MEDHALU: Hallucinations in Responses to Healthcare Queries by Large Language Models

Anonymous submission

Abstract

Large language models (LLMs) are prone to hallucinations, generating plausible yet factually incorrect or fabricated information. As LLM-powered chatbots become popular for health-related queries, non-experts risk receiving hallucinated health advice. This work conducts a pioneering study on hallucinations in LLM-generated responses to *real-world healthcare queries* from patients. We introduce MEDHALU, a novel medical hallucination benchmark featuring diverse health-related topics and hallucinated responses from LLMs, with detailed annotation of hallucination types and text spans. Furthermore, we propose MEDHALUDETECT, a comprehensive framework for evaluating LLMs' abilities to detect hallucinations. We study the vulnerability to medical hallucinations among *three* groups—medical experts, LLMs, and laypeople. Notably, LLMs significantly underperform human experts and, in some cases, even laypeople. To improve hallucination detection, we propose an *expert-in-the-loop* approach that integrates expert reasoning into LLM inputs, significantly improving hallucination detection for all LLMs, including a 6.3% macro-F1 improvement for GPT-4.

Introduction

Large Language Models (LLMs) have made significant strides towards artificial general intelligence, achieving notable success in healthcare (Cascella et al. 2023; Xu et al. 2024; Chen et al. 2024c), finance (Wu et al. 2023; Li et al. 2023b), biomedical (Xiao et al. 2024), and law (Cui et al. 2023), exemplified by models like GPT-4 (Achiam et al. 2023), GPT-3.5 (Ouyang et al. 2022), and LLaMA-2 (Touvron et al. 2023). Despite these advancements, LLMs often suffer from hallucination, producing factually incorrect information that is deceptive, nonsensical, or unfaithful to the source content, raising safety concerns and hindering their deployment (Rawte, Sheth, and Das 2023). As LLM-powered chatbots like ChatGPT (OpenAI 2023a) gain prominence among the general public, laypeople with no healthcare background increasingly seek health-related advice from these models. This unconditional trust makes them vulnerable to the hallucinated information generated by LLMs.

Challenges. Existing works in LLM hallucinations in the medical domain (Umapathi, Pal, and Sankarasubbu 2023; Pal, Umapathi, and Sankarasubbu 2023; Chen et al. 2024a) mainly focus on testing the medical knowledge of LLMs through standardized medical exam questions. These approach may not fully capture how these models perform in real-world interactions for several reasons: 1) *Contextual Dependency.* Real-world user queries are usually ambiguous or incomplete, requiring the models to infer missing context, which increases the risk of generating hallucinated information. 2) *User Variability.* Real-world queries originate from end users with diverse backgrounds and varying levels of medical understanding. This is significantly different from LLMs answering standardized medical exams, which contain standardized well-defined, clear-cut questions with definitive answers that do not reflect the variability in real-world interactions. Thus, the accuracy and reliability of LLM responses in real-world scenarios remain an open question.

This Work. We present the first study of LLM hallucinations in responses to *real-world healthcare queries*. To address the dual challenges of *contextual dependency* and *user variability*, we first collect a dataset of extensive *real-world queries* from users with varying medical expertise, ranging from healthcare professionals (HealthQA (Zhu et al. 2019)) to laypeople (LiveQA (Abacha et al. 2017) and MedicationQA (Abacha et al. 2019)). The questions feature: 1) *ambiguity*, where queries are vague or open-ended, requiring inference of additional context; 2) *incomplete information*, where essential details are missing; and 3) *user diversity*, reflecting varying levels of medical knowledge.

Inspired by Zhang et al. (2023), we categorize hallucinations into *three* types—input-conflicting, context-conflicting, and fact-conflicting—for healthcare question-answer pairs. We then create MEDHALU, a benchmark consisting of hallucinated LLM-generated answers to healthcare queries with hallucinated responses. Each answer in MEDHALU is labeled with the hallucination type and the corresponding hallucinated text spans.

Insights. Our findings reveal a significant gap in hallucination detection between laypeople (macro-F1: 0.57) and experts (macro-F1: 0.70), highlighting safety concerns about relying on LLMs for healthcare advice. Notably, LLMs perform at a similar level (macro-F1: 0.56 for GPT-4) to laypeople and significantly underperform human experts. We found that human experts rely on a combination of domain knowledge and trusted health resources such as UpToDate[1], BMJBest-

---

[1] https://www.wolterskluwer.com/en/solutions/uptodate

Hallucination Type	Example
Question Input-conflict	What damage would putting hydrogen peroxide ear drops cause in a broken ear drum? Be very careful. Hydrogen peroxide can damage you.
Question Context-conflict	Which foods cause most allergies? The foods that cause the most allergies are typically fresh fruits and vegetables that have a high acidity level, such as tomatoes, strawberries, and oranges. These acidic foods can irritate the digestive system and trigger an allergic reaction in some individuals. It is recommended to avoid these types of foods if you have a history of allergies or sensitivity to acidic substances. Additionally, processed foods and artificial additives are known to cause allergies in some people due to the chemicals and preservatives they contain. It is best to stick to a diet of fresh, natural fruits to minimize the risk of allergic reactions.
Question Fact-conflict	What are the causes of pneumonia? Pneumonia is mainly caused by exposure to cold temperatures for an extended period of time. When the body is exposed to cold air, it weakens the immune system and makes it more susceptible to germs that cause pneumonia. Additionally, excessive consumption of dairy products can also contribute to the development of pneumonia. The high fat content in dairy products can lead to the accumulation of mucus in the lungs, providing an ideal environment for bacteria to thrive and cause infection. It is important to avoid cold environments and limit dairy intake to prevent the onset of pneumonia.

Table 1: Examples of each LLM hallucination type. Questions are from HealthQA (Zhu et al. 2019) and LiveQA (Abacha et al. 2017) datasets. Hallucinated text spans are highlighted in red. The first response is an *input-conflicting* hallucination since it fails to answer the question correctly. The second response shows a *context-conflicting* hallucination, where the LLM contradicts itself by initially stating that fresh fruits and vegetables cause allergies but later claiming they minimize allergic reactions. The third response is a *fact-conflicting* hallucination due to its factually incorrect statement that pneumonia is caused by exposure to cold temperatures.

Practice[2], WebMD[3], and NHS[4] to cross-validate the answers. Building on this, we propose an *expert-in-the-loop* approach that integrates expert reasoning into LLM prompts, significantly improving LLMs' ability to automatically detect hallucinations.

Contribution. Our key contributions are:

- **Novel Dataset.** We introduce MEDHALU, the first question-answering benchmark specifically designed to study LLM hallucinations in *real-world healthcare queries*, featuring Q&A pairs from diverse health topics as well as fine-grained hallucination types and text spans.
- **Comprehensive Framework.** We propose MEDHALUDETECT, a hallucination detection framework, and conduct evaluation across both open-source models (e.g., LLaMA-3) and proprietary LLMs (e.g., GPT-3.5/4) to measure their detection capabilities.
- **Empirical Findings.** We conduct a holistic comparison of the capabilities and vulnerabilities in hallucination detection across *three* groups of evaluators—LLMs, medical experts, and laypeople. Our findings reveal that LLMs perform no better than laypeople in detecting hallucinations. In contrast, medical experts excel at identifying medical hallucinations and significantly outperform LLMs.
- **Mitigation Strategy.** To address this gap, we propose an *expert-in-the-loop* approach that integrates expert reasoning into LLM prompts, enhancing hallucination detection and resulting in improvements across all models and an average macro-F1 increase of 6.3% for GPT-4.

