Langlotz, and Dina Demner-Fushman. 2021. Overview of the mediqa 2021 shared task on summarization in the medical domain. In *Proceedings of the 20th Workshop on Biomedical Language Processing*, pages 74–85.
- Alfred V. Aho and Jeffrey D. Ullman. 1972. *The Theory of Parsing, Translation and Compiling*, volume 1. Prentice-Hall, Englewood Cliffs, NJ.
- American Psychological Association. 1983. *Publications Manual*. American Psychological Association, Washington, DC.
- Rie Kubota Ando and Tong Zhang. 2005. [A framework for learning predictive structures from multiple tasks and unlabeled data](#). *Journal of Machine Learning Research*, 6:1817–1853.
- Galen Andrew and Jianfeng Gao. 2007. [Scalable training of \$L_1\$ -regularized log-linear models](#). In *Proceedings of the 24th International Conference on Machine Learning*, pages 33–40.
- Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. [MultiFC: A real-world multi-domain dataset for evidence-based fact checking of claims](#). 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 4685–4697, Hong Kong, China. Association for Computational Linguistics.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). 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 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Ashok K. Chandra, Dexter C. Kozen, and Larry J. Stockmeyer. 1981. [Alternation](#). *Journal of the Association for Computing Machinery*, 28(1):114–133.
- Jacob Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and Psychological Measurement*, 20(1):37–46.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- James W. Cooley and John W. Tukey. 1965. [An algorithm for the machine calculation of complex Fourier series](#). *Mathematics of Computation*, 19(90):297–301.
- Michael Crawshaw. 2020. Multi-task learning with deep neural networks: A survey. *CoRR*, abs/2009.09796.
- Dina Demner-Fushman, Yassine Mrabet, and Asma Ben Abacha. 2019. [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.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. 2020a. [ERASER: A benchmark to evaluate rationalized NLP models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4443–4458, Online. Association for Computational Linguistics.

- Jay DeYoung, Eric Lehman, Benjamin Nye, Iain Marshall, and Byron C. Wallace. 2020b. [Evidence inference 2.0: More data, better models](#). In *Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing*, pages 123–132, Online. Association for Computational Linguistics.
- Eurostat. 2020. [One in two eu citizens look for health information online](#). *Survey on the use of ICT in households and by individuals*, 2020.
- JL Fleiss. 1971. [Measuring nominal scale agreement among many raters](#). *Psychological bulletin*, 76(5):378–382.
- Susannah Fox and Maeve Duggan. 2013. [Health online 2013](#). *Health*, 2013:1–55.
- Gowri Gopalakrishna, Reem A. Mustafa, Clare Davenport, Rob J.P.M. Scholten, Christopher Hyde, Jan Brozek, Holger J. Schünemann, Patrick M.M. Bossuyt, Mariska M.G. Leeflang, and Miranda W. Langendam. 2014. [Applying grading of recommendations assessment, development and evaluation \(grade\) to diagnostic tests was challenging but doable](#). *Journal of Clinical Epidemiology*, 67(7):760–768.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. [A survey on automated fact-checking](#). *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Ashim Gupta and Vivek Srikumar. 2021. X-fact: A new benchmark dataset for multilingual fact checking. *arXiv preprint arXiv:2106.09248*.
- Dan Gusfield. 1997. *Algorithms on Strings, Trees and Sequences*. Cambridge University Press, Cambridge, UK.
- Andreas Hanselowski, Christian Stab, Claudia Schulz, Zile Li, and Iryna Gurevych. 2019. A richly annotated corpus for different tasks in automated fact-checking. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 493–503.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. [DeBERTa: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.
- Jeremy Howick, Despina Koletsi, John PA Ioannidis, Claire Madigan, Nikolaos Pandis, Martin Loef, Harald Walach, Sebastian Sauer, Jos Kleijnen, Jadbinder Seehra, et al. 2022. Most health-care interventions tested in cochrane reviews are not effective according to high quality evidence: a systematic review and meta-analysis. *Journal of clinical epidemiology*, 148:160–169.
- Xuming Hu, Zhijiang Guo, GuanYu Wu, Aiwei Liu, Lijie Wen, and Philip Yu. 2022. [CHEF: A pilot Chinese dataset for evidence-based fact-checking](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3362–3376, Seattle, United States. Association for Computational Linguistics.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of hallucination in natural language generation](#). *ACM Comput. Surv.*, 55(12).
- 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.
- Mael Jullien, Marco Valentino, Hannah Frost, Paul O’Regan, Dónal Landers, and Andre Freitas. 2023. [NLI4CT: Multi-evidence natural language inference for clinical trial reports](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 16745–16764, Singapore. Association for Computational Linguistics.
- Bernd Kerschner, Jörg Wipplinger, Irma Klerings, and Gerald Gartlehner. 2015. Wie evidenzbasiert berichten print- und online-medien in Österreich? eine quantitative analyse. *Z. Evid. Fortbild. Qual. Gesundheitswes.*, 109(4-5):341–349.
- Maria Kleinstäuber, Michael Witthöft, Andrés Stefanowski, Harm Van Marwijk, Wolfgang Hiller, and Michael J Lambert. 1996. *Cochrane database of systematic reviews*.
- 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)*, pages 7740–7754, Online. Association for Computational Linguistics.
- Anastasia Krithara, Anastasios Nentidis, Konstantinos Bougiatiotis, and Georgios Paliouras.

