

Evaluation of an LLM in Identifying Logical Fallacies

Gionnieve Lim

gionnievelim@gmail.com

Singapore University of Technology and Design

Singapore

Simon T. Perrault

perrault.simon@gmail.com

Singapore University of Technology and Design

Singapore

ABSTRACT

There is increasing interest in the adoption of LLMs in HCI research. However, LLMs may often be regarded as a panacea because of their powerful capabilities with an accompanying oversight on whether they are suitable for their intended tasks. We contend that LLMs should be adopted in a critical manner following rigorous evaluation. Accordingly, we present the evaluation of an LLM in identifying logical fallacies that will form part of a digital misinformation intervention. By comparing to a labeled dataset, we found that GPT-4 achieves an accuracy of 0.79, and for our intended use case that excludes invalid or unidentified instances, an accuracy of 0.90. This gives us the confidence to proceed with the application of the LLM while keeping in mind the areas where it still falls short. The paper describes our evaluation approach, results and reflections on the use of the LLM for our intended task.

CCS CONCEPTS

• **Human-centered computing** → Interactive systems and tools; Empirical studies in HCI.

KEYWORDS

large language models, technical evaluation, classification, logical fallacies

1 BACKGROUND

LLMs have been rather successfully employed in various fields and for various applications. Their low barrier to entry, where only a prompt is needed at the highest level, makes them an appealing technology for adoption. Due to their powerful natural language capabilities, LLMs have been used for a variety of tasks including question answering, machine translation, generation, prediction and classification. Focusing on classification tasks, the applications are diverse. LLMs can be used to classify textual [20], tabular [7] and image [16] data. And they are employed for various operations like data annotation [11], sentiment analysis [21] and information extraction [5]. Leveraging on the classification capabilities of LLMs, we considered if they can be used to identify logical fallacies in texts, a function that would serve as part of a digital misinformation intervention.

While a simple test of prompting an LLM with various texts and asking it to identify logical fallacies led to *seemingly correct* results, we could not be convinced that the LLM was truly capable of this task. Aligned with the principles of trustworthy AI, particularly on the transparency of the performance of the AI [10], a more rigorous evaluation was called for. To that end, we looked at verifying the

suitability of the LLM for identifying logical fallacies by using a more rigorous approach of testing against a labeled dataset. The goal was to understand to what extent the LLM performs for the task and to obtain an accuracy metric that could be conveyed to users so that they are aware of the performance of the LLM and can adjust their expectations for the digital misinformation intervention.

2 THE LOGICAL FALLACIES

To understand whether LLMs are suitable for the task of identifying logical fallacies, we conducted an evaluation of the classification performance of an LLM on a limited set of logical fallacies. Since there are over 300 logical fallacies [1], we focused only on a reduced set which is more practical and replicable. We referred to work that classified, identified or used logical fallacies in online discussions [14, 15, 17] and misinformation [2, 4, 8, 9, 12, 13, 18, 19] literature and selected the most commonly mentioned fallacies (*Against the Person*, *Appeal to Authority*, *Appeal to Popularity*, *Appeal to Emotion*, *Hasty Generalization*, *Questionable Cause*, and *Red Herring*). Table 1 shows the definitions of these logical fallacies.

3 TECHNICAL EVALUATION

We used the LOGIC dataset by Jin et al. [9] that contained 2,449 logical fallacy instances across 13 logical fallacy types to evaluate the LLM. To select only the relevant data, we removed data that (1) were duplicates, (2) did not pertain to our set of logical fallacies, (3) only defined or described logical fallacies, (4) contained Latin phrases, and (5) contained quizzing phrases. The last was due to the LOGIC dataset being crawled from student quiz websites. This resulted in a reduced dataset ($N = 801$).

OpenAI's GPT-4 model¹ was used as the LLM for identifying logical fallacies. We used the Chat Completions API with the gpt-4 model, a temperature setting of 0, and the role: You are a critical thinker. Few-shot prompting [3] was used for the logical fallacies classification task where three examples of each of the seven logical fallacies were provided. These instances ($n = 21$) were taken from the dataset and excluded from the evaluation of the performance of the LLM. The breakdown of the full data used for evaluation ($N = 780$) was thus as follows: 157 *Against the Person*, 74 *Appeal to Authority*, 133 *Appeal to Popularity*, 41 *Appeal to Emotion*, 165 *Hasty Generalization*, 126 *Questionable Cause*, and 84 *Red Herring*. The prompt used for identifying logical fallacies is shown in Appendix A.

4 RESULTS

We present the multi-class classification performance metrics [6] of the LLM in the identification of logical fallacies in Table 2. When

1st HEAL Workshop at CHI Conference on Human Factors in Computing Systems, May 12, Honolulu, HI, USA 2024.