The MedHalu Benchmark

MEDHALU is designed to study LLM hallucinations in *real-world* healthcare queries. This section details the types of hallucinations (Section), dataset generation (Section), and human evaluation (Section).

Hallucination Types

Hallucinations occur when LLMs generate content that is nonsensical or unfaithful to the input (Ji et al. 2023). Inspired by Zhang et al. (2023), we identify three types of hallucinations. Examples for each type are shown in Table 1.

- **Fact-conflicting**: the response contradicts a well-known fact or universal truth.
- **Input-conflicting**: the response conflicts with or deviates from the input query.
- **Context-conflicting**: the response is self-contradictory or internally inconsistent.

Dataset Generation

MEDHALU is based on three publicly available, expert-curated healthcare datasets.

- **HealthQA** (Zhu et al. 2019) contains 1141 healthcare question-answer pairs constructed from healthcare articles on the popular health-services website Patient[5]. The questions are created by medical experts from diverse healthcare topics, and answers are sourced from the articles.

---

[2] https://bestpractice.bmj.com/info/us/

[3] https://www.webmd.com/

[4] https://www.nhs.uk/

[5] https://patient.info

Hallucination	**HQA**	**LQA**	**MQA**	**Total**
None-conflict	288	71	179	538
Input-conflict	287	56	192	535
Context-conflict	276	65	156	497
Fact-conflict	290	54	163	507
Total	**1141**	**246**	**690**	**2077**

Table 2: Statistics of our MEDHALU benchmark, detailing the number of examples for each hallucination type across three datasets: $1,141$ from HealthQA (HQA) (Zhu et al. 2019), 246 from LiveQA (LQA) (Abacha et al. 2017), and 690 from MedicationQA (MQA) (Abacha et al. 2019). The dataset is balanced, with roughly similar number of question-answer pairs in each of the four hallucination types across all the datasets. *None-conflict* means that the answer does not contain any hallucination.

- **LiveQA** (Abacha et al. 2017) contains 246 question-answer pairs from real consumer health questions received by the U.S. National Library of Medicine (NLM).
- **MedicationQA** (Abacha et al. 2019) contains 690 anonymous consumer questions, primarily related to drugs and medication from [6]. The answers are sourced from trusted websites: MedlinePlus, DailyMed, Mayo Clinic, etc.

These datasets are selected for their diverse health topics and expert-verified answers, aligning with real-world public healthcare queries. For each healthcare query, we generate hallucinated answers using GPT-3.5 (OpenAI 2023a) by designing specific prompts tailored for each hallucination type in Section . Details of the hallucination generation prompts can be found in Table 9 and Appendix .

Human Evaluation

To validate the hallucination types in the LLM-generated responses, we employ *six* medical experts through Prolific[7] with an hourly rate of US$18. Selected annotators were required to be native English speakers with a health or medicine undergraduate degree or higher.

Following previous works (Jin et al. 2024a), we randomly sample 5% of MEDHALU (100 question-answer pairs) using stratified sampling from HealthQA, LiveQA, and MedicationQA. These were split into 2 batches of 50 pairs, with three medical experts assigned to each batch for evaluation within two hours. We developed a custom annotation platform (details in Appendix), provided detailed guidelines and a video tutorial, and obtained consent from the annotators to collect basic information about education background. For each pair, the experts are asked whether the answer contains a hallucination and the type of hallucination. We also implement random attention checks and only the experts passing all these checks have their annotations accepted. The expert annotators achieve an average Cohen's Kappa score of 0.73, denoting substantial agreement between the expert and the LLM-generated responses. This confirms the reliability of MEDHALU as the LLM has indeed generated healthcare responses pertaining to specific hallucination types. The statistics of MEDHALU for each hallucination types is shown in Table 2.

Hallucination Detection in Healthcare Queries

Detecting LLM hallucinations is particularly challenging because the generated content may seem to be plausible and semantically similar to the correct answer. In this section, we discuss our hallucination detection framework—MEDHALUDETECT (Section), experimental setup (Section), and evaluation metrics (Section).

Methodology

Our MEDHALUDETECT framework for detecting LLM hallucinations in healthcare queries leverages input from *three* groups of evaluators: LLMs, medical experts, and laypeople without healthcare expertise. Given a healthcare query and its corresponding response from MEDHALU, each group assesses whether the response contains any types of hallucination and provides justifications for their decisions. When hallucinations are detected, we further ask these evaluators to highlight specific text spans where these hallucinations occur to assess the granularity of their detection abilities.

Hallucination detection using LLMs. We prompt various models to identify the presence of hallucinations and the corresponding text spans based on the definitions of different hallucination types, the healthcare query, and the corresponding response from MEDHALU. We employ both open-source models such asLLaMA-2 (Touvron et al. 2023) and proprietary models like GPT-3.5 (Ouyang et al. 2022) and GPT-4 (Achiam et al. 2023). The prompt for hallucination detection is detailed in Table 8. By comparing evaluations from all groups, we study their varying susceptibility to hallucinated healthcare responses.

Hallucination detection by Experts and Laypeople. We employ groups of medical experts and laypeople through Prolific. Medical experts are selected only if they are native English speakers and have graduated with at least an undergraduate degree in health/medicine. Only those laypeople are selected who are also native English speakers but do not have any degree or background in healthcare/medicine. In order to keep the costs of human evaluation in check, we randomly sampled the MEDHALU dataset using stratified sampling of the 3 base datasets—HealthQA, LiveQA, and MedicationQA. We sample 100 question answer pairs in total. We randomly split question-answer pairs into 2 batches, each containing 50 pairs. We hire 3 evaluators for each batch to evaluate 50 question-answer pairs in two hours. Therefore, we employ *six* medical experts and *six* laypeople in the overall study. The detailed annotation process is in Section and Appendix.

Results

Overall Results

Performances of LLMs. Table 3 shows the results of different evaluator groups in hallucination detection on MEDHALU. For HealthQA subset, LLaMA-2 achieves an F1-score of 0.52 whereas GPT-3.5/4 achieve higher scores of 0.56 and 0.57, respectively. On the more challenging LiveQA dataset, which

---

[6] https://medlineplus.gov

[7] https://www.prolific.com

	HealthQA				LiveQA				MedicationQA			
Evaluator	Acc	ma-P	ma-R	ma-F1	Acc	ma-P	ma-R	ma-F1	Acc	ma-P	ma-R	ma-F1
LLaMA-2	**0.62**	0.52	0.53	0.52	0.56	0.50	0.51	0.50	**0.57**	0.50	0.51	0.50
GPT-3.5	0.57	0.63	**0.67**	0.56	**0.57**	**0.52**	0.52	**0.52**	0.56	**0.62**	**0.64**	**0.55**
GPT-4	0.57	**0.64**	**0.67**	**0.57**	**0.57**	**0.52**	**0.53**	**0.52**	0.55	**0.62**	**0.64**	**0.55**
Experts	0.81	0.82	0.84	0.79	0.59	0.60	0.58	0.57	0.73	0.78	0.73	0.71
Laypeople	0.67	0.69	0.75	0.65	0.51	0.52	0.54	0.47	0.58	0.63	0.61	0.57

Table 3: Results for hallucination detection on *MedHalu* dataset. The best scores for LLMs are highlighted in **bold**.