2023. [Bioasq-qa: A manually curated corpus for biomedical question answering](#). *Scientific Data*, 10:170.
- Shafreena Kühn and Ulrich M Rieger. 2017. Health is a state of complete physical, mental and social well-being and not merely absence of disease or infirmity. *Surg. Obes. Relat. Dis.*, 13(5):887.
- Yanis Labrak, Adrien Bazoge, Emmanuel Morin, Pierre-Antoine Gourraud, Mickael Rouvier, and Richard Dufour. 2024. [Biomistral: A collection of open-source pretrained large language models for medical domains](#).
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pretrained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. 2023. [Expertqa: Expert-curated questions and attributed answers](#).
- Iain J Marshall and Byron C Wallace. 2019. Toward systematic review automation: a practical guide to using machine learning tools in research synthesis. *Systematic reviews*, 8:1–10.
- Isabelle Mohr, Amelie Wüthrl, and Roman Klinger. 2022. [Covert: A corpus of fact-checked biomedical covid-19 tweets](#). In *Proceedings of the Language Resources and Evaluation Conference*, pages 244–257, Marseille, France. European Language Resources Association.
- Preslav Nakov, David P. A. Corney, Maram Hasanain, Firoj Alam, Tamer Elsayed, Alberto Barr'on-Cedeno, Paolo Papotti, Shaden Shaar, and Giovanni Da San Martino. 2021. Automated fact-checking for assisting human fact-checkers. In *International Joint Conference on Artificial Intelligence*.
- Gordon Pennycook, Jonathon McPhetres, Yunhao Zhang, Jackson G. Lu, and David G. Rand. 2020. [Fighting covid-19 misinformation on social media: Experimental evidence for a scalable accuracy-nudge intervention](#). *Psychological Science*, 31(7):770–780. PMID: 32603243.
- Steven Piantadosi. 2017. *Clinical trials: a methodological perspective*. John Wiley & Sons.
- Lisa Pucknat, Maren Pielka, and Rafet Sifa. 2021. Detecting contradictions in german text: A comparative study. In *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 01–07. IEEE.
- Mohammad Sadegh Rasooli and Joel R. Tetreault. 2015. [Yara parser: A fast and accurate dependency parser](#). *Computing Research Repository*, arXiv:1503.06733. Version 2.
- Anna Rogers, Matt Gardner, and Isabelle Augenstein. 2023. [Qa dataset explosion: A taxonomy of nlp resources for question answering and reading comprehension](#). *ACM Comput. Surv.*, 55(10).
- Alexey Romanov and Chaitanya Shivade. 2018. Lessons from natural language inference in the clinical domain. *arXiv preprint arXiv:1808.06752*.
- Arkadiy Saakyan, Tuhin Chakrabarty, and Smaranda Muresan. 2021. [COVID-fact: Fact extraction and verification of real-world claims on COVID-19 pandemic](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2116–2129, Online. Association for Computational Linguistics.
- 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*, pages 3499–3512, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mandeep Sekhon, Martin Cartwright, and Jill J Francis. 2017. Acceptability of healthcare interventions: an overview of reviews and development of a theoretical framework. *BMC health services research*, 17(1):1–13.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Perry Payne, Stephen Pfohl, Martin Seneviratne, Paul Gamble, Christopher Kelly, Abubakr Abdelrazig Hassan Babiker, Nathanael Schaeferli, Aakanksha Chowdhery, Philip Mansfield, Dina Demner-Fushman, Blaise Aguerre-Arcas, Dale Webster, Greg Corrado, Yossi Matias, Katherine Chou, Juraj Gottweis, Nenad Tomašev, Yun Liu, Alvin Rajkomar, Joelle Barral, Christopher Semurs, Alan Karthikesalingam, and Vivek Nataraian. 2023. [Large language models encode clinical knowledge](#). *Nature*.
- James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2018. [The fact extraction and VERification \(FEVER\) shared task](#). In *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*, pages 1–9, Brussels, Belgium. Association for Computational Linguistics.

- Jörg Tiedemann and Santhosh Thottingal. 2020. OPUS-MT — Building open translation services for the World. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation (EAMT)*, Lisbon, Portugal.
- 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, Stephen Meisenbacher, and Florian Matthes. 2022. [TUM sebis at GermEval 2022: A hybrid model leveraging Gaussian processes and fine-tuned XLM-RoBERTa for German text complexity analysis](#). In *Proceedings of the GermEval 2022 Workshop on Text Complexity Assessment of German Text*, pages 51–56, Potsdam, Germany. Association for Computational Linguistics.
- 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)*, pages 7534–7550, Online. Association for Computational Linguistics.
- David Wadden, Kyle Lo, Lucy Wang, Arman Cohan, Iz Beltagy, and Hannaneh Hajishirzi. 2022. [MultiVerS: Improving scientific claim verification with weak supervision and full-document context](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 61–76, Seattle, United States. Association for Computational Linguistics.
- Somin Wadhwa, Vivek Khetan, Silvio Amir, and Byron Wallace. 2023. [RedHOT: A corpus of annotated medical questions, experiences, and claims on social media](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 809–827, Dubrovnik, Croatia. Association for Computational Linguistics.
- Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. [Pmc-llama: Towards building open-source language models for medicine](#).
- Amelie Wührl, Yarik Menchaca Resendiz, Lara Grimminger, and Roman Klinger. 2024. [What makes medical claims \(un\)verifiable? analyzing entity and relation properties for fact verification](#).
- John Zarocostas. 2020. How to fight an infodemic. *The lancet*, 395(10225):676.
- Guangtao Zeng, Wenmian Yang, Zeqian Ju, Yue Yang, Sicheng Wang, Ruisi Zhang, Meng Zhou, Jiaqi Zeng, Xiangyu Dong, Ruoyu Zhang, et al. 2020. Meddialog: Large-scale medical dialogue datasets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9241–9250.
- Sendong Zhao, Ting Liu, Sicheng Zhao, and Fei Wang. 2019. A neural multi-task learning framework to jointly model medical named entity recognition and normalization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 817–824.