

Assessing Automated Fact-Checking for Medical LLM Responses with Knowledge Graphs

Shasha Zhou^{1,2}, Mingyu Huang¹, Jack Cole², Charles Britton³, Ming Yin⁴, Jan Wolber⁵, Ke Li²

¹ Department of Computer Science and Engineering, University of Electronic Science and Technology of China

² Department of Computer Science, University of Exeter

³ College of Medicine and Health, University of Exeter

⁴ Department of Computer Science, Purdue University

⁵ GE HealthCare

{sz484, m.huang, jc1417, k.li}@exeter.ac.uk, charles.britton@nhs.scot, mingyin@purdue.edu, jan.wolber@gehealthcare.com

Abstract

The recent proliferation of large language models (LLMs) holds the potential to revolutionize healthcare, with strong capabilities in diverse medical tasks. Yet, deploying LLMs in high-stakes healthcare settings requires rigorous verification and validation to understand any potential harm. This paper investigates the reliability and viability of using medical knowledge graphs (KGs) for the automated factuality evaluation of LLM-generated responses. To ground this investigation, we introduce FAITH, a framework designed to systematically probe the strengths and limitations of this KG-based approach. FAITH operates without reference answers by decomposing responses into atomic claims, linking them to a medical KG, and scoring them based on evidence paths. Experiments on diverse medical tasks with human subjective evaluations demonstrate that KG-grounded evaluation achieves considerably higher correlations with clinician judgments and can effectively distinguish LLMs with varying capabilities. It is also robust to textual variances. The inherent explainability of its scoring can further help users understand and mitigate the limitations of current LLMs. We conclude that while limitations exist, leveraging KGs is a prominent direction for automated factuality assessment in healthcare.

Code — <https://github.com/COLA-Laboratory/FAITH>

Appendix — <https://zenodo.org/records/17603819>

Introduction

Large language models (LLMs) have emerged as powerful tools with potential to revolutionize healthcare. These models demonstrate strong capabilities across diverse medical tasks (Singhal et al. 2023; Saab *et al.* 2024; Nori et al. 2023; Van Veen et al. 2024; Liu et al. 2025; McDuff et al. 2025). However, deploying LLMs in the high-stakes healthcare domain, where precise and trustworthy information is critical, calls for rigorous evaluation. Due to inherent characteristics of LLMs like the vast output scale, automated evaluation methods have become increasingly necessary. These methods complement traditional validation approaches like clinical studies (Singhal et al. 2023; Wang et al. 2023b), which are often too slow or limited in scope to keep pace

with rapidly evolving LLM technologies. Foremost among these evaluation concerns is ensuring the factual integrity of LLM-generated medical content (Wang et al. 2023a; Thirunavukarasu et al. 2023; Lee, Bubeck, and Petro 2023; Shen et al. 2023; Augenstein et al. 2024). This is crucial as LLMs can produce plausible yet dangerously inaccurate information, known as hallucinations (Farquhar et al. 2024; Ji et al. 2023). Such inaccuracies undermine stakeholder’s trust and pose significant barriers to clinical adoption and regulatory approval of LLM applications.

Traditionally, automated evaluation of LLM outputs has relied on natural language processing (NLP) metrics such as BLEU (Papineni et al. 2002) and BERTScore (Zhang et al. 2020). Yet, the usefulness of these metrics in specialized domains like healthcare is limited as they require manually crafted reference responses that are rarely available in practice. In addition, it has been reported in various recent studies that they often correlate poorly with clinician judgments (Singhal et al. 2023; Van Veen et al. 2024; Liu et al. 2025), probably because these metrics focus on *overall* lexical or semantic similarities with reference texts. In contrast, human experts prioritize evaluating the factual accuracy of *specific* claims within LLM-generated responses (e.g., the statement “dry cough is a symptom of bronchiectasis”, Fig. 1a-b) stylistic variations or peripheral details.

Alternative approaches for assessing factuality involve using other LLMs or supervised models as evaluators (Chern et al. 2023; Min et al. 2023; Wang et al. 2023c). However, relying on LLMs for factuality evaluation can also be problematic, as they are susceptible to data-poisoning attacks that compromise their reliability (Alber et al. 2025), and these evaluators themselves can suffer from hallucinations. Supervised models detect factual errors by identifying patterns from historical training data, assuming future data will follow the same patterns. However, this assumption often fails in unforeseen error types, reducing their generalizability.

In recent years, various medical knowledge graphs (KGs) have been developed to represent diverse validated medical knowledge as collections of nodes (entities) and edges (relationships) (Chen et al. 2025; Poulain and Beheshti 2024; Hamza et al. 2025). They have been successfully applied to assist the generation of medical responses via retrieval-augmented generation (RAG) (Gao et al. 2023). Inspired

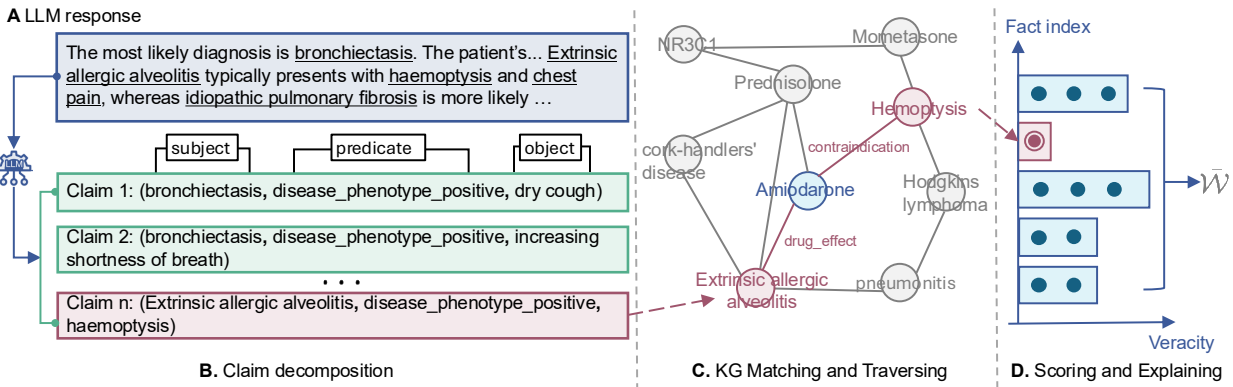


Figure 1: Overview of FAITH. (a) An example medical content generated by an LLM. (b) FAITH processes structured claims relating different medical entities in the response, which are automatically extracted by an LLM. (c) Extracted claims are matched with nodes (maroon) in a medical KG. Paths (maroon) and intermediate nodes (blue) linking the entities are identified. (d) A factuality score for each claim is computed based on the path characteristics and edge semantics. The response score is aggregated from individual claim scores.

by this, in this paper we hypothesize that these same medical KGs may also be helpful in assisting the fact-checking of LLM-generated medical contents. Yet, unlike generation tasks, fact-checking with KGs can be more complex as it requires multi-hop exploration through complex relational paths within the KG. Additionally, effective fact-checking must quantify how well extracted claims align with structured medical knowledge, making simple lookups inadequate.

To systematically assess the potential use of KG in medical fact-checking for LLMs, we developed FAITH (fact-aware evaluation of LLM-generated contents in healthcare), an unsupervised, reference-free framework based on medical KGs. FAITH first decomposes an LLM-generated response into atomic medical claims (Fig. 1b). It then uses entity resolution techniques (Christophides et al. 2021) to map the medical entities within these claims to their corresponding nodes in a medical KG (Fig. 1c). Next, FAITH traverses the KG to identify the *evidence paths* connecting the linked entities, based on which the factuality of each claim will be assessed. Finally, an overall factuality score for the entire response is aggregated from the individual claim scores, while preserving per-claim scores for nuanced interpretability (Fig. 1c-d).

We evaluated the effectiveness of FAITH and various existing fact-checking approaches on four established medical question-answering (QA) tasks with both quantitative and subjective evaluations. We also investigated whether FAITH is readily usable to improve the safety of LLM-based medical systems by employing FAITH as a thresholding mechanism. We additionally assessed FAITH’s broader utility by applying it to medical summarization and fact verification tasks and conducted ablation analysis on its core components. With these experiments, we observed several features of this KG-based fact-checker compared to conventional ones:

- It operates without reference answers or supervised training, making it highly suitable for real-world settings.
- It correlates significantly better with clinician judgments.
- It is robust to noise in the input texts, and can effectively distinguish LLMs of varying capabilities.