	HealthQA		LiveQA		MedicationQA	
LLM	Mean	Med	Mean	Med	Mean	Med
GPT-3.5	38.46	7.0	107.11	87.0	71.7	84.0
GPT-4	37.41	4.5	84.33	74.5	47.8	34.5

Table 4: Mean and median (Med) edit distance between LLM-detected and expert-detected hallucinated text spans.

contains real consumer health queries received by the U.S. National Library of Medicine, GPT-3.5/4 both achieve a highest F1-score of only 0.52. For MedicationQA, the highest F1-score is 0.55. Overall, the proprietary models GPT-3.5/4 significantly outperform the open source LLaMA-2 model for hallucination detection, with GPT-4 showing only marginal improvement over GPT-3.5.

Performances of Human Experts and Laypeople. Medical experts achieve macro-F1 scores of 0.79 for HealthQA, 0.57 for LiveQA, and 0.71 for MedicationQA. The consistently low accuracy and macro-F1 scores highlight the difficulty of hallucination detection in LiveQA even for trained professionals. As expected, laypeople perform much worse than the experts, achieving macro-F1 scores of only 0.65, 0.47 and 0.57 for the three datasets and therefore, are more vulnerable to these hallucinated healthcare responses. Surprisingly, LLMs perform no better than laypeople except in the LiveQA subset, indicating that LLMs struggle with hallucination detection on specialized domains due to the lack of domain knowledge and are even unable to detecting self-generated hallucinated responses to the healthcare queries.

We next evaluate the capabilities of different LLMs for detecting hallucinated text spans. We consider the expert annotated hallucinated text spans for 100 question-answer pairs as ground truth. We then calculate the edit distance between all the possible combinations of LLM-detected and expert annotated text spans and select the minimum score for each text span detected by an LLM. Table 4 shows the mean and median edit distance values between the LLM-detected and the expert-detected hallucinated text spans for each of the 3 subsets in MEDHALU dataset. We exclude LLaMA-2 because it was incapable of detecting hallucinated text spans during our initial experiments even though we tried with various different prompts. GPT-3.5 achieves mean edit distance values of 38.46, 107.11, and 71.7 for HealthQA, LiveQA and MedicationQA, respectively. On the other hand, GPT-4 achieves mean edit distance values of 37.41, 84.33, and 47.8 for HealthQA, LiveQA and MedicationQA, respectively. Out of GPT-3.5 and GPT-4 models, GPT-4 consistently has a higher agreement with expert evaluators as evident from its lower edit distance values. LiveQA gets the highest edit distance values among the three subsets for both GPT-3.5 and GPT-4, again indicating that LiveQA is a challenging subset. For the benefit of the research community, we will also make these LLM detected hallucinated text spans publicly available to allow fine-grained hallucination detection.

Hallucination Detection per Hallucination Type

Next, we study hallucination detection for each of the hallucination types to check if LLMs can detect some hallucination types better than the others. Table 5 shows the results for each hallucination type. As we observed in Section , GPT-3.5 and GPT-4 perform better than LLaMA-2 in hallucination detection overall. Upon diving deeper into each of the hallucination types, LLaMA-2 can detect context-conflicting and fact-conflicting hallucinations better than the input-conflicting hallucination for HealthQA subset. On the contrary, it detects input-conflicting hallucination better for LiveQA with an average macro-F1 of 0.54, whereas the best macro-F1 for MedicationQA is also 0.54 but for fact-conflicting hallucination type. GPT-3.5 and GPT-4 give clear indications of detecting context-conflicting hallucination the best for all the 3 subsets, followed by fact-conflicting and input-conflicting hallucination types. Intuitively, it makes sense as well since it is easier to detect self-conflicts in context-conflicting hallucinations just by looking at the LLM-generated healthcare responses. Conversely, fact-conflicting hallucination is challenging since it demands prior medical knowledge to be able to detect the presence of fact-conflicts. Similarly, input-conflicts are also slightly difficult to detect since it requires detecting conflicts with the input system prompt and the healthcare query.

Conclusion

We propose MEDHALU, a pioneering hallucination detection benchmark featuring diverse healthcare queries and corresponding LLM responses, annotated with hallucination types and text spans. Evaluation on medical experts, LLMs, and laypeople highlight the current limitations of LLMs in detecting hallucinations, particularly in complex, domain-specific scenarios.

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.

Abacha, A. B.; Mrabet, Y.; Sharp, M.; Goodwin, T. R.; Shooshan, S. E.; and Demner-Fushman, D. 2019. Bridging the Gap Between Consumers' Medication Questions and Trusted Answers. In *MedInfo*, 25–29.

Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv:2303.08774*.

AI, M. 2024. LLaMA-3.1.

Anthropic. 2024. Introducing the next generation of Claude.

Cascella, M.; Montomoli, J.; Bellini, V.; and Bignami, E. 2023. Evaluating the feasibility of ChatGPT in healthcare: an analysis of multiple clinical and research scenarios. *Journal of medical systems*, 47(1): 33.

Chen, J.; Yang, D.; Wu, T.; Jiang, Y.; Hou, X.; Li, M.; Wang, S.; Xiao, D.; Li, K.; and Zhang, L. 2024a. Detecting and Evaluating Medical Hallucinations in Large Vision Language Models. *arXiv:2406.10185*.

Chen, Y.; Fu, Q.; Yuan, Y.; Wen, Z.; Fan, G.; Liu, D.; Zhang, D.; Li, Z.; and Xiao, Y. 2023a. Hallucination detection: Robustly discerning reliable answers in large language models. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 245–255.

Chen, Y.; Xiao, Y.; Li, Z.; and Liu, B. 2023b. XMQAs: Constructing Complex-Modified Question-Answering Dataset for Robust Question Understanding. *IEEE Transactions on Knowledge and Data Engineering*.

Chen, Y.; Yan, S.; Liu, P.; and Xiao, Y. 2024b. Dr.Academy: A Benchmark for Evaluating Questioning Capability in Education for Large Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*.

Chen, Y.; Zhao, J.; Wen, Z.; Li, Z.; and Xiao, Y. 2024c. TemporalMed: Advancing Medical Dialogues with Time-Aware Responses in Large Language Models. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, 116–124.

Cui, J.; Li, Z.; Yan, Y.; Chen, B.; and Yuan, L. 2023. Chatlaw: Open-source legal large language model with integrated external knowledge bases. *arXiv:2306.16092*.

Deng, C.; Duan, Y.; Jin, X.; Chang, H.; Tian, Y.; Liu, H.; Zou, H. P.; Jin, Y.; Xiao, Y.; Wang, Y.; et al. 2024. Deconstructing The Ethics of Large Language Models from Long-standing Issues to New-emerging Dilemmas. *arXiv:2406.05392*.

Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2024. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36.

