

This is a repository copy of *Do LLMs provide consistent answers to health-related questions across languages*?

White Rose Research Online URL for this paper: <u>https://eprints.whiterose.ac.uk/225646/</u>

Version: Accepted Version

Proceedings Paper:

Schlicht, I.B. orcid.org/0000-0002-5037-2203, Zhao, Z. orcid.org/0000-0002-3060-269X, Sayin, B. orcid.org/0000-0001-6804-127X et al. (2 more authors) (2025) Do LLMs provide consistent answers to health-related questions across languages? In: Hauff, C., Macdonald, C., Jannach, D., Kazai, G., Nardini, F.M., Pinelli, F., Silvestri, F. and Tonellotto, N., (eds.) Advances in Information Retrieval: 47th European Conference on Information Retrieval, ECIR 2025, Lucca, Italy, April 6–10, 2025, Proceedings, Part III. 47th European Conference on Information Retrieval, ECIR 2025, 06-10 Apr 2025, Lucca, Italy. Springer Nature Switzerland , pp. 314-322. ISBN 9783031887130

https://doi.org/10.1007/978-3-031-88714-7_30

© 2025 The Authors. Except as otherwise noted, this author-accepted version of a journal article published in Advances in Information Retrieval: 47th European Conference on Information Retrieval, ECIR 2025, Lucca, Italy, April 6–10, 2025, Proceedings, Part III is made available via the University of Sheffield Research Publications and Copyright Policy under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: https://creativecommons.org/licenses/

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



Do LLMs Provide Consistent Answers to Health-Related Questions across Languages?

Ipek Baris Schlicht^{1,2}, Zhixue Zhao³, Burcu Sayin⁴, Lucie Flek^{2,5}, and Paolo Rosso^{1,6}

¹ Universitat Politècnica de València, Spain {ibarsch@doctor,prosso@dsic}.upv.es
² Bonn-Aachen International Center for IT, Germany {ibarissc,flek@bit}.uni-bonn.de
³ University of Sheffield, UK zhixue.zhao@sheffield.ac.uk
⁴ University of Trento, Italy burcu.sayin@unitn.it
⁵ Lamarr Institute for Machine Learning and Artificial Intelligence, Germany

⁶ ValgrAI Valencian Graduate School and Research Network of Artificial Intelligence,

 Spain

Abstract. Equitable access to reliable health information is vital for public health, but the quality of online health resources varies by language, raising concerns about inconsistencies in Large Language Models (LLMs) for healthcare. In this study, we examine the consistency of responses provided by LLMs to health-related questions across English, German, Turkish, and Chinese. We largely expand the HealthFC dataset by categorizing health-related questions by disease type and broadening its multilingual scope with Turkish and Chinese translations. We reveal significant inconsistencies in responses that could spread healthcare misinformation. Our main contributions are 1) a multilingual health-related inquiry dataset with meta-information on disease categories, and 2) a novel prompt-based evaluation workflow that enables sub-dimensional comparisons between two languages through parsing. Our findings highlight key challenges in deploying LLM-based tools in multilingual contexts and emphasize the need for improved cross-lingual alignment to ensure accurate and equitable healthcare information.

Keywords: Multilingual Q&A \cdot Healthcare Misinformation \cdot Consistency Evaluation \cdot Large Language Models

1 Introduction

LLMs have become powerful tools for NLP tasks and become widely accessible through chat interfaces to non-technical users. Their applications also extend to healthcare, where users consult LLM-based chat applications for health-related questions. LLMs are primarily trained on data from online sources. However, the quality of health information online varies by language, reflecting differences

in national healthcare policies and practices. For example, non-English online health content often has lower quality [20,11,5]. This could result in inconsistencies in health-related information provided by LLM-based applications across different languages.

In this paper, we examine inconsistencies between responses provided in English and lower-resourced languages. These inconsistencies can introduce factual errors and may lead to biased medical guidance, potentially harming specific communities by perpetuating misinformation in one language while omitting it in another. We particularly investigate the following research questions: **RQ1**. How do responses generated by an LLM differ when a question is asked in a non-English language compared to English?, and **RQ2**. Are there specific diseaserelated topics where LLM responses differ significantly across languages? To address these RQs, we first expand a health-related information-seeking dataset [19] (originally available in English (EN) and German (DE)) to include Turkish (TR) and Chinese (ZH), categorizing responses by disease type. Next, we introduce a prompt-based evaluation framework, illustrated in Figure 1, which segments the long-form responses of LLMs into informative parts to pinpoint areas with the highest occurrence of inconsistencies. Our source code is available online⁷.