- It offers better interpretability than existing approaches by pinpointing which particular claim is susceptible.
- Yet, the reliability of FAITH places high demands on the KG quality, which is expected to improve as newer medical KGs come out in the future.

Related Work

Factuality Evaluation of LLMs in Medicine

Methods for evaluating the factuality of LLM-generated medical contents generally fall into two categories.

The first line of work compares LLM outputs to reference answers using NLP metrics. Examples include BLEU (Papineni et al. 2002) and ROUGE (Lin 2004) for lexical overlap, and BERTScore (Zhang et al. 2020) for contextual similarity in embedding space. To further incorporate domain-specific knowledge into evaluation, MEDCON (Yim et al. 2023) uses QuickUMLS (Soldaini and Goharian 2016)—a tool for extracting biomedical concepts—to allow for a more informed evaluation. imapScore (Wang et al. 2024) adopts the similar idea, but instead uses an LLM for medical concept extraction. However, in practice, human references are hardly available, the application of these metrics is therefore largely limited. In contrast, FAITH, is reference-free and much more flexible.

Another line of work employs supervised machine learning methods to detect factual errors. Earlier ones like FED (Mehri and Eskénazi 2020) and DEB (Sai et al. 2020) adopted BERT-like architectures. Yet, these approaches hardly generalize beyond their original training data. Later works used more advanced GPT-series models for evaluation, e.g., GPTScore (Fu et al. 2024), G-Eval (Liu et al. 2023), and GPT-judge (Lin, Hilton, and Evans 2022). More recently, LLM-Eval (Lin and Chen 2023) adopts a multi-aspect evaluation schema for more comprehensive evaluation. Specifically in the medical domain, Vladika, Hacajová, and Matthes leverage multi-turn interactions with LLMs to reach a binary decision on a single claim. However, these LLM judges can also suffer from hallucinations. In addition, all these model-based approaches

lack sufficient explainability to the evaluation. In contrast, FAITH is unsupervised, grounded on medical facts, and offers inherent explainability to pinpoint the errors.

Also related are those works regarding hallucination detection of LLMs (Ji et al. 2023; Zhang et al. 2023), which is a broader topic. Our work specifically targets one type of hallucination that stems from factual errors, and complements research on other types, e.g., confabulations (Farquhar et al. 2024) or instruction misalignment (Zeng et al. 2024). Combining FAITH with other hallucination-detection methods can enable more systematic safety oversight.

Knowledge Graph-Based Fact-Checking

In general domains, external knowledge base like KGs, have been widely used for fact-checking (Gad-Elrab et al. 2019; Ciampaglia et al. 2015; Shiralkar et al. 2017; Fionda and Pirrò 2018; Shi and Weninger 2016; Kim et al. 2023; Dammu et al. 2024; Huang et al. 2025; Yu et al. 2025; Huang, Zhou, and Li 2024). Early approaches assessed the “quality” of paths linking entities within each claim. For instance, KL (Ciampaglia et al. 2015) suggests that paths involving generic entities provide weaker support and uses node degrees for evaluation. KL-REL (Shiralkar et al. 2017) extends this by considering the semantic similarity of edge labels in the path to the predicate in the original claim. Other methods like TransE (Bordes et al. 2013) use global representations of entities and relationships to better capture hierarchical structures in KGs. These methods often assume that a single fact is being checked each time and that its entities are directly available in the KG, making them not directly applicable to free-text outputs.

The FAITH Evaluation Framework

Preliminaries

Claims of fact. A *claim of fact* is a triplet $t = (s, r, o)$, where s, o denote the subject and object, respectively, and r is the predicate (i.e., relation between s and o). For example, the triplet (Barack Obama, *was_born_in*, Hawaii) expresses the claim that Barack Obama was born in Hawaii.

Knowledge graphs (KGs). A *knowledge graph* $\mathcal{G} = (\mathcal{E}, \mathcal{R})$ is composed of a set of entities (nodes) \mathcal{E} and relations (edges) \mathcal{R} among them. It can be regarded as a collection of factual claims defined as $\mathcal{T} = \mathcal{E} \times \mathcal{R} \times \mathcal{E}$.

Knowledge paths. A *knowledge path* p between entity e_1 and e_k is a sequence of alternating entities and relations, denoted as $p = (e_1, r_1, e_2, r_2, \dots, r_{k-1}, e_k)$, where $e_i \in \mathcal{E} (1 \leq i \leq k)$, $r_i \in \mathcal{R} (1 \leq i < k)$, and $(e_i, r_i, e_{i+1}) \in \mathcal{T}$ for $1 \leq i < k$. This path consists of k entities and $k - 1$ relations. The path length is defined as the number of relations it contains, $k - 1$.

Computational fact checking. Given an LLM-generated response \mathcal{D} containing a set of medical factual claims $\{t_i\}_{i=1}^n$ and a KG \mathcal{G} , the goal of *computational fact checking* is to determine the factuality of each claim triplet $t_i = (s_i, r_i, o_i)$ in \mathcal{D} by leveraging the knowledge paths p_i in \mathcal{G} .

Key Modules of FAITH

As shown in Fig. 1, FAITH comprises 4 key modules:

Medical claim extraction. To extract medical claims $\{t_i\}_{i=1}^n$ from an LLM response \mathcal{D} , we use GPT-4o for this purpose

due to its strong extraction capability (Dunn et al. 2022; Polak and Morgan 2024). Specifically, we use a multi-phase prompting strategy (Appendix B.1) that first identifies all medical entities in \mathcal{D} and then determines relationships among them. To enhance performance, we employ in-context learning with 5 clinician-crafted examples. We also adopt a multi-round prompting (Kuratov et al. 2024; Edge et al. 2024) to increase recall, whilst a redundancy-inducing prompt is used to let the model critically reanalyze the extracted claims with suspicion to mitigate hallucination (Polak and Morgan 2024).

Medical entity matching. In classic fact-checking literature (Ciampaglia et al. 2015; Shiralkar et al. 2017), benchmarking datasets are often generated from the KG used in the algorithm. Thus, entities in this data can always match with KG nodes. Yet, in real-world applications, entities in the LLM responses often cannot directly match with nodes in the KG (e.g., ‘haemoptysis’ and ‘Hemoptysis’ in Fig. 1c). To resolve these discrepancies, we use entity resolution techniques (Binette and Steorts 2022) via the Unified Medical Language System (UMLS) (Bodenreider 2004). UMLS is a comprehensive ontology linking synonymous medical terms from over 200 different biomedical vocabulary sources via concept unique identifiers (CUIs). We use the UMLS API to map entities from both claims and nodes in \mathcal{G} to standardized CUIs. This provides a strong foundation of reliable, consensual knowledge, which is essential for a fact-checking system. Claims with unmatched entities are labeled as ‘unverifiable’ and excluded from further evaluation. This represents a conservative design choice, ensuring that FAITH only scores claims for which it has a high-confidence knowledge base.

KG traversal and factual evaluation. After entity matching, we traverse the KG to extract the knowledge paths $\mathcal{P} = \{p_i\}_{i=1}^n$ connecting the resulting nodes, where each path p is a sequence $(e_1, r_1, e_2, r_2, \dots, r_{k-1}, e_k)$. Given the vast scale and dense connectivity of medical KGs, an exhaustive search of all possible paths between two entities is computationally intractable. Therefore, we constrain our search to the shortest path, which represents the most direct and salient connection. In addition to computational feasibility, the shortest paths are also considered optimal for maximizing the information content of the connection (Ciampaglia et al. 2015). The obtained path p_i can then serve as evidence to evaluate the factuality of the corresponding medical claim t_i .

It is often considered that longer paths indicate weaker factuality (Ciampaglia et al. 2015; Shiralkar et al. 2017), but path length alone is insufficient. This is because the semantics of relations between entities can significantly affect the factuality of the claim. For example, the relations ‘*indication*’ and ‘*contraindication*’ have totally different meanings when directly linking two entities. To include this, we introduce a relation semantic similarity $\mathcal{S}(p, t)$ to assess the congruence between the semantics of the knowledge path p and the medical claim t . We use the mean cosine similarity between the embeddings of the relations in p and the relation in t :

$$\mathcal{S}(p, t) = \frac{1}{k-1} \times \sum_{i=1}^{k-1} \cos_sim(r_i, \hat{r}), \quad (1)$$

where $k - 1$ is the length of p , r_i represents the i -th relation in p , and \hat{r} is the predicate in t .