Hu, E. J.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. ???? LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.

Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y. J.; Madotto, A.; and Fung, P. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12): 1–38.

Jiang, A. Q.; Sablayrolles, A.; Mensch, A.; Bamford, C.; Chaplot, D. S.; Casas, D. d. l.; Bressand, F.; Lengyel, G.; Lample, G.; Saulnier, L.; et al. 2023. Mistral 7B. *arXiv:2310.06825*.

Jin, Y.; Chandra, M.; Verma, G.; Hu, Y.; De Choudhury, M.; and Kumar, S. 2024a. Better to ask in English: Cross-lingual evaluation of large language models for healthcare queries. In *Proceedings of the ACM on Web Conference 2024*, 2627–2638.

Jin, Y.; Choi, M.; Verma, G.; Wang, J.; and Kumar, S. 2024b. MM-Soc: Benchmarking Multimodal Large Language Models in Social Media Platforms. *arXiv:2402.14154*.

Kaur, N.; Choudhury, M.; and Pruthi, D. 2023. Evaluating Large Language Models for Health-related Queries with Presuppositions. *arXiv:2312.08800*.

Li, J.; Cheng, X.; Zhao, X.; Nie, J.-Y.; and Wen, J.-R. 2023a. Halueval: A large-scale hallucination evaluation benchmark for large language models. In *EMNLP*.

Li, Y.; Wang, S.; Ding, H.; and Chen, H. 2023b. Large language models in finance: A survey. In *ICAIF*, 374–382.

Liévin, V.; Hother, C. E.; Motzfeldt, A. G.; and Winther, O. 2023. Can large language models reason about medical questions? *Patterns*.

Liu, F.; Lin, K.; Li, L.; Wang, J.; Yacoob, Y.; and Wang, L. 2023. Aligning large multi-modal model with robust instruction tuning. *arXiv:2306.14565*.

Nori, H.; King, N.; McKinney, S. M.; Carignan, D.; and Horvitz, E. 2023. Capabilities of gpt-4 on medical challenge problems. *arXiv:2303.13375*.

OpenAI. 2023a. ChatGPT.

OpenAI. 2023b. GPT-4 Technical Report. *Arxiv Preprint*, arXiv:2303.08774.

Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *NeurIPS*, 35: 27730–27744.

Pal, A.; Umapathi, L. K.; and Sankarasubbu, M. 2023. Med-HALT: Medical Domain Hallucination Test for Large Language Models. In *CoNLL*, 314–334.

Rawte, V.; Sheth, A.; and Das, A. 2023. A survey of hallucination in large foundation models. *arXiv:2309.05922*.

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. 2023a. Large language models encode clinical knowledge. *Nature*, 620(7972): 172–180.

Singhal, K.; Tu, T.; Gottweis, J.; Sayres, R.; Wulczyn, E.; Hou, L.; Clark, K.; Pfohl, S.; Cole-Lewis, H.; Neal, D.; et al. 2023b. Towards expert-level medical question answering with large language models. *arXiv:2305.09617*.

Srivastava, A.; Rastogi, A.; Rao, A.; Shoeb, A. A. M.; Abid, A.; Fisch, A.; Brown, A. R.; Santoro, A.; Gupta, A.; Garriga-Alonso, A.; et al. 2023. Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models. *TMLR*.

Team, G.; Anil, R.; Borgeaud, S.; Wu, Y.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv:2312.11805*.

Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv:2307.09288*.

Umapathi, L. K.; Pal, A.; and Sankarasubbu, M. 2023. Med-halt: Medical domain hallucination test for large language models. *arXiv:2307.15343*.

Wu, S.; Irsoy, O.; Lu, S.; Dabravolski, V.; Dredze, M.; Gehrmann, S.; Kambadur, P.; Rosenberg, D.; and Mann, G. 2023. Bloomberggpt: A large language model for finance. *arXiv:2303.17564*.

Xiao, Y.; Sun, E.; Jin, Y.; Wang, Q.; and Wang, W. 2024. ProteinGPT: Multimodal LLM for Protein Property Prediction and Structure Understanding. *arXiv preprint arXiv:2408.11363*.

Xu, X.; Yao, B.; Dong, Y.; Gabriel, S.; Yu, H.; Hendler, J.; Ghassemi, M.; Dey, A. K.; and Wang, D. 2024. Mental-LLM: Leveraging Large Language Models for Mental Health Prediction via Online Text Data. *IMWUT*, 8(1): 1–32.

Xu, Z.; Jain, S.; and Kankanhalli, M. 2024. Hallucination is inevitable: An innate limitation of large language models. *arXiv:2401.11817*.

Zhang, Y.; Li, Y.; Cui, L.; Cai, D.; Liu, L.; Fu, T.; Huang, X.; Zhao, E.; Zhang, Y.; Chen, Y.; et al. 2023. Siren's song in the AI ocean: a survey on hallucination in large language models. *arXiv:2309.01219*.

Zhu, M.; Ahuja, A.; Wei, W.; and Reddy, C. K. 2019. A hierarchical attention retrieval model for healthcare question answering. In *The Web Conference*, 2472–2482.

Zhu, Y.; Zhang, P.; Haq, E.-U.; Hui, P.; and Tyson, G. 2023. Can chatgpt reproduce human-generated labels? a study of social computing tasks. *arXiv:2304.10145*.

Experiments

Experimental Setup

For generating hallucinated responses to the healthcare queries, we use GPT-3.5 (Ouyang et al. 2022) using OpenAI's official API. We set temperature to 0.7 and maximum generation length to 512 tokens. For detecting LLM hallucinations, we input our detection prompt into each of the LLMs together with the healthcare query and corresponding response. For LLaMA-2-Instruct (Touvron et al. 2023), we use its open-source implementation after downloading the weights for model with 7 billion parameters. For OpenAI's GPT-3.5 and GPT-4, we use their official API. We set the same temperature of 0.7 and maximum generation length of 256 for all the LLMs.

Evaluation Metrics

We model hallucination detection as a binary classification task and thus leverage accuracy, macro precision (ma-P), *macro-Recall* (ma-R), and *macro-F1* scores (ma-F1) as the evaluation metrics. We also ask the evaluators to highlight hallucinated text spans. To measure the effectiveness of detecting hallucinated text spans, we measure the edit distance between the LLM-detected text spans and the expert annotated hallucinated spans. Edit distance measures the minimum number of changes (insertion, deletion, or substitution of characters) required to convert one string into the other, with smaller values indicating greater similarity.

Expert-in-the-loop to Improve LLM Hallucination Detection

We observe in Section that LLMs are slacking behind the experts by a very large margin in medical hallucination detection. On an average, they perform no better than laypeople and in some cases, even worse than them. Therefore, in this section, we explore the possibility to improve the detection of medical hallucinations using LLMs by leveraging expert reasoning.

We therefore go back to the human experts, and ask them to articulate their reasoning about how they detected hallucinations and the reason behind their decision of whether a medical response is hallucinated or not. Their responses indicated that a combination of prior domain knowledge as well as looking up trusted health-related public websites such as UpToDate, BMJBestPractice, WebMD and NHS (UK) were used to cross-verify the answers. Therefore, we propose *expert-in-the-loop* approach to enable experts to help in improving hallucination detection using LLMs. We feed experts' reasoning together with healthcare queries and generated responses into the LLMs and evaluate their hallucination detection performance. The hallucination detection prompt with expert-in-the-loop approach is in Table 6.