2 Related Works

Various studies have examined how prompt variations impact the factual consistency of LLMs in EN medical Q&A. Zuccon and Koopman [10] showed that the accuracy of ChatGPT can vary widely based on prompt phrasing, with supporting or contradicting evidence in the prompt potentially introducing biases that affect answer correctness. Similarly, Sayin et al. [17] explored the potential of LLMs to assist and correct physicians in medical decision-making. They

⁷ http://bit.ly/4gQDJzT



Fig. 2. Distribution of 11 disease categories, each with at least 20 samples.

Disease Category	\mathbf{Count}
Symptoms, Signs & Abnormal Findings	83
Neoplasms (Cancer)	61
Mental & Behavioral Disorders	46
Diseases of the circulatory system	42
Diseases of the respiratory system	36
Diseases of the nervous system	31
Musculoskeletal & Connective Tissue Diseases	29
Etiology/Emergency use	28
Injury, Poisoning & External Causes	26
Diseases of the digestive system	22
Endocrine & Metabolic Diseases	22

Fig. 1. The proposed framework to automatically measure the consistency of health-related answers in EN, ZH, TR and DE.

highlighted the role of prompt engineering in enhancing LLM interactions with medical experts, demonstrating its impact on LLMs' ability to correct physician errors, explain medical reasoning, adapt to physician input, and improve overall performance. Kaur et al. [9] examined how presuppositions within prompts affect the factual accuracy and consistency of responses across LLMs such as ChatGPT and GPT-4. They showed that while LLMs rarely contradict established health facts, they often do not challenge false claims. Concerning multilingual Q&A, Jin et al. [8] assessed ChatGPT's response accuracy and consistency in multiple languages, translating non-EN responses into EN for evaluation. However, this method risks losing language-specific nuances due to its reliance on EN translations. While Jin et al. [8] evaluated consistency within individual languages, we assess relative consistency across languages, focusing on EN (a high-resource language in multilingual LLM pretraining) compared to other languages.

3 Methodology

To evaluate LLM performance across prompts in different languages, we first use Named Entity Recognition (NER) to categorize samples by disease. We then prompt LLMs in EN, DE, ZH, and TR on these categorized corpora.

3.1 Information Seeker Queries

We used the HealthFC dataset [19], available in DE and EN, containing 750 question-form claims along with corresponding evidence. HealthFC lacks specified disease categories for the claims, and some claims cover other health-related topics. To address this, we identified disease-related claims and enriched HealthFC with disease categories. Additionally, we expanded it by adding TR and ZH translations to support multilingual experiments.

Disease Categorization. To identify disease-related claims in the dataset, we used HunFlair2 [16] from flairNLP [2] on EN claims. We selected samples containing disease entities, resulting in 508 samples. We used the ICD10 Coding [21] for the categorization. We manually reviewed each entity, assigned the corresponding ICD10 codes, and then tagged them with their main disease category through the ICD10 look-up table [22]. For example, the ICD10 code of skin cancer is C44.90 and hence its main disease category is Neoplasms according to the table. After tagging, we kept only samples whose category occurred more than 20 times to maintain balance in disease categories (see Table 2).

Query Translations HealthFC includes samples in DE and EN. To expand it, we added TR and ZH samples by translating the EN samples with Google Translate. Native TR and ZH speakers from our team reviewed 50 random translations to ensure fluency and accuracy to preserve the original meaning. Specifically, they checked if the translated sentences and medical terms were correct and the translations retained the original meaning. Despite noting areas for improvement—such as using folk language instead of medical terms or adjusting culturally specific references (e.g., replacing 'pig meat' with 'sucuk' in TR), 84% of the TR translations and 98% of the ZH translations were satisfactory.

3.2 Prompt-based Evaluation Workflow

LLMs generate complex long-form answers, unlike simpler formats like multiplechoice responses. Long answers may include not only direct responses but also background information, examples, limitations, and related content, making direct comparisons challenging. To address this, we developed a prompt-based evaluation workflow that automatically analyzes bilingual answer nuances by segmenting long-form responses into informative parts. This workflow includes a parsing prompt and a consistency-check prompt.