However, this is still inadequate to fully capture the factuality of the medical claim t . Ciampaglia et al. (2015) suggest that paths involving generic entities provide weaker support to a claim. Moreover, Shiralkar et al. (2017) emphasize the importance of considering the co-occurrence of each relationship r_i in p with t 's predicate to enhance the accuracy of factuality assessment. Building on these insights, FAITH incorporates the entity centrality (via PageRank, PR (Brin and Page 1998)) and relationship co-occurrence ($u(r_i, \hat{r})$) derived from contracted line graphs. Putting these together, the factuality score for a given medical claim t related to a knowledge path p is given by:

$$\mathcal{W}(p, t) = \mathcal{S}(p, t) \left[\sum_{i=2}^{k-1} \frac{e^{\alpha PR(e_i)}}{u(r_{i-1}, \hat{r})} + \frac{1}{u(r_{k-1}, \hat{r})} \right]^{-1}, \quad (2)$$

where e_i and r_i are the i -th entity and relation in p , respectively, and α is a scaling constant set to 100 in this study (see Appendix B.2). The term $u(r_i, \hat{r})$ is calculated as the co-occurrence between two relations based on contracted line graph, with higher values indicating tighter relations.

Scoring and interpretation. Finally, the overall FAITH score for the whole LLM-generated response \mathcal{D} , which contains the set of claims $\{t_i\}_{i=1}^n$ is aggregated from the individual factuality scores for each claim involved:

$$\bar{\mathcal{W}}(\mathcal{P}, \mathcal{D}) = \frac{1}{n} \times \sum_{p_i \in \mathcal{P}, t_i \in \mathcal{D}} \mathcal{W}(p_i, t_i). \quad (3)$$

FAITH generates factuality scores in the range of $[-1, 1]$, where positive values indicate alignment with established medical knowledge while negative values suggest contradictions. The absolute value reflects how strongly entities in a claim are interconnected. Our scoring system ensures interpretability by explicitly linking each assessment to the KG paths examined (see examples in Appendix D.1).

Experimental Setup

Baseline methods. We compare FAITH against the following baselines: **► NLP metrics:** This includes BLEU-4 (Papineni et al. 2002) and ROUGE-L (Lin 2004). We also included BERTScore (Zhang et al. 2020) which leverages contextual embedding model (here voyage-3-large, a top-performer on MTEB leaderboard (Muennighoff et al. 2023)) to evaluate the semantic similarity between reference and input. Additionally, we considered MEDCON (Yim et al. 2023) as it is specifically designed for medical scenarios. **► Computational fact-checking methods:** These include 3 established methods: KL (Ciampaglia et al. 2015) and KL-REL (Shiralkar et al. 2017), and TransE (Bordes et al. 2013). **► LLM-based method:** We consider FActScore (Min et al. 2023) and imapScore (Wang et al. 2024) to represent this category. Note that as KL, KL-REL, and TransE cannot be directly applied to free texts, we preprocess their inputs with our entity matching module. NLP metrics and imapScore are provided with reference responses for evaluation.

LLMs. To compare the efficacy of our proposed approach with the baselines, we considered 5 representative LLMs. This includes 2 established proprietary models, GPT-4o-mini

and GPT-4o (Hurst et al. 2024). We also included two open-source models, Llama 3-8B and Llama 3.1-8B (Dubey et al. 2024). We additionally included a specialized model, OpenBioLLM, which is fine-tuned on medical texts and specifically designed to answer medical questions in an informative way.

Medical tasks. To assess FAITH's efficacy, we focused on medical QA tasks as they purely demand a high factual accuracy on medical knowledge. Other tasks, such as clinical decision making, further require advanced reasoning capabilities. While these are useful for evaluating the overall medical capabilities of LLMs, they can introduce additional confounding factors to our evaluation, which focuses on factuality. In this spirit, we selected 4 established datasets: *i*) MedQA (Jin et al. 2020) *ii*) MMLU (Hendrycks et al. 2021), *iii*) MS-AKT, and *iv*) LiveQA (Abacha et al. 2017). These span both multiple-choice (MedQA, MMLU, MS-AKT) and open-ended questions (LiveQA), with varying levels of complexity and coverage. For more details, see Appendix C.3.

Medical KG. While many medical KGs exist (Lu et al. 2025; Bodenreider 2004; Loudon 2020; Cheung et al. 1996), many of them are tailored to focus on specific medical problems, whereas our fact-checking would desire an encyclopedic KG that covers a wide range of medical information (see discussion in Appendix B.3). To this end, we consider the UMLS (Bodenreider 2004) 2025AA version for this paper. It integrates knowledge from over 200 standard biomedical vocabulary sources, (e.g., MeSH, SNOMED CT and Human Phenotype Ontology), incorporating more than 23 million relationships among 3.4 million concepts.

Experimental Results

Reliability and Validity of FAITH

An effective evaluation metric for LLM-generated responses in healthcare should satisfy several key desiderata from practitioners. To be specific, it should: *i*) clearly distinguish responses produced by LLMs with different levels of capabilities; *ii*) be robust to the variance or noise in the responses that are not relevant to the evaluation; and *iii*) align with the consensus of experienced clinicians. We thus evaluate FAITH on these 3 dimensions.

Distinguishing LLM capabilities. We first probed FAITH's ability to differentiate 5 LLMs with varying medical capabilities (Zhang et al. 2024; Boggavarapu et al. 2024; Ke et al. 2025; Chen et al. 2024) by evaluating their responses on 4 medical QA tasks. We find that metrics such as ROUGE-L, BERTScore, and TransE on average rated all 5 models similarly (paired t -test $p > 0.05$; Fig. 2a, b, e, and g). While other metrics like BLEU-4, MEDCON, and FActScore achieved statistically significant differences between some of the models, this discrepancy does not persist across all models and tasks. For example, FActScore fails to distinguish between Llama 3.1 and GPT-4o-mini on MedQA (Fig. 2a). In contrast, FAITH consistently distinguished responses from all 5 LLMs ($p < 0.004$). This indicates that the approach is sensitive enough to capture meaningful differences in the factual quality of outputs from models with varying capabilities.

Robustness to variance. We then test FAITH's robustness to variations in response phrasing. To do this, we used GPT-4o