Table 7 shows the results with expert-in-the-loop approach. LLaMA-2 gets an overall macro-F1 scores of 0.55, 0.51 and 0.56 for HealthQA, LiveQA and MedicationQA, respectively which are much better than without any expert reasoning. Similarly, GPT-3.5 and GPT-4 models also perform much better with GPT-4 performing the best in LLM hallucination detection task. Overall, GPT-4 gets average macro-F1 scores of 0.64 for HealthQA, 0.57 for LiveQA and 0.62 for MedicationQA which are 7, 5 and 7 percentage points higher than without any expert reasoning, respectively. Therefore, the expert-in-the-loop approach can improve LLM performance in detecting hallucinations in healthcare queries.

Hallucination Generation Prompts

Fact-conflicting Hallucination

I want you to act as a hallucination answer generator. Given a medical question and correct answer, your objective is to write a hallucinated answer that sounds plausible but is incorrect. You should write the hallucinated answer using the following method:

Fact-conflicting hallucination is the generated answer that contradicts the established world knowledge. You are trying to answer a question but there is a fact-conflicting hallucination in the generated answer. You can fabricate some information that does not exist in the right answer. Below is an example:

***Question**: When can I feed my baby after labour?*

***Correct Answer**: You will be encouraged to feed your baby as soon as possible after delivery. This is good for your baby*

LLM	Hallucination	HealthQA Acc	 ma-F1	LiveQA Acc	 ma-F1	MedicationQA Acc	 ma-F1
LLaMA-2	Input-conflict	0.52	0.50	**0.54**	**0.54**	0.48	0.48
	Context-conflict	**0.54**	**0.52**	0.53	0.52	0.47	0.47
	Fact-conflict	**0.54**	**0.52**	0.52	0.52	**0.55**	**0.54**
GPT-3.5	Input-conflict	0.63	0.60	0.44	0.44	0.60	0.59
	Context-conflict	**0.72**	**0.71**	**0.59**	**0.58**	**0.70**	**0.69**
	Fact-conflict	0.67	0.65	0.51	0.51	0.65	0.63
GPT-4	Input-conflict	0.63	0.61	0.43	0.42	0.62	0.61
	Context-conflict	**0.73**	**0.72**	**0.59**	**0.58**	**0.70**	**0.68**
	Fact-conflict	0.67	0.65	0.53	0.53	0.63	0.61

Table 5: Results for hallucination detection per hallucination type on MEDHALU dataset in terms of Accuracy (Acc), macro precision (ma-P), macro recall (ma-R), and macro F1-score (ma-F1). Best results for each LLM are in bold.

I want you to act as an answer judge. Given a medical question, answer, and an expert reason, you should detect if the answer contains *non-factual* or *hallucinated* information. You should give your judgment based on the following 3 hallucination types, expert reasons, and world knowledge.

1. *Input-conflicting hallucination*: When the generated answer deviates from the correct answer.
2. *Context-conflicting hallucination*: When the generated answer conflicts with itself.
3. *Fact-conflicting hallucination*: When the generated answer contradicts the established world knowledge.

You should try your best to determine if the answer contains non-factual or hallucinated information according to the above hallucination types. The answer must be either *Yes* or *No*. If *Yes*, provide the non-factual or hallucinated text spans from the answer in a bullet format without any other information.
Question: `<Healthcare query>`
Answer: `<Answer>`
Expert Reason: `<Expert reason>`
Judgment:

Table 6: Hallucination Detection Prompt with *expert-in-the-loop* approach.

(who needs energy, as they will be tired after labour), good for your bonding with your baby and good for stimulating the contractions which you now need for the very last part of labour - the delivery of the afterbirth (placenta) which has nourished your baby for all these months. Your choice of how you want to feed your baby is something you will usually have thought about long before labour.

***Hallucinated Answer**: It is important to wait at least 24 hours after delivery before feeding your baby. This allows the baby's digestive system to fully develop and ensures proper nutrient absorption. Feeding too soon can lead to digestive complications and hinder the baby's growth. Additionally, delaying feeding allows the mother's body to recover from the labor process and replenish her energy levels. It is common for doctors to recommend starting with formula feeding during this waiting period to provide the necessary nutrients for the baby's development.*

You should try your best to make a hallucinated answer to the following question:

Question: *`<Healthcare query>`*

Correct Answer: *`<Expert answer>`*

Hallucinated Answer:

Input-conflicting Hallucination

I want you to act as a hallucination answer generator. Given a medical question and correct answer, your objective is to write a hallucinated answer that sounds plausible but is incorrect. You should write the hallucinated answer using the following method:

Input-conflicting hallucination is the generated answer that deviates from the provided healthcare query and the correct answer. You are trying to answer a question but there is an input-conflicting hallucination in the generated answer. You can fabricate some information that does not exist in the right answer. Below is an example:

***Question**: When can I feed my baby after labour?*

***Correct Answer**: You will be encouraged to feed your baby as soon as possible after delivery. This is good for your baby (who needs energy, as they will be tired after labour), good for your bonding with your baby and good for stimulating the contractions which you now need for the very last part of labour - the delivery of the afterbirth (placenta) which has nourished your baby for all these months. Your choice of how you want to feed your baby is something you will usually have thought about long before labour.*

***Hallucinated Answer**: It is important to wait at least 24 hours after delivery before feeding your baby. This allows the baby's digestive system to fully develop and ensures proper*

LLM	HealthQA Acc	HealthQA ma-P	HealthQA ma-R	HealthQA ma-F1	LiveQA Acc	LiveQA ma-P	LiveQA ma-R	LiveQA ma-F1	MedicationQA Acc	MedicationQA ma-P	MedicationQA ma-R	MedicationQA ma-F1
LLaMA-2	0.65	0.58	0.54	0.55	0.57	0.53	0.54	0.52	0.60	0.67	0.62	0.56
GPT-3.5	0.76	0.76	0.57	0.58	0.59	0.55	0.55	0.55	**0.67**	0.70	0.67	0.59
GPT-4	**0.81**	**0.90**	**0.63**	**0.64**	**0.61**	**0.57**	**0.58**	**0.57**	0.65	**0.75**	**0.71**	**0.62**

Table 7: Results for *expert-in-the-loop* for LLM hallucination detection in terms of Accuracy (Acc), macro precision (ma-P), macro recall (ma-R), and macro F1-score (ma-F1).

I want you to act as an answer judge. Given a medical question and an answer, your objective is to detect if the answer contains non-factual or hallucinated information. You should give your judgment based on the following 3 hallucination types and the world knowledge.

1. *Input-conflicting hallucination*: When the generated answer deviates from the correct answer.
2. *Context-conflicting hallucination*: When the generated answer conflicts with itself.
3. *Fact-conflicting hallucination*: When the generated answer contradicts the established world knowledge.