Parsing Prompt. This prompt parses the answer based on a discourse structure, inspired by prior work by Xu et al. [24] and adapted for our use case. The discourse ontology comprises the following elements: (i) Answer-Summary (AS): The part of the answer addressing the question, excluding sentences elaborating on the summary or providing extra context, (ii) Health Benefits and Outcomes (HBO): Describes the positive effects or results of a medical intervention or behavior, (iii) Clinical Guidelines and Evidence (CGE): Refers to established guidelines or research that support the medical recommendations, (iv) Individual Considerations/Caveats (ICC): Highlights individual variability and emphasizes the need for personalized advice, and (v) Public Health/Professional Advice (PHPA): Emphasizes consulting healthcare professionals and following public health recommendations. A parsing example is given in Figure 3.

Consistency Comparison Prompt. This prompt compares parsed EN answers with those in other languages, assessing whether they are consistent. Consistency is evaluated based on the following annotation schema: (i) *Consistent:* The EN answer and the answer in the other language are fully consistent and semantically aligned, (ii) *Partially Consistent:* The EN answer and the answer in the other language partially agree, overlap, or support each other, though with some irrelevant or contradictory content, (iii) *Contradict:* Answers contradict each other, and (iv) *Irrelevant:* Answers address different topics and are unrelated. Empty responses are also included in this irrelevant category.

We refined prompts iteratively and evaluated the final parsing prompt using a point-based system across five answers per language. Elements were scored: 2



Fig. 3. The answers to the question "Does discontinuation of denosumab increase the risk of vertebral fractures?": the EN answer on the left and the TR answer on the right. Both were parsed by GPT-40 using the proposed discourse ontology.

points for correct classification, 1 point for partial parsing, with a maximum of 10. Average scores were 8.4 (EN, TR), 9.4 (ZH), and 7.2 (DE). We used final prompts ⁸to parse and annotate the answers, and randomly selected 50 samples. The team annotated the parsed answers in their native languages based on the inconsistency schema. We obtained 250 annotations for DE and TR, and 240 annotations for ZH⁹. We compared human annotations with LLM responses and observed disagreements between human and model annotations, especially within the inconsistent labels (partially inconsistent, irrelevant, contradiction). For automated evaluation, we consolidated these into a single 'inconsistent' label, which resulted in higher Kappa Scores [13] between human and model evaluations across languages. The final Kappa scores showed substantial agreement for TR (K=0.66) and ZH (K=0.71), and moderate agreement for DE (K=0.50).

4 Experiments and Results

We included four general-purpose, multilingual LLMs in the experiments: the closed-source models ChatGPT-40, continuously updated model [14], and GPT4o from OpenAI [1], and the open-source models Llama3-70B-Instruction [6] from Meta and CommandR+ from Cohere [4]. To generate answers, we set the LLMs' temperature to 0 and the response context limit to 2048 tokens. Prompts were given in EN, as the models tend to follow task descriptions more effectively in EN than in other languages [12,3], and EN is more cost-effective due to the larger token size required by some languages. No system prompts were used, except for Llama and CommandR+ when generating responses in TR and ZH, as these models default to EN. We accessed the closed models via the official OpenAI API. We used the Huggingface Inference Endpoint for Llama3 and the Transformers library [23] for CommandR+. Additionally, GPT-40 was used for parsing and comparing responses, chosen for its strong performance on various NLP tasks and its support for structured outputs essential to our analysis. (**RQ1**) To evaluate inconsistencies across different sections of the responses, we

Table 1. % of inconsistent responses by information units (AS: Answer Summary, HBO: Health Benefits and Outcomes, CGE: Clinical Guidelines and Evidence, ICC: Individual Considerations/Caveat, PHPA: Public Health/Professional Advice.), with frequent variations when questions are asked in non-EN languages.

	Chinese				Turkish				German			
	ChatGPT	GPT4c	Llama3	$\mathbf{CommandR}+$	ChatGPT	GPT40	Llama3	$\mathbf{CommandR}+$	ChatGPT	GPT40	Llama3	$\mathbf{CommandR}+$
AS	19.95	24.18	31.22	37.56	17.61	28.64	31.46	35.21	14.55	23.24	25.59	27.23
нво	36.15	40.61	57.51	61.74	44.37	46.48	61.74	62.91	33.10	30.75	51.88	54.23
CGE	47.65	42.02	69.01	71.83	53.05	49.06	77.93	73.00	40.61	38.03	73.94	65.26
ICC	66.43	62.91	82.16	81.22	65.73	69.25	84.98	86.39	64.08	60.09	71.83	77.70
PHPA	45.07	43.43	71.13	76.76	54.93	47.89	81.92	83.10	50.47	47.42	65.73	80.75
Average	43.05	42.63	62.21	65.82	47.34	48.26	67.61	68.12	40.56	39.91	57.79	61.03

⁸ Final prompts can be accessed in our source code.