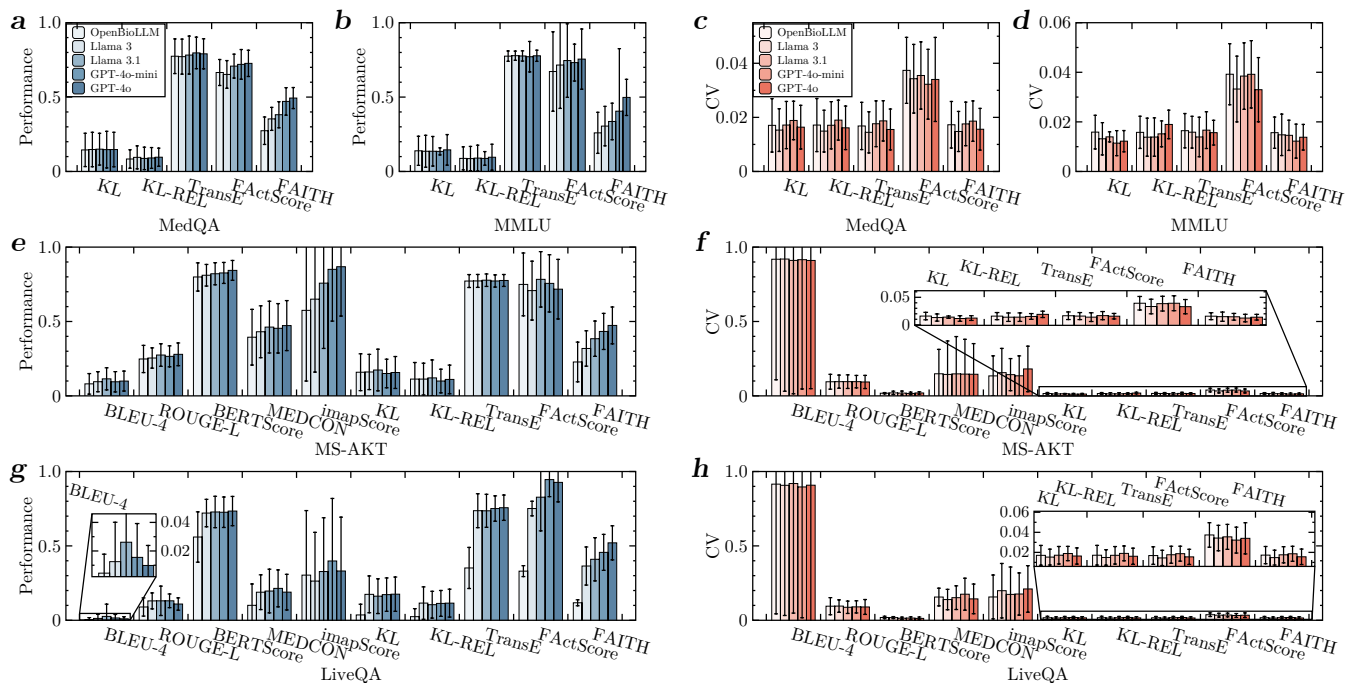


Figure 2: FAITH effectively distinguishes between LLMs and is robust to noise. This figure shows the mean factuality scores assigned by FAITH and baseline metrics to responses from five LLMs on four datasets: MedQA (a), MMLU (b), MS-AKT (e), and LiveQA (g). A reliable metric should assign distinguishable scores across different models. Panels c, d, f, and h display the corresponding coefficients of variation (CV) of these scores under noisy conditions.

to generate 10 paraphrased versions for each response from the 5 LLMs (Appendix C.5). The paraphrasing instructions emphasized modifying sentence structures, pronouns, and prepositions while preserving medical terminology and relationships. An ideal metric should yield low variance across such paraphrased responses. We selected the coefficient of variation (CV) for this purpose, as it offers a measure of relative variability. Across all scenarios, FAITH exhibits consistently low CV (Fig. 2c, d, f, and h), with the average \pm s.d. being 0.014 ± 0.005 . In contrast, metrics like BLEU-4 showed much higher variability (0.910 ± 0.862). This suggests that FAITH is more robust to the variations in phrasing.

Correlation with clinician judgments. Finally, to determine the alignment of FAITH with human clinicians, we recruited 20 UK-based clinicians to manually grade the responses generated by LLMs. Each clinician holds a medical degree and possesses at least 5 years of clinical practice experience. In close consultation with these experts, we identified 3 key evaluation dimensions: *i*) factuality, *ii*) relevance, and *iii*) potential harm. Each dimension was rated on a 5-point Likert scale (with 1 being the lowest and 5 the highest; Appendix C.6), and an additional overall assessment also used the same scale. Following these criteria, each clinician assessed 5 responses, which were randomly drawn from a candidate question pool of 16 questions from MS-AKT dataset. These responses could either originate from Llama 3.1, GPT-4o, or official answers. To ensure rating consistency, each response was evaluated by at least two clinicians, with Cohen’s

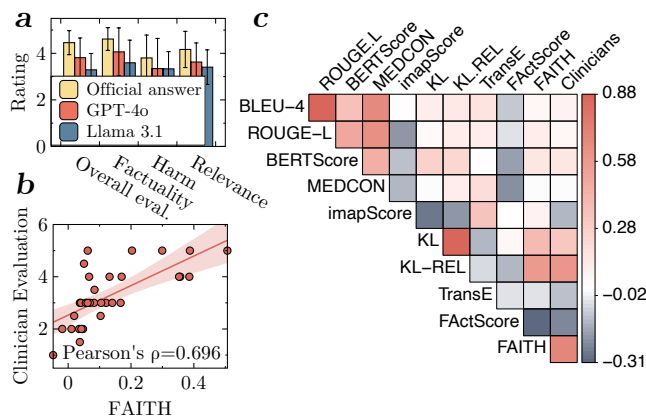


Figure 3: FAITH exhibits highest correlation with clinician judgments. a, Clinician evaluation scores for answers to 16 questions generated by GPT-4o, Llama 3.1, and official answers. b, Scatter plots correlating clinician judgments with scores from FAITH. Linear regression fits are shown with 95% confidence intervals. c, Pearson correlation between clinicians and scores from FAITH and various baselines.

$\kappa = 0.64$, indicating substantial inter-rater agreement.

As shown in Fig. 3a, clinician assessments consistently ranked official answers highest across all evaluation axes, followed by GPT-4o, and then Llama 3.1. To quantify alignment between automated metrics and these human judgments,

we computed Pearson’s correlation coefficient ρ . The results show that FAITH scores achieve a considerably higher correlation with clinicians (Pearson’s $\rho = 0.696$; Fig. 3b, c) compared to all baselines. This provides strong evidence for the effectiveness of the KG-based evaluation approach. Notably, baselines like BLEU-4 only has $\rho = 0.081$, consistent with previous findings (Van Veen et al. 2024).

Explainability of FAITH

We evaluate FAITH’s interpretability with 2 aspects: *i*): the faithfulness of its explanations to medical consensus, and *ii*): the utility of these explanations for analyzing LLM behaviors. **Faithfulness to medical consensus.** For the first aspect, we additionally asked clinicians to annotate the one most incorrect medical claim in each justification they had previously rated the factuality score lower than 5. These annotations were manually converted to fact triplets to be comparable with the factuality explanations generated by FAITH. We evaluated FAITH’s ability to pinpoint these specific annotations, achieving a precision of 0.65, recall of 0.59, and F1 score of 0.62 (Fig. 4a). It is important to consider that the selection of the most incorrect claim by clinicians can involve subjective considerations (e.g., specific expertise, perceived claim importance). We therefore evaluated its performance in localizing these claims and found that in 83.6% of these justifications, the erroneous statement pinpointed by clinicians was ranked by FAITH among the top-5 lowest-scoring. These findings suggest that FAITH’s explanations could be faithful to current medical consensus, offering valuable, granular insights into response inaccuracies.

Utility in analyzing LLM limitations. Such explanations offered by FAITH can be potentially utilized to offer insights into limitations of LLMs. To show this, we investigated the top-5 most frequent edge types in the KG that are associated with the most incorrect claims made in GPT-4o’s responses, as identified by FAITH. As shown in Fig. 4b, nearly half of the identified inaccuracies involve phenotypical features of diseases. Such errors are particularly prominent when LLMs attempt to determine a disease from patient signs, which underscores their current challenges in the rigorous clinical task of differentiating between numerous potential pathologies for a given initial complaint (Hager et al. 2024). Furthermore, frequent errors in establishing causal relationships between diseases and exposures highlight LLM tendencies to overstate associations or assert causality without sufficient evidence (Griot et al. 2025; Pawitan and Holmes 2024). This granular error typology (Appendix D.3), is vital for guiding target LLM refinement in complex medical applications.

Practical Utility of FAITH for Safeguarding

Beyond evaluation, FAITH can be used to improve the factuality and safety of LLM-generated medical responses. We show this in two intervention scenarios. In the first one, we use FAITH scores below a set threshold to activate a Reject-to-Answer (RTA) protocol, filtering out responses with low factuality (Singhal et al. 2023; Wen et al. 2024). In the second one, we augment the LLMs with a RAG module (Gao et al. 2023) that retrieves relevant information from the medical KG when potential factual errors are detected. Model




Figure 4: Explainability of FAITH via faithful error identification and LLM limitation analysis. a, Alignment between clinician-identified incorrect claims and FAITH’s lowest-scoring claims in LLM responses. b, Distribution of the top-5 most frequent KG relation types linked to incorrect claims in GPT-4o’s responses, as identified by FAITH.