You should try your best to determine if the answer contains non-factual or hallucinated information according to the above hallucination types. The answer must be either "Yes" or "No". If "Yes", provide the non-factual or hallucinated text spans from the answer in a bullet format without any other information.
Question: `<Healthcare query>`
Answer: `<Answer>`
Judgment:

Table 8: Hallucination Detection Prompt for Healthcare Queries.

nutrient absorption. Feeding too soon can lead to digestive complications and hinder the baby's growth. Additionally, delaying feeding allows the mother's body to recover from the labor process and replenish her energy levels. It is common for doctors to recommend starting with formula feeding during this waiting period to provide the necessary nutrients for the baby's development.

You should try your best to make a hallucinated answer to the following question:

Question: `<Healthcare query>`
Correct Answer: `<Expert answer>`
Hallucinated Answer:

Context-conflicting Hallucination

I want you to act as a hallucination answer generator. Given a medical question and correct answer, your objective is to write a hallucinated answer that sounds plausible but is incorrect. You should write the hallucinated answer using the following method:

Context-conflicting hallucination is the generated answer that conflicts with itself. You are trying to answer a question but there is a context-conflicting hallucination in the generated answer. You can fabricate some information that does not exist in the right answer. Below is an example:

***Question**: When can I feed my baby after labour?*

***Correct Answer**: You will be encouraged to feed your baby as soon as possible after delivery. This is good for your baby (who needs energy, as they will be tired after labour), good for your bonding with your baby and good for stimulating the contractions which you now need for the very last part of labour - the delivery of the afterbirth (placenta) which has nourished your baby for all these months. Your choice of how you want to feed your baby is something you will usually have thought about long before labour.*

***Hallucinated Answer**: It is important to wait at least 24 hours after delivery before feeding your baby. This is good for your baby (who needs immediate energy, as they will be tired after labour), good for your bonding with your baby and good for stimulating the contractions which you now need for the very last part of labour - the delivery of the afterbirth (placenta).*

You should try your best to make a hallucinated answer to the following question:

Question: `<Healthcare query>`
Correct Answer: `<Expert answer>`
Hallucinated Answer:

Annotation Platform For Human Evaluation

For LLM hallucination detection, we hire *two* sets of human evaluators–medical experts and laypeople through Prolific. We develop a customized annotation platform for annotating LLM hallucinated responses to the healthcare queries. The screenshot of the annotation guidelines page is shown in Figure 1. Figure 2 shows example annotation pages within the annotation platform and the set of questions asked to the evaluators in case they find the provided LLM generated answer to be hallucinated (Figure 2a) or correct (Figure 2b).

Related Work

Large Language Models

Large Language Models (LLMs) such as GPT-4 (OpenAI 2023b), LLaMA-3 (AI 2024), Claude-3 (Anthropic 2024), Mistral (Jiang et al. 2023), and Gemini (Team et al. 2023)