⁹ Two examples were skipped by the annotator due to an error by Google Form

calculated the percentages of consistent and inconsistent samples as classified by GPT-40. Although LLMs generally produced consistent answer summaries, we observed significant inconsistencies in other parts of the responses (see Table 1). Consistent answer sections are mainly related to individual considerations, public health, or professional advice. (**RQ2**) To pinpoint the disease categories where LLMs produced the most inconsistent answers, we analyzed the top three categories for each information unit with inconsistencies. LLMs generated inconsistent answers across languages, about digestive systems, and endocrine and metabolic diseases. Inconsistencies are notable in an average of 71.97% of answers referencing guidelines and evidence, and in 59.09% of sentences within the latter disease category referencing health benefits and outcomes.

Table 2. Inconsistent Examples between the English answers and the answers in otherlanguages with their English translations.





Fig. 4. The LLMs tend to produce lengthy answers to the EN and DE questions.



Fig. 5. Answer occurence across the categories. TR answers omit some categories.

Finally, we performed a qualitative analysis on 50 samples per language used for human annotation. EN responses often included more detailed information, such as references to research organizations and studies, specific findings, and precise examples, compared to other languages. Interestingly, this level of detail appeared even in cases deemed semantically consistent, which was not accounted for in our evaluation schema. Moreover, when answers included statistical or numerical details (e.g., research years or quantities), the values in EN responses often differed from those provided in other languages. We also noted contradictory/irrelevant claims within a single response or between responses across different languages. Examples from the analysis are provided in Table 2.

Discussion. As shown in Figure 4, LLMs, particularly ChatGPT and Llama3, tend to generate longer answers in EN and DE. In contrast, TR answers often omit key elements like CGE, ICC, and PHPA, likely due to differences in training data and communication culture. For example, this may be affected by differences in the information available in different languages on Wikipedia [15] which is a common training source, or EN and DE often include more contextual detail due to their use in low-context cultures [7]. The closed models deliver more consistent answers, possibly due to regular updates and diverse human feedback.

Limitations. The framework currently does not assess factual accuracy, as ground-truth data is only available in EN and DE. Creating such data requires medical expertise because treatments for some diseases (e.g., depression [18]) can vary across languages. In addition, parsing and labeling rely on GPT-40, which may affect reproducibility. Despite human evaluation, errors in parsing and labeling could arise, particularly for information outside the current ontology. Furthermore, the single-question evaluation style restricts broader applicability, and GPT-40 struggles to detect fine-grained inconsistencies compared to human annotations. Finally, the analysis is limited to three non-EN languages, though the framework could be extended to others with expert support.

5 Conclusion

Our evaluation framework revealed notable differences between the answers in EN and other languages when addressing identical questions. These differences highlight the need for careful consideration of fairness implications when LLMs are deployed for medical applications, as they may affect equity in health communication and the potential for bias. In future work, we will explore open LLMs instead of GPT-40 to assess fine-grained inconsistencies and enhance the evaluation schema with detailed comparisons of information levels across languages.

Acknowledgments. The work of IBS and LF was supported by DynSoDA (funded by BMBF) and the Lamarr Institute for Machine Learning and Artificial Intelligence. The work of PR was supported by MCIN/AEI and by ERDF/EU (Grant No. PID2021-124361OB-C31). BS acknowledges support from the Horizon Europe project TANGO (Grant Agreement no. 101120763). Funded by the European Union.

Disclosure of Interests. Views and opinions expressed are however those of the author(s) only and do not reflect those of the funding organizations, EU or HaDEA.