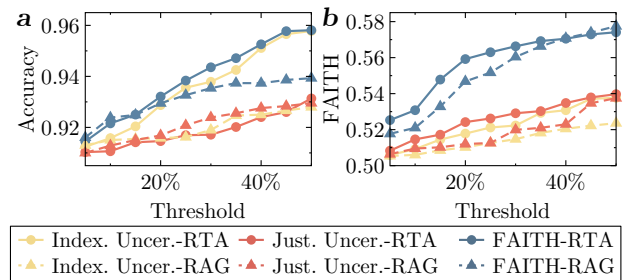


Figure 5: FAITH enhances LLM factuality via selective intervention. GPT-4o performance on MedQA using Reject-to-Answer (RTA) or Retrieval-Augmented Generation (RAG). Interventions triggered by either FAITH scores or a model uncertainty baseline. Metrics plotted against percentage of responses selected for intervention.

uncertainty served as a baseline triggering mechanism (see Appendix B.4). For both scenarios, intervention thresholds were determined by percentiles of the FAITH or uncertainty score distribution in the dataset (ranging from 5% to 50%).

Our findings show that by applying the RTA protocol, both the accuracy of answers and the factuality of the justification generated by GPT-4o could be significantly improved (Fig. 5). While uncertainty-based thresholding could also lead to boosted results, the improvements are not as significant as those achieved by FAITH. Likewise, RAG could also enhance both metrics, particularly the factuality of the justification. Instead of employing RAG for all questions, the use of FAITH as a thresholding mechanism can offer a more cost-effective solution in practice. This advantage also persists if we consider other LLMs (Fig. A16).

Broader Applicability of FAITH

To interrogate FAITH’s utility beyond medical QA, we further assessed its performance on medical summarization (Shaib et al. 2023; Tang et al. 2023; Thawakar et al. 2024; Van Veen et al. 2024) and medical fact verification (MFV) (Vladika, Schneider, and Matthes 2024; Mohr, Wüthrich, and Klinger 2022) (details available in Appendix C.4). For medical summarization, we used the FactPICO benchmark (Joseph et al. 2024), which comprises expert

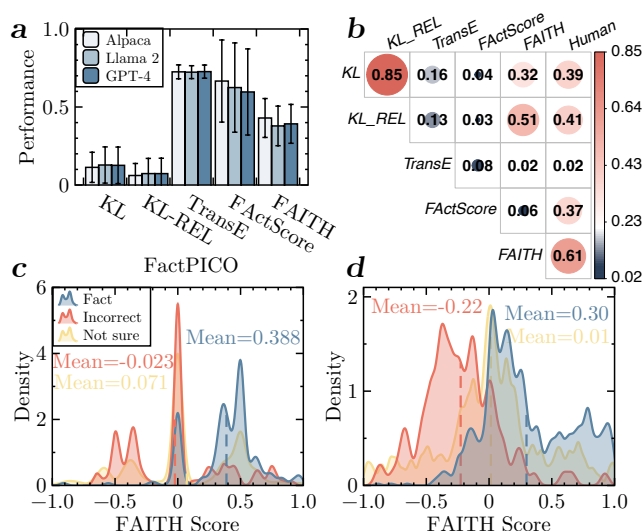


Figure 6: Broad applicability of FAITH to medical summarization and fact verification (MFV). a, Factuality scores assigned by FAITH to LLM-created medical summaries from the FactPICO benchmark, showing alignment with model performance hierarchies. b, Pearson correlations of FAITH and baselines with expert judgements on FactPICO. c, d, Distribution of FAITH scores for claims of known veracity from the HealthFC (c) and BEAR-FACT (d) MFV benchmarks, indicating differentiation between true and false statements.

factuality assessments of 345 LLM-generated summaries of 115 randomized controlled trials (RCTs) across the Alpaca, Llama-2 and GPT-4 models. Results shown in Fig. 6a indicate that FAITH aligns closely with the expert-rated model rankings (Alpaca > GPT-4 > Llama-2). More importantly, FAITH achieves the highest Pearson’s correlation with expert judgments ($\rho = 0.61$; Fig. 6b). For MFV task, we assessed FAITH on the HealthFC (Vladika, Schneider, and Matthes 2024) and BEAR-FACT (Wühlrl et al. 2024) datasets, which consist of 750 and 300 medical claims, respectively. The results show that FAITH effectively distinguished true from false claims (Fig. 6c, d, Fig. A17, and Fig. A18).

Ablation Studies

Sensitivity to KG Choice and Integrity

Compatibility with different KGs. To validate this, we integrated FAITH with other established KGs of varying scales, PrimeKG (Chandak, Huang, and Zitnik 2023) and OGBL-biokg (Hu et al. 2020) (Appendix B.3), and repeated our core experiments. The results show that while relative ranking of LLMs remains consistent, the performance will be affected by different KGs (Fig. A11).

Robustness against KG noise. Given the real-world KGs inevitably contain noise (e.g., missing/erroneous triplets), we assessed FAITH’s robustness by perturbing the UMLS. We applied 3 types of significant (Albert, Jeong, and Barabási 2000) noise: *i*) 20% random edge deletion, *ii*) 20% random node deletion, and *iii*) 20% random noisy edge insertion.

The results in Fig. A12 showed that deleting entire nodes or inserting noisy edges cause performance degradation. This implies that high-quality KGs remain a crucial prerequisite.

Reliability of Medical Claim Extraction Module

To assess the medical claim extraction pipeline in FAITH, we compare various extractor implementations against QuickUMLS (Soldaini and Goharian 2016) on the BioRED (Luo et al. 2022) and DDI13 (Segura-Bedmar, Martínez, and Herrero-Zazo 2013) datasets (details in Appendix C.7).

Impact of prompting strategies. We first tested how our multi-round and critical analysis prompting affect extraction performance. We thus implemented the LLM extractor in FAITH with 4 prompting strategies: **► Base prompt:** a basic extraction prompt without advanced techniques. **► Base + critical:** the base prompt with critical analysis. **► Base + multi:** the base prompt with multi-round conversation. **► FAITH:** FAITH’s full prompting strategy. As shown in Fig. A13, all LLM-based variants surpass QuickUMLS on both datasets. Multi-round conversation mainly boosts recall, while the critical analysis prompts helps ensure the precision, making their combination most effective.

Sensitivity to LLM choice. We then examine the sensitivity of the claim extraction performance to the specific choice of LLM. For this, we replaced the GPT-4o model with 4 other LLMs and repeated the experiments on BioRED and DDI13. The results in Fig. A14 showed that while changing to cheap open-source models like Llama 3 and Llama 3.1 leads to a slight decrease, all LLM extractors still yield considerable improvements over QuickUMLS. On the other hand, the performance gain of GPT-4o over GPT-4o-mini is marginal, so GPT-4o model is already well-suited for the task.

Discussion and Conclusion

This paper investigated the reliability of medical KGs as a foundation for evaluating LLM factuality in healthcare. Through our framework, we demonstrate that leveraging KGs is a powerful and reliable approach for factuality evaluation.

However, our work also illuminates the some limitations of this approach. The first one is its reliance on the KG’s quality and coverage. Factual claims involving knowledge absent from the KG cannot be verified, which can lead to potential false negatives. Secondly, the entire pipeline is dependent on the performance of the upstream medical claim extraction module. While our LLM-based extractor proved effective, errors at this stage can propagate through the system. Future work should focus on mitigating these limitations, perhaps by integrating multiple KGs or developing methods to handle claims that are out-of-scope for the KG.

Social impact. FAITH enhances the safety and reliability of LLMs (Zhou et al. 2024; Zhou, Li, and Min 2022b,a) by automatically assessing the factuality of their responses. It serves as a crucial “guardrail” against misinformation, fostering trustworthy adoption by clinicians and regulators while guiding developers with targeted feedback for model improvement. The core approach is also adaptable for fact-checking in other critical domains like law (Cui et al. 2023; Hamdani et al. 2024) and finance (Xie et al. 2023; Kang and Liu 2023) where factuality is paramount.

Acknowledgements

We sincerely thank all the reviewers for their encouraging and constructive feedback. This work was supported by the UKRI Future Leaders Fellowship under Grant MR/S017062/1 and MR/X011135/1; in part by NSFC under Grant 62376056 and 62076056; in part by the Royal Society Faraday Discovery Fellowship (FDF/S2/251014), BBSRC Transformative Research Technologies (UKRI1875), Royal Society International Exchanges Award (IES/R3/243136), Kan Tong Po Fellowship (KTP/R1/231017); and the Amazon Research Award and Alan Turing Fellowship.