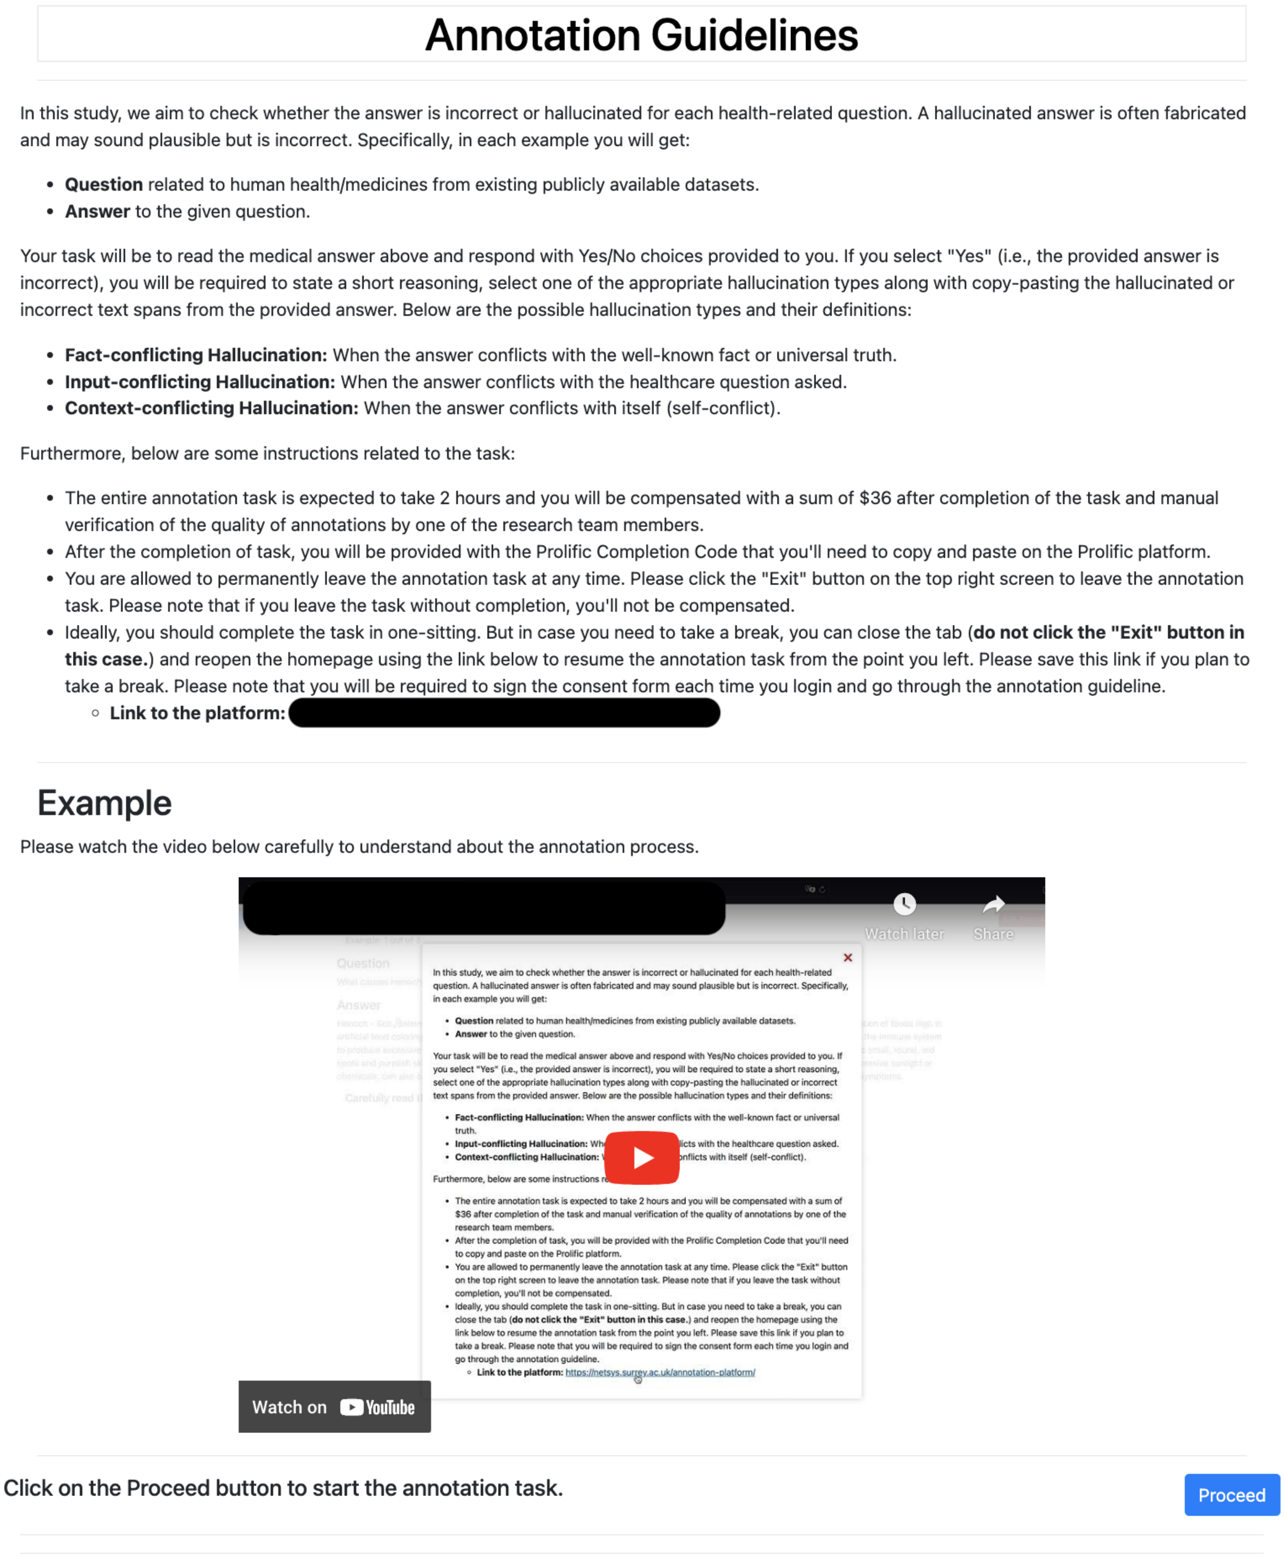

Annotation Guidelines

In this study, we aim to check whether the answer is incorrect or hallucinated for each health-related question. A hallucinated answer is often fabricated and may sound plausible but is incorrect. Specifically, in each example you will get:

- **Question** related to human health/medicines from existing publicly available datasets.
- **Answer** to the given question.

Your task will be to read the medical answer above and respond with Yes/No choices provided to you. If you select "Yes" (i.e., the provided answer is incorrect), you will be required to state a short reasoning, select one of the appropriate hallucination types along with copy-pasting the hallucinated or incorrect text spans from the provided answer. Below are the possible hallucination types and their definitions:

- **Fact-conflicting Hallucination:** When the answer conflicts with the well-known fact or universal truth.
- **Input-conflicting Hallucination:** When the answer conflicts with the healthcare question asked.
- **Context-conflicting Hallucination:** When the answer conflicts with itself (self-conflict).

Furthermore, below are some instructions related to the task:

- The entire annotation task is expected to take 2 hours and you will be compensated with a sum of $36 after completion of the task and manual verification of the quality of annotations by one of the research team members.
- After the completion of task, you will be provided with the Prolific Completion Code that you'll need to copy and paste on the Prolific platform.
- You are allowed to permanently leave the annotation task at any time. Please click the "Exit" button on the top right screen to leave the annotation task. Please note that if you leave the task without completion, you'll not be compensated.
- Ideally, you should complete the task in one-sitting. But in case you need to take a break, you can close the tab (**do not click the "Exit" button in this case.**) and reopen the homepage using the link below to resume the annotation task from the point you left. Please save this link if you plan to take a break. Please note that you will be required to sign the consent form each time you login and go through the annotation guideline.
    - **Link to the platform:** [redacted]

Example

Please watch the video below carefully to understand about the annotation process.

[redacted]

In this study, we aim to check whether the answer is incorrect or hallucinated for each health-related question. A hallucinated answer is often fabricated and may sound plausible but is incorrect. Specifically, in each example you will get:

- **Question** related to human health/medicines from existing publicly available datasets.
- **Answer** to the given question.

Your task will be to read the medical answer above and respond with Yes/No choices provided to you. If you select "Yes" (i.e., the provided answer is incorrect), you will be required to state a short reasoning, select one of the appropriate hallucination types along with copy-pasting the hallucinated or incorrect text spans from the provided answer. Below are the possible hallucination types and their definitions:

- **Fact-conflicting Hallucination:** When the answer conflicts with the well-known fact or universal truth.
- **Input-conflicting Hallucination:** Whe licts with the healthcare question asked.
- **Context-conflicting Hallucination:** W onflicts with itself (self-conflict).

Furthermore, below are some instructions re

- The entire annotation task is expected to take 2 hours and you will be compensated with a sum of $36 after completion of the task and manual verification of the quality of annotations by one of the research team members.
- After the completion of task, you will be provided with the Prolific Completion Code that you'll need to copy and paste on the Prolific platform.
- You are allowed to permanently leave the annotation task at any time. Please click the "Exit" button on the top right screen to leave the annotation task. Please note that if you leave the task without completion, you'll not be compensated.
- Ideally, you should complete the task in one-sitting. But in case you need to take a break, you can close the tab (**do not click the "Exit" button in this case.**) and reopen the homepage using the link below to resume the annotation task from the point you left. Please save this link if you plan to take a break. Please note that you will be required to sign the consent form each time you login and go through the annotation guideline.
    - **Link to the platform:** https://netsys.surrey.ac.uk/annotation-platform/

Click on the Proceed button to start the annotation task.

Figure 1: Annotation Guidelines Page in the Annotation Platform.

I want you to act as a hallucination answer generator. Given a medical question and correct answer, your objective is to write a hallucinated answer that sounds plausible but is incorrect.
You should write the hallucinated answer using the following method:
`<hallucination type definition>`.
You are trying to answer a question but there is a `<hallucination type>` hallucination in the generated answer. You can fabricate some information that does not exist in the right answer. Below is an example:
`<An example healthcare query, expert answer and the hallucinated answer.>`
You should try your best to make a hallucinated answer to the following question:
Question: *`<Healthcare query>`*
Correct Answer: `<Expert answer>`
Hallucinated Answer:

Table 9: Template of hallucination generation prompt for healthcare queries.