¹<https://platform.openai.com/docs/models/gpt-4>

Table 1: Common logical fallacies found in online discussions and misinformation.

Logical Fallacy (<i>in Latin</i>)	Definition
Against the Person (<i>Argumentum Ad Hominem</i>)	Attacking the person or some aspect of the person making the argument instead of addressing the argument directly.
Appeal to Authority (<i>Argumentum Ad Verecundiam</i>)	Using an alleged authority who is not really an authority on the facts relevant to the argument as evidence.
Appeal to Popularity (<i>Argumentum Ad Populum</i>)	Affirming that something is real or better because the majority in general or of a particular group thinks so.
Appeal to Emotion (<i>Argumentum Ad Passiones</i>)	Manipulating the reader’s emotions in order to win the argument in place of a valid reason.
Hasty Generalization (<i>Secundum Quid</i>)	Drawing a conclusion about all or many instances of a phenomenon on the basis of one or a few instances of that phenomenon.
Questionable Cause (<i>Non Causa Pro Causa</i>)	Concluding that one thing caused another simply because they are regularly associated.
Red Herring (<i>Ignoratio Elenchi</i>)	Attempting to divert the reader’s attention from the original argument by offering a different point.

prompting, we allowed the LLM to answer *None* if no logical fallacy was identified in a text. In rare instances where the LLM answered with a logical fallacy outside our considered set, we also post-processed the result to *None*. As such, we further present the performance metrics on the subset ($N = 685$) of the data where instances predicted or post-processed as *None* ($n = 95$) were removed.

Table 2: Multi-class classification performance of the LLM in identifying logical fallacies for the full data ($N = 780$) and subset data ($N = 685$, where 95 instances of *None* were removed).

Data	Accuracy	Macro Precision	Macro Recall	Macro F1 Score
Full	0.79	0.76	0.66	0.71
Subset	0.90	0.87	0.89	0.88

The LLM achieved an average accuracy of 0.79 in classifying logical fallacies for the full data. Many of the inaccuracies occurred in failing to identify a valid logical fallacy, e.g., classifying a text as *None* (Table 3). Most strikingly, the LLM performed worst in identifying instances of the *Appeal to Emotion* fallacy, failing to do so nearly 40% of the time. As LLMs struggle with extracting fine-grained structured sentiment and opinion information [21], this may explain the poor performance for *Appeal to Emotion* which demands stronger nuanced interpretations of texts from the LLM.

When excluding the instances classified as *None*, the LLM achieved an average accuracy of 0.90 on the subset data. Figure 1 shows the normalized classification matrix of the subset data which provides a clearer picture of the misclassifications of one logical fallacy as another. The LLM performed best for *Against the Person* with an accuracy of 0.94 and generally had good performance across the other fallacies, with the lowest accuracy being 0.84. The main misclassifications of logical fallacies occurred for *Appeal to Emotion*

where 12% were classified as *Red Herring* and for *Questionable Cause* where 11% were classified as *Hasty Generalization*. While the inadequacies of LLMs in making nuanced interpretations may explain the former, the latter could have been due to both fallacies being logically different yet sharing semantically similar definitions (see the definitions in Appendix A).

5 DISCUSSION

We describe our views of the LLM for our use case of identifying logical fallacies and reflect on the design of the evaluation procedure.

5.1 Evaluate to Inform Design Decisions

While the LLM failed to identify valid logical fallacies for more than a tenth of a time (Table 3), it performed well in classifying logical fallacies (Figure 1). Since our planned use case prioritizes categorizing the *right* logical fallacy and not surfacing *all* of them, we consider the LLM to be capable and suitable for our task definition.

Without thorough evaluation of the LLM, we could not have had the same level of confidence in choosing to adopt the LLM for our use case. While the results turned out to be in our favor this time, it would not be difficult to imagine how unfavorable results would impact our next steps instead. For instance, if our intervention prioritized finding *all* logical fallacies, we would have made the decision to exclude the *Appeal to Emotion* fallacy going forward due to the subpar performance of the LLM in spotting it.

Such decisions are not naive, underscoring the importance of evaluating LLMs rigorously if they are used as a key component of the research project to inform the design decisions of the artifact or use case where the LLM is adopted.

5.2 Considerations in Evaluating LLMs

There were several considerations we faced in the design of the evaluation procedure and we share the more prominent ones here.

Table 3: Breakdown of instances in the full data ($N = 780$) that were classified as *None*.