References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
- Akbik, A., Bergmann, T., Blythe, D., Rasul, K., Schweter, S., Vollgraf, R.: FLAIR: an easy-to-use framework for state-of-the-art NLP. In: NAACL-HLT (Demonstrations). pp. 54–59. Association for Computational Linguistics (2019)
- Bareiß, P., Klinger, R., Barnes, J.: English prompts are better for nli-based zeroshot emotion classification than target-language prompts. In: WWW (Companion Volume). pp. 1318–1326. ACM (2024)
- Cohere: Command r plus microsoft azure. https://cohere.com/blog/ command-r-plus-microsoft-azure
- 5. Davaris, M., Barnett, S., Abouassaly, R., Lawrentschuk, N., et al.: Thoracic surgery information on the internet: a multilingual quality assessment. Interactive Journal of Medical Research 6(1), e6732 (2017)
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al.: The llama 3 herd of models. arXiv preprint arXiv:2407.21783 (2024)
- 7. Hall, E.T.: Understanding cultural differences. Intercultural P (1990)
- Jin, Y., Chandra, M., Verma, G., Hu, Y., Choudhury, M.D., Kumar, S.: Better to ask in english: Cross-lingual evaluation of large language models for healthcare queries. In: WWW. pp. 2627–2638. ACM (2024)
- Kaur, N., Choudhury, M., Pruthi, D.: Evaluating large language models for healthrelated queries with presuppositions. In: ACL (Findings). pp. 14308–14331. Association for Computational Linguistics (2024)
- Koopman, B., Zuccon, G.: Dr ChatGPT tell me what I want to hear: How different prompts impact health answer correctness. In: EMNLP. pp. 15012–15022. Association for Computational Linguistics (2023)
- Lawrentschuk, N., Abouassaly, R., Hewitt, E., Mulcahy, A., Bolton, D., Jobling, T.: Health information quality on the internet in gynecological oncology: a multilingual evaluation. Eur. J. Gynaecol. Oncol 37(4), 2016
- Lupo, L., Magnusson, O., Hovy, D., Naurin, E., Wängnerud, L.: How to use large language models for text coding: The case of fatherhood roles in public policy documents. arXiv preprint arXiv:2311.11844 (2023)
- McHugh, M.L.: Interrater reliability: the kappa statistic. Biochemia medica 22(3), 276–282 (2012)
- OpenAI: Openai documentation (2024), https://platform.openai.com/docs/ models/continuous-model-upgrades,
- Roy, D., Bhatia, S., Jain, P.: Information asymmetry in wikipedia across different languages: A statistical analysis. Journal of the Association for Information Science and Technology 73(3), 347–361 (2022)
- Sänger, M., Garda, S., Wang, X.D., Weber-Genzel, L., Droop, P., Fuchs, B., Akbik, A., Leser, U.: Hunflair2 in a cross-corpus evaluation of biomedical named entity recognition and normalization tools. Bioinformatics 40(10), btae564 (2024)
- Sayin, B., Minervini, P., Staiano, J., Passerini, A.: Can llms correct physicians, yet? investigating effective interaction methods in the medical domain. In: ClinicalNLP@NAACL. pp. 218–237. Association for Computational Linguistics (2024)
- Skianis, K., Pavlopoulos, J., Doğruöz, A.S.: Severity prediction in mental health: Llm-based creation, analysis, evaluation of a novel multilingual dataset. arXiv preprint arXiv:2409.17397 (2024)

- Vladika, J., Schneider, P., Matthes, F.: HealthFC: Verifying health claims with evidence-based medical fact-checking. pp. 8095–8107. ELRA and ICCL, Torino, Italia (May 2024)
- Weissenberger, C., Jonassen, S., Beranek-Chiu, J., Neumann, M., Müller, D., Bartelt, S., Schulz, S., Mönting, J., Henne, K., Gitsch, G., et al.: Breast cancer: patient information needs reflected in english and german web sites. British Journal of Cancer 91(8), 1482–1487 (2004)
- 21. WHO: The ICD-10 classification of mental and behavioural disorders: diagnostic criteria for research, vol. 2. World Health Organization (1993)
- 22. WHO: Icd10 guideline (2019), https://icd.who.int/browse10/2019/en,
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T.L., Gugger, S., Drame, M., Lhoest, Q., Rush, A.M.: Transformers: State-of-the-art natural language processing. In: EMNLP (Demos). pp. 38–45. Association for Computational Linguistics (2020)
- Xu, F., Li, J.J., Choi, E.: How do we answer complex questions: Discourse structure of long-form answers. In: ACL (1). pp. 3556–3572. Association for Computational Linguistics (2022)