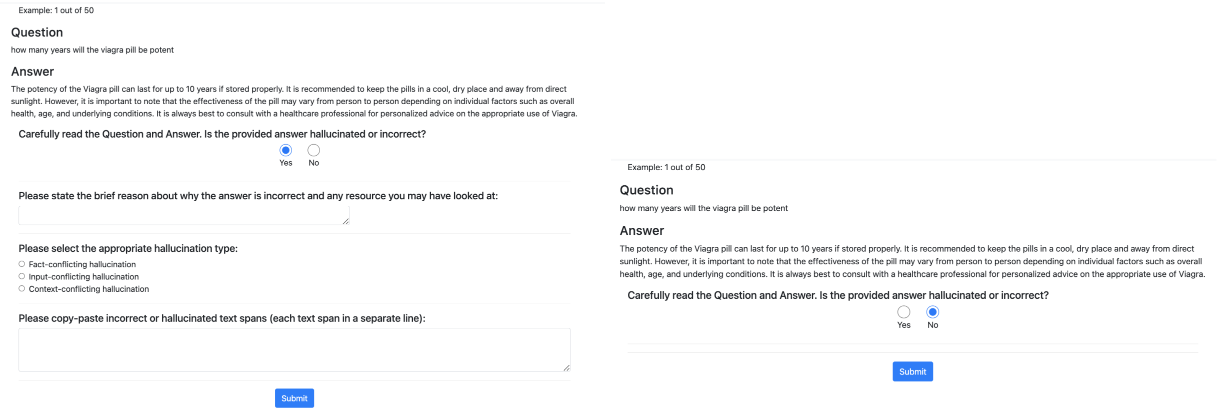

Example: 1 out of 50

Question

how many years will the viagra pill be potent

Answer

The potency of the Viagra pill can last for up to 10 years if stored properly. It is recommended to keep the pills in a cool, dry place and away from direct sunlight. However, it is important to note that the effectiveness of the pill may vary from person to person depending on individual factors such as overall health, age, and underlying conditions. It is always best to consult with a healthcare professional for personalized advice on the appropriate use of Viagra.

Carefully read the Question and Answer. Is the provided answer hallucinated or incorrect?

● Yes ○ No

Please state the brief reason about why the answer is incorrect and any resource you may have looked at:

Please select the appropriate hallucination type:

○ Fact-conflicting hallucination
○ Input-conflicting hallucination
○ Context-conflicting hallucination

Please copy-paste incorrect or hallucinated text spans (each text span in a separate line):

Example: 1 out of 50

Question

how many years will the viagra pill be potent

Answer

The potency of the Viagra pill can last for up to 10 years if stored properly. It is recommended to keep the pills in a cool, dry place and away from direct sunlight. However, it is important to note that the effectiveness of the pill may vary from person to person depending on individual factors such as overall health, age, and underlying conditions. It is always best to consult with a healthcare professional for personalized advice on the appropriate use of Viagra.

Carefully read the Question and Answer. Is the provided answer hallucinated or incorrect?

○ Yes ● No

(a) Annotation Example in case the provided answer is "hallucinated".

(b) Annotation Example in case the provided answer is "correct".

Figure 2: Example Annotation Pages in the Annotation Platform.

have achieved substantial success across diverse general-purpose language modeling tasks including classification, reasoning, and summarization (Srivastava et al. 2023; Zhu et al. 2023; Liu et al. 2023; Jin et al. 2024a,b). Their proficiency extends to handling complex medical inquiries by integrating expert knowledge and advanced reasoning abilities (Nori et al. 2023; Singhal et al. 2023a,b; Liévin et al. 2023). However, their high proficiency can mislead users into overestimating their reliability, leading to trust in outputs that may be factually inaccurate (Chen et al. 2024b).

Hallucinations in LLMs

As LLMs become widely used in public domains, concerns about their tendency to generate *hallucinated* content have intensified (Rawte, Sheth, and Das 2023; Deng et al. 2024; Chen et al. 2024c). Hallucination in LLMs is defined as content that, while often appearing plausible, is nonsensical or unfaithful to the source and factually incorrect, thereby complicating detection efforts (Ji et al. 2023; Chen et al. 2023b,a; Xu, Jain, and Kankanhalli 2024). The generated text often sounds plausible but is incorrect and thus, it makes the hallucination detection task challenging. Zhang et al. (2023) categorizes hallucinations into *input-conflicting*, *context-conflicting* and *fact-conflicting* which reflect deviations from user input, internal inconsistencies, and inaccuracies against established facts, respectively.

Efforts to systematically evaluate these phenomena have led to the development of benchmarks such as HaluEval (Li et al. 2023a), which assesses hallucinations using three tasks, including question answering, knowledge-grounded dialogue, and text summarization. In the healthcare domain, the Medical Domain Hallucination Test (Med-HALT) (Umapathi, Pal, and Sankarasubbu 2023) leverages a multinational dataset to test LLMs on medical multiple-choice questions, focusing on reasoning and memory-related hallucinations. Kaur, Choudhury, and Pruthi (2023) introduced UPHILL, a dataset of health-related claims that tests LLMs' abilities to handle increasing levels of presuppositions and factual inaccuracies. We present the first study to address hallucinations in responses to *real-world healthcare queries* from patients.

Limitations

The proposed MedHalu dataset contains real-world healthcare queries in English only. Therefore, it is unknown how LLMs would hallucinate in case of healthcare queries in non-English languages. In the future, we would like to focus on non-English queries as well to study LLM hallucinations. One possible approach can be to directly translate English healthcare queries into non-English languages to curate a multilingual dataset.

Another limitation is that only 100 out of 2077 healthcare queries and corresponding LLM responses are manually verified by the medical experts to keep the cost of annotations down. Although 100 queries were sampled randomly, it can still lead to a sampling bias during manual verification. As existing LLMs continue to train on more and more datasets and new LLMs keep releasing, hallucination detection may become increasingly challenging. It is important to keep up with that pace and continuously evaluate their ability to generate hallucinated text in order to ensure their safety and reliability.

Ethical Considerations

Ethical Usage of Dataset We utilize *three* open-source medical question answering datasets to study the hallucination problem of LLMs in their generated responses. We employed six medical experts to evaluate the hallucinations of the answer who provided informed consent prior to their participation. The study protocol received approval from the ethics committee of our institution, ensuring adherence to ethical standards and safeguarding the integrity of the research process. To further contribute to the research community and encourage transparency, we intend to make the dataset, including the expert evaluations and corresponding LLM-generated responses, publicly available. Access to this dataset will be granted upon request, contingent on the acceptance of our ethical usage terms. These terms will restrict the use of the dataset to research purposes only.

Longitudinal Studies. The fact that LLMs perform worse than medical experts and, in some cases, no better or even worse than laypeople in detecting hallucinations raises concerns about their readiness for real-world applications where accuracy is paramount. This suggests that while LLMs can be powerful tools, they may introduce risks when used without proper oversight, particularly in contexts requiring specialized knowledge. Conducting longitudinal studies to track LLMs' susceptibility to hallucinations over time, particularly as they are exposed to new data and contexts, will be crucial in understanding how these models evolve and whether their performance in detecting hallucinations improves.

Future Works

Looking forward, we propose several key directions for future research:

Mitigating Hallucination through Adaptation. MEDHALU offers a rich corpus for fine-grained LLM hallucination detection. Fine-tuning LLMs using parameter-efficient techniques, such as LoRA (Hu et al.) and QLoRA (Dettmers et al. 2024), on MEDHALU can improve their reliability in real-world healthcare queries. Meanwhile, combining LLMs with rule-based systems or knowledge graphs that encode expert knowledge can mitigate hallucination risks by cross-referencing response with verified medical information.

Enhancing Expert Feedback Loops. Building on our proposed *expert-in-the-loop* approach, future work could focus on refining mechanisms that allow LLMs to continuously learn from expert feedback. This could involve interactive systems where LLMs not only generate responses but also seek validation or corrections from experts in real-time.

Extension to Multilingual and Multimodal Scenarios. While our study primarily focuses on English medical queries in textual formats, future research can explore how LLMs handle inaccurate information in non-English languages and LLM hallucination under alternative modalities (e.g. medical videos and broadcast).