References

- Abacha, A. B.; Agichtein, E.; Pinter, Y.; and Demner-Fushman, D. 2017. Overview of the medical question answering task at TREC 2017 LiveQA. In *TREC*, 1–12.
- Alber, D. A.; Yang, Z.; Alyakin, A.; Yang, E.; Rai, S.; Valiani, A. A.; Zhang, J.; Rosenbaum, G. R.; Amend-Thomas, A. K.; Kurland, D. B.; Kremer, C. M.; Eremiev, A.; Negash, B.; Wiggan, D. D.; Nakatsuka, M. A.; Sangwon, K. L.; Neifert, S. N.; Khan, H. A.; Save, A. V.; Palla, A.; Grin, E. A.; Hedman, M.; Nasir-Moin, M.; Liu, X. C.; Jiang, L. Y.; Mankowski, M. A.; Segev, D. L.; Aphinyanaphongs, Y.; Rina, H. A.; Golfinos, J. G.; Orringer, D. A.; Kondziolka, D.; and Oermann, E. K. 2025. Medical large language models are vulnerable to data-poisoning attacks. *Nat. Med.*
- Albert, R.; Jeong, H.; and Barabási, A.-L. 2000. Error and attack tolerance of complex networks. *nature*, 406(6794): 378–382.
- Augenstein, I.; Baldwin, T.; Cha, M.; Chakraborty, T.; Ciampaglia, G. L.; Corney, D. P. A.; DiResta, R.; Ferrara, E.; Hale, S.; Halevy, A. Y.; Hovy, E. H.; Ji, H.; Menczer, F.; Míguez, R.; Nakov, P.; Scheufele, D.; Sharma, S.; and Zagni, G. 2024. Factuality challenges in the era of large language models and opportunities for fact-checking. *Nat. Mac. Intell.*, 6(8): 852–863.
- Binette, O.; and Steorts, R. C. 2022. (Almost) all of entity resolution. *Science Advances*, 8(12): eabi8021.
- Bodenreider, O. 2004. The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Res.*, 32(suppl_1): 267–270.
- Boggavarapu, L.; Srivastava, V.; Varanasi, A. M.; Lu, Y.; and Bhaumik, R. 2024. Evaluating Enhanced LLMs for Precise Mental Health Diagnosis from Clinical Notes. *medRxiv*, 2024–12.
- Bordes, A.; Usunier, N.; García-Durán, A.; Weston, J.; and Yakhnenko, O. 2013. Translating Embeddings for Modeling Multi-relational Data. In *NIPS'13: Proc. of the Advances in Neural Information Processing Systems 26*, 2787–2795.
- Brin, S.; and Page, L. 1998. The Anatomy of a Large-Scale Hypertextual Web Search Engine. *Comput. Networks*, 30(1-7): 107–117.
- Chandak, P.; Huang, K.; and Zitnik, M. 2023. Building a knowledge graph to enable precision medicine. *Scientific Data*, 10(1): 67.
- Chen, H.; Shen, X.; Wang, J.; Wang, Z.; Lv, Q.; He, J.; Wu, R.; Wu, F.; and Ye, J. 2025. Knowledge Graph Finetuning Enhances Knowledge Manipulation in Large Language Models. In *ICLR'25: The Thirteenth International Conference on Learning Representations*.
- Chen, S.; Gao, M.; Sasse, K.; Hartvigsen, T.; Anthony, B.; Fan, L.; Aerts, H. J.; Gallifant, J.; and Bitterman, D. S. 2024. The Risks of Medical Misinformation Generation in Large Language Models. Available at SSRN 5020664.
- Chern, I.; Chern, S.; Chen, S.; Yuan, W.; Feng, K.; Zhou, C.; He, J.; Neubig, G.; and Liu, P. 2023. FacTool: Factuality Detection in Generative AI - A Tool Augmented Framework for Multi-Task and Multi-Domain Scenarios. *CoRR*, abs/2307.13528.
- Cheung, K.-H.; Nadkarni, P.; Silverstein, S.; Kidd, J. R.; Pakstis, A. J.; Miller, P.; and Kidd, K. K. 1996. PhenoDB: an integrated client/server database for linkage and population genetics. *Computers and biomedical research*, 29(4): 327–337.
- Christophides, V.; Efthymiou, V.; Palpanas, T.; Papadakis, G.; and Stefanidis, K. 2021. An Overview of End-to-End Entity Resolution for Big Data. *ACM Comput. Surv.*, 53(6): 127:1–127:42.
- Ciampaglia, G. L.; Shiralkar, P.; Rocha, L. M.; Bollen, J.; Menczer, F.; and Flammini, A. 2015. Computational fact checking from knowledge networks. *PLoS one*, 10(6): e0128193.
- Cui, J.; Li, Z.; Yan, Y.; Chen, B.; and Yuan, L. 2023. ChatLaw: Open-Source Legal Large Language Model with Integrated External Knowledge Bases. *CoRR*, abs/2306.16092.
- Dammu, P. P. S.; Naidu, H.; Dewan, M.; Kim, Y.; Roosta, T.; Chadha, A.; and Shah, C. 2024. ClaimVer: Explainable Claim-Level Verification and Evidence Attribution of Text Through Knowledge Graphs. In *EMNLP'24: Findings of the Association for Computational Linguistics: EMNLP 2024*, 13613–13627. Association for Computational Linguistics.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Dunn, A.; Dagdelen, J.; Walker, N.; Lee, S.; Rosen, A. S.; Ceder, G.; Persson, K. A.; and Jain, A. 2022. Structured information extraction from complex scientific text with fine-tuned large language models. *CoRR*, abs/2212.05238.
- Edge, D.; Trinh, H.; Cheng, N.; Bradley, J.; Chao, A.; Mody, A.; Truitt, S.; and Larson, J. 2024. From Local to Global: A Graph RAG Approach to Query-Focused Summarization. *CoRR*, abs/2404.16130.
- Farquhar, S.; Kossen, J.; Kuhn, L.; and Gal, Y. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017): 625–630.
- Fionda, V.; and Pirrò, G. 2018. Fact Checking via Evidence Patterns. In *IJCAI'18: Proc. of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, 3755–3761. ijcai.org.

- Fu, J.; Ng, S.; Jiang, Z.; and Liu, P. 2024. GPTScore: Evaluate as You Desire. In *NAACL'24: Proc. of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics*, 6556–6576. Association for Computational Linguistics.
- Gad-Elrab, M. H.; Stepanova, D.; Urbani, J.; and Weikum, G. 2019. ExFaKT: A Framework for Explaining Facts over Knowledge Graphs and Text. In *WSDM'19: Proc. of the Twelfth ACM International Conference on Web Search and Data Mining*, 87–95. ACM.
- Gao, Y.; Xiong, Y.; Gao, X.; Jia, K.; Pan, J.; Bi, Y.; Dai, Y.; Sun, J.; Guo, Q.; Wang, M.; and Wang, H. 2023. Retrieval-Augmented Generation for Large Language Models: A Survey. *CoRR*, abs/2312.10997.
- Griot, M.; Hemptinne, C.; Vanderdonckt, J.; and Yuksel, D. 2025. Large language models lack essential metacognition for reliable medical reasoning. *Nature communications*, 16(1): 642.
- Hager, P.; Jungmann, F.; Holland, R.; Bhagat, K.; Hubrecht, I.; Knauer, M.; Vielhauer, J.; Makowski, M.; Braren, R.; Kaissis, G.; et al. 2024. Evaluation and mitigation of the limitations of large language models in clinical decision-making. *Nature medicine*, 30(9): 2613–2622.
- Hamdani, R. E.; Bonald, T.; Malliaros, F. D.; Holzenberger, N.; and Suchanek, F. M. 2024. The Factuality of Large Language Models in the Legal Domain. In *CIKM'24: Proc. of the 33rd ACM International Conference on Information and Knowledge Management*, 3741–3746. ACM.
- Hamza, A.; Abdullah, A.; Ahn, Y. H.; Lee, S.; and Kim, S. T. 2025. LLaVA Needs More Knowledge: Retrieval Augmented Natural Language Generation with Knowledge Graph for Explaining Thoracic Pathologies. In *AAAI'25: Sponsored by the Association for the Advancement of Artificial Intelligence*, 3311–3319. AAAI Press.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021. Measuring Massive Multi-task Language Understanding. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Hu, W.; Fey, M.; Zitnik, M.; Dong, Y.; Ren, H.; Liu, B.; Catasta, M.; and Leskovec, J. 2020. Open Graph Benchmark: Datasets for Machine Learning on Graphs. In *NeurIPS'20: Advances in Neural Information Processing Systems*.
- Huang, M.; Zhou, S.; Chen, Y.; and Li, K. 2025. Conversational Exploration of Literature Landscape with LitChat. In *IJCAI'25: Proc. of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, 11058–11061. ijcai.org.
- Huang, M.; Zhou, S.; and Li, K. 2024. Mapping Literature Landscapes with Data-Driven Discovery: A Case Study on MOEA/D. *CoRR*, abs/2404.14228.
- Hurst, A.; Lerer, A.; Goucher, A. P.; and et al. 2024. GPT-4o System Card. *CoRR*, abs/2410.21276.
- Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y.; Madotto, A.; and Fung, P. 2023. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.*, 55(12): 248:1–248:38.
- Jin, D.; Pan, E.; Oufattole, N.; Weng, W.; Fang, H.; and Szolovits, P. 2020. What Disease does this Patient Have? A Large-scale Open Domain Question Answering Dataset from Medical Exams. *CoRR*, abs/2009.13081.
- Joseph, S.; Chen, L.; Trienes, J.; Göke, H. L.; Coers, M.; Xu, W.; Wallace, B. C.; and Li, J. J. 2024. FactPICO: Factuality Evaluation for Plain Language Summarization of Medical Evidence. In *ACL'24: Proc. of the 62nd Annual Meeting of the Association for Computational Linguistics*, 8437–8464. Association for Computational Linguistics.
- Kang, H.; and Liu, X. 2023. Deficiency of Large Language Models in Finance: An Empirical Examination of Hallucination. *CoRR*, abs/2311.15548.
- Ke, Y.; Jin, L.; Elangovan, K.; Abdullah, H. R.; Liu, N.; Sia, A. T. H.; Soh, C. R.; Tung, J. Y. M.; Ong, J. C. L.; Kuo, C. F.; Wu, S.; Kovacheva, V. P.; and Ting, D. S. W. 2025. Retrieval augmented generation for 10 large language models and its generalizability in assessing medical fitness. *npj Digit. Medicine*, 8(1).
- Kim, J.; Park, S.; Kwon, Y.; Jo, Y.; Thorne, J.; and Choi, E. 2023. FactKG: Fact Verification via Reasoning on Knowledge Graphs. In *ACL'23: Proc. of the 61st Annual Meeting of the Association for Computational Linguistics ACL 2023*, 16190–16206. Association for Computational Linguistics.
- Kurатов, Y.; Bulatov, A.; Anokhin, P.; Sorokin, D.; Sorokin, A.; and Burtsev, M. 2024. In search of needles in a 10m haystack: Recurrent memory finds what llms miss. *arXiv preprint arXiv:2402.10790*.
- Lee, P.; Bubeck, S.; and Petro, J. 2023. Benefits, Limits, and Risks of GPT-4 as an AI Chatbot for Medicine. *New England Journal of Medicine*, 388(13): 1233–1239.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Lin, S.; Hilton, J.; and Evans, O. 2022. TruthfulQA: Measuring How Models Mimic Human Falsehoods. In *ACL'22: Proc. of the 60th Annual Meeting of the Association for Computational Linguistics*, 3214–3252. Association for Computational Linguistics.
- Lin, Y.; and Chen, Y. 2023. LLM-Eval: Unified Multi-Dimensional Automatic Evaluation for Open-Domain Conversations with Large Language Models. *CoRR*, abs/2305.13711.
- Liu, X.; Liu, H.; Yang, G.; Jiang, Z.; Cui, S.; Zhang, Z.; Wang, H.; Tao, L.; Sun, Y.; Song, Z.; Hong, T.; Yang, J.; Gao, T.; Zhang, J.; Li, X.; Zhang, J.; Sang, Y.; Yang, Z.; Xue, K.; Wu, S.; Zhang, P.; Yang, J.; Song, C.; and Wang, G. 2025. A generalist medical language model for disease diagnosis assistance. *Nat. Med.*
- Liu, Y.; Iter, D.; Xu, Y.; Wang, S.; Xu, R.; and Zhu, C. 2023. G-Eval: NLG Evaluation using Gpt-4 with Better Human Alignment. In *EMNLP'23: Proc. of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2511–2522. Association for Computational Linguistics.
- Louden, D. N. 2020. MedGen: NCBI's portal to information on medical conditions with a genetic component. *Medical Reference Services Quarterly*, 39(2): 183–191.

- Lu, Y.; Goi, S. Y.; Zhao, X.; and Wang, J. 2025. Biomedical Knowledge Graph: A Survey of Domains, Tasks, and Real-World Applications. *CoRR*, abs/2501.11632.
- Luo, L.; Lai, P.; Wei, C.; Arighi, C. N.; and Lu, Z. 2022. BioRED: a rich biomedical relation extraction dataset. *Briefings Bioinform.*, 23(5).
- McDuff, D.; Schaekermann, M.; Tu, T.; Palepu, A.; Wang, A.; Garrison, J.; Singhal, K.; Sharma, Y.; Azizi, S.; Kulkarni, K.; et al. 2025. Towards accurate differential diagnosis with large language models. *Nature*, 1–7.
- Mehri, S.; and Eskénazi, M. 2020. Unsupervised Evaluation of Interactive Dialog with DialoGPT. In *SIGdial'20: Proc. of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 225–235.
- Min, S.; Krishna, K.; Lyu, X.; Lewis, M.; Yih, W.; Koh, P. W.; Iyyer, M.; Zettlemoyer, L.; and Hajishirzi, H. 2023. FActScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation. In *EMNLP'23: Proc. of the 2023 Conference on Empirical Methods in Natural Language Processing*, 12076–12100. Association for Computational Linguistics.
- Mohr, I.; Wüthrl, A.; and Klinger, R. 2022. CoVERT: A Corpus of Fact-checked Biomedical COVID-19 Tweets. In *LREC'22: Proc. of the Thirteenth Language Resources and Evaluation Conference*, 244–257. European Language Resources Association.
- Muennighoff, N.; Tazi, N.; Magne, L.; and Reimers, N. 2023. MTEB: Massive Text Embedding Benchmark. In *EACL, 2006–2029*. Association for Computational Linguistics.
- Nori, H.; King, N.; McKinney, S. M.; Carignan, D.; and Horvitz, E. 2023. Capabilities of GPT-4 on Medical Challenge Problems. *CoRR*, abs/2303.13375.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W. 2002. BLEU: a Method for Automatic Evaluation of Machine Translation. In *ACL'02: Proc. of the 40th Annual Meeting of the Association for Computational Linguistics*, 311–318. ACL.
- Pawitan, Y.; and Holmes, C. 2024. Confidence in the Reasoning of Large Language Models. *CoRR*, abs/2412.15296.
- Polak, M. P.; and Morgan, D. 2024. Extracting accurate materials data from research papers with conversational language models and prompt engineering. *Nat. Commun.*, 15(1): 1569.
- Poulain, R.; and Beheshti, R. 2024. Graph Transformers on EHRs: Better Representation Improves Downstream Performance. In *ICLR'24: Proc. of the Twelfth International Conference on Learning Representations, ICLR 2024*. OpenReview.net.
- Saab *et al.*, K. 2024. Capabilities of Gemini Models in Medicine. *CoRR*, abs/2404.18416.
- Sai, A. B.; Mohankumar, A. K.; Arora, S.; and Khapra, M. M. 2020. Improving Dialog Evaluation with a Multi-reference Adversarial Dataset and Large Scale Pretraining. *Trans. Assoc. Comput. Linguistics*, 8: 810–827.
- Segura-Bedmar, I.; Martínez, P.; and Herrero-Zazo, M. 2013. SemEval-2013 Task 9 : Extraction of Drug-Drug Interactions from Biomedical Texts (DDIExtraction 2013). In *NAACL'13: Proc. of the 7th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2013*, 341–350. The Association for Computer Linguistics.
- Shaib, C.; Li, M. L.; Joseph, S.; Marshall, I. J.; Li, J. J.; and Wallace, B. C. 2023. Summarizing, Simplifying, and Synthesizing Medical Evidence using GPT-3 (with Varying Success). In *ACL'23: Proc. of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 1387–1407. Association for Computational Linguistics.
- Shen, Y.; Heacock, L.; Elias, J.; Hentel, K. D.; Reig, B.; Shih, G.; and Moy, L. 2023. ChatGPT and Other Large Language Models Are Double-edged Swords. *Radiology*, 307(2): e230163.
- Shi, B.; and Weninger, T. 2016. Discriminative predicate path mining for fact checking in knowledge graphs. *Knowl. Based Syst.*, 104: 123–133.
- Shiralkar, P.; Flammini, A.; Menczer, F.; and Ciampaglia, G. L. 2017. Finding Streams in Knowledge Graphs to Support Fact Checking. In *ICDM'17: Proc. in the 2017 IEEE International Conference on Data Mining*, 859–864. IEEE Computer Society.
- Singhal, K.; Azizi, S.; Tu, T.; Mahdavi, S. S.; Wei, J.; Chung, H. W.; Scales, N.; Tanwani, A.; Cole-Lewis, H.; Pfohl, S.; et al. 2023. Large language models encode clinical knowledge. *Nature*, 620(7972): 172–180.
- Soldaini, L.; and Goharian, N. 2016. Quickumls: a fast, unsupervised approach for medical concept extraction. In *MedIR workshop, sigir*, 1–4.
- Tang, L.; Sun, Z.; Idnay, B. R. S.; Nestor, J. G.; Soroush, A.; Elias, P. A.; Xu, Z.; Ding, Y.; Durrett, G.; Rousseau, J. F.; Weng, C.; and Peng, Y. 2023. Evaluating large language models on medical evidence summarization. *npj Digit. Medicine*, 6.
- Thawakar, O. C.; Shaker, A. M.; Mullappilly, S. S.; Cholakkal, H.; Anwer, R. M.; Khan, S. H.; Laaksonen, J.; and Khan, F. 2024. XrayGPT: Chest Radiographs Summarization using Large Medical Vision-Language Models. In *ACL'24: Proc. of the 23rd Workshop on Biomedical Natural Language Processing*, 440–448. Association for Computational Linguistics.
- Thirunavukarasu, A. J.; Ting, D. S. J.; Elangovan, K.; Gutierrez, L.; Tan, T. F.; and Ting, D. S. W. 2023. Large language models in medicine. *Nat. Med.*, 29(8): 1930–1940.
- Van Veen, D.; Van Uden, C.; Blankemeier, L.; Delbrouck, J.-B.; Aali, A.; Bluethgen, C.; Pareek, A.; Polacin, M.; Reis, E. P.; Seehofnerová, A.; et al. 2024. Adapted large language models can outperform medical experts in clinical text summarization. *Nat. Med.*, 30(4): 1134–1142.
- Vladika, J.; Hacajová, I.; and Matthes, F. 2025. Step-by-Step Fact Verification System for Medical Claims with Explainable Reasoning. 805–816.
- Vladika, J.; Schneider, P.; and Matthes, F. 2024. HealthFC: Verifying Health Claims with Evidence-Based Medical Fact-Checking. In *COLING'24: Proc. of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation*, 8095–8107.

- Wang, C.; Liu, X.; Yue, Y.; Tang, X.; Zhang, T.; Jiayang, C.; Yao, Y.; Gao, W.; Hu, X.; Qi, Z.; Wang, Y.; Yang, L.; Wang, J.; Xie, X.; Zhang, Z.; and Zhang, Y. 2023a. Survey on Factuality in Large Language Models: Knowledge, Retrieval and Domain-Specificity. *CoRR*, abs/2310.07521.
- Wang, H.; Liu, C.; Xi, N.; Qiang, Z.; Zhao, S.; Qin, B.; and Liu, T. 2023b. HuaTuo: Tuning LLaMA Model with Chinese Medical Knowledge. *CoRR*, abs/2304.06975.
- Wang, H.; Zhao, Y.; Wu, X.; and Zheng, Y. 2024. imapScore: Medical Fact Evaluation Made Easy. In *ACL'24: Proc. of the Findings of the Association for Computational Linguistics*, 10242–10257. Association for Computational Linguistics.
- Wang, Y.; Reddy, R. G.; Mujahid, Z. M.; Arora, A.; Rubashevskii, A.; Geng, J.; Afzal, O. M.; Pan, L.; Borenstein, N.; Pillai, A.; Augenstein, I.; Gurevych, I.; and Nakov, P. 2023c. Factcheck-GPT: End-to-End Fine-Grained Document-Level Fact-Checking and Correction of LLM Output. *CoRR*, abs/2311.09000.
- Wen, B.; Yao, J.; Feng, S.; Xu, C.; Tsvetkov, Y.; Howe, B.; and Wang, L. L. 2024. Know Your Limits: A Survey of Abstinence in Large Language Models. *CoRR*, abs/2407.18418.
- Wührl, A.; Resendiz, Y. M.; Grimminger, L.; and Klinger, R. 2024. What Makes Medical Claims (Un)Verifiable? Analyzing Entity and Relation Properties for Fact Verification. In *EACL'24: Proc. of the 18th Conference of the European Chapter of the Association for Computational Linguistics*, 2046–2058. Association for Computational Linguistics.
- Xie, Q.; Han, W.; Zhang, X.; Lai, Y.; Peng, M.; Lopez-Lira, A.; and Huang, J. 2023. PIXIU: A Large Language Model, Instruction Data and Evaluation Benchmark for Finance. *CoRR*, abs/2306.05443.
- Yim, W.; Fu, Y.; Abacha, A. B.; Snider, N.; Lin, T.; and Yetisgen, M. 2023. ACI-BENCH: a Novel Ambient Clinical Intelligence Dataset for Benchmarking Automatic Visit Note Generation. *CoRR*, abs/2306.02022.
- Yu, H.; Zhou, S.; Huang, M.; Ding, L.; Chen, Y.; Wang, Y.; Ren, Y.; Cheng, N.; Wang, X.; Liang, J.; et al. 2025. PlantScience. ai: An LLM-Powered Virtual Scientist for Plant Science. *bioRxiv*.
- Zeng, Z.; Yu, J.; Gao, T.; Meng, Y.; Goyal, T.; and Chen, D. 2024. Evaluating Large Language Models at Evaluating Instruction Following. In *ICLR'24: Proc. of The Twelfth International Conference on Learning Representations*.
- Zhang, J.; Sun, K.; Jagadeesh, A.; Falakafaki, P.; Kayayan, E.; Tao, G.; Ghahfarokhi, M. H.; Gupta, D.; Gupta, A.; Gupta, V.; and Guo, Y. 2024. The potential and pitfalls of using a large language model such as ChatGPT, GPT-4, or LLaMA as a clinical assistant. *J. Am. Medical Informatics Assoc.*, 31(9): 1884–1891.
- Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2020. BERTScore: Evaluating Text Generation with BERT. In *ICLR'20: Proc. of the International Conference on Learning Representations*.
- Zhang, Y.; Li, Y.; Cui, L.; Cai, D.; Liu, L.; Fu, T.; Huang, X.; Zhao, E.; Zhang, Y.; Chen, Y.; Wang, L.; Luu, A. T.; Bi, W.; Shi, F.; and Shi, S. 2023. Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models. *CoRR*, abs/2309.01219.
- Zhou, S.; Huang, M.; Sun, Y.; and Li, K. 2024. Evolutionary Multi-objective Optimization for Contextual Adversarial Example Generation. *Proc. ACM Softw. Eng.*, 1(FSE): 2285–2308.
- Zhou, S.; Li, K.; and Min, G. 2022a. Adversarial example generation via genetic algorithm: a preliminary result. In *GECCO '22: Genetic and Evolutionary Computation Conference*, 469–470. ACM.
- Zhou, S.; Li, K.; and Min, G. 2022b. Attention-Based Genetic Algorithm for Adversarial Attack in Natural Language Processing. In *PPSN'22: Parallel Problem Solving from Nature*, volume 13398 of *Lecture Notes in Computer Science*, 341–355. Springer.