	Against the Person	Appeal to Authority	Appeal to Popularity	Appeal to Emotion	Hasty Generalization	Questionable Cause	Red Herring	Total
<i>None</i>	7	4	14	16	27	17	10	95
All	157	74	133	41	165	126	84	780
Percentage	0.04	0.05	0.11	0.39	0.16	0.13	0.12	0.12

**Figure 1: Normalized confusion matrix of the LLM in identifying logical fallacies for the subset data ($N = 685$).**

Evaluation Strategy. The first is with choosing the evaluation strategy. Our choice to evaluate the LLM against a labeled dataset was inspired by research from the NLP community where factors like ease of access to existing datasets and convenience in the evaluation process were taken into account. Other evaluation strategies can also be gleaned from the NLP community, which we mention briefly for researchers interested in conducting their own evaluations. A strategy is to engage raters, either humans or automated, to rate the outputs of the LLM according to pre-specified metrics based on theoretical frameworks, typically using Likert scales. Another strategy is to find existing benchmarks that are applicable to the intended use case and to run the outputs of the LLM against them.

Dataset. The second is with finding the appropriate dataset. The dataset we used contained mostly instances with one or two sentences and occasionally those with a paragraph. Yet, for our planned use of the LLM, we intend to pass online discussion comments that can range from a few sentences to a few paragraphs. While our evaluation shows that the LLM performs acceptably for shorter texts, the same may not be expected for longer texts which

are likely to contain more logical fallacies. This would merit further testing, although to the best of our knowledge, there is no existing labeled dataset on logical fallacies for long texts which would necessitate manually creating one. Nonetheless, we will acknowledge this as a limitation of the evaluation of the LLM when we report about its adoption in the digital misinformation intervention.

Prompt Engineering. The last is with settling on the prompt. While not reported here, several iterations of the prompt were made before the final version. While prompts can be easily created, what is more difficult is deciding when to *stop* engineering. While we initially focused on crafting a prompt that would achieve the best performance, we reflected on our priorities as HCI researchers and felt that more so than the best performance, what is important is to design the interactions of the artifact with an awareness of the capabilities and insufficiencies of the LLM for the task. Refocusing, we aimed for a prompt that would deliver adequate performance, as presented in this work, and to gain a comprehensive understanding of the extent of its performance. Here, we also want to highlight how designing and evaluating LLMs prompts became an iterative

process for us where an established evaluation protocol lended towards knowing when to call a halt.

6 CONCLUSION

While the use of LLMs for HCI research is gaining interest, the question of whether LLMs are suitable for their intended tasks remains. HCI researchers should go beyond incorporating the LLM merely because it “seems to work” and instead conduct a thorough evaluation for its specific use. In line with this, we present an evaluation of the use of an LLM in identifying logical fallacies, describing our approach, results and reflections. Keeping in mind that the LLM performs adequately, though not perfectly, for this task, we intend to proceed with using the LLM for our research, being transparent about its performance with the users.

ACKNOWLEDGMENTS

This work was supported by resources from the KAIST Interaction Lab (KIXLAB).

REFERENCES

- [1] B. Bennett. 2012. *Logically Fallacious: The Ultimate Collection of Over 300 Logical Fallacies (Academic Edition)*. Ebookit.com.
- [2] Claire Bonial, Austin Blodgett, Taylor Hudson, Stephanie M. Lukin, Jeffrey Micher, Douglas Summers-Stay, Peter Sutor, and Clare Voss. 2022. The Search for Agreement on Logical Fallacy Annotation of an Infodemic. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, Nicoletta Calzolari, Frédéric Béchet, Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Jan Odijk, and Stelios Piperidis (Eds.). European Language Resources Association, Marseille, France, 4430–4438. <https://aclanthology.org/2022.lrec-1.471>
- [3] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS'20)*. Curran Associates Inc., Red Hook, NY, USA, Article 159, 25 pages.
- [4] John Cook, Peter Ellerton, and David Kinkead. 2018. Deconstructing climate misinformation to identify reasoning errors. *Environmental Research Letters* 13, 2 (feb 2018), 024018. <https://doi.org/10.1088/1748-9326/aaa49f>
- [5] Akshay Goel, Almog Gueta, Omry Gilon, Chang Liu, Sofia Erell, Lan Huong Nguyen, Xiaohong Hao, Bolous Jaber, Shashir Reddy, Rupesh Kartha, Jean Steiner, Itay Laish, and Amir Feder. 2023. LLMs Accelerate Annotation for Medical Information Extraction. In *Proceedings of the 3rd Machine Learning for Health Symposium (Proceedings of Machine Learning Research, Vol. 225)*, Stefan Hegselmann, Antonio Parziale, Divya Shanmugam, Shengpu Tang, Mercy Nyamewaa Asiedu, Serina Chang, Tom Hartvigsen, and Harvineet Singh (Eds.). PMLR, 82–100. <https://proceedings.mlr.press/v225/goel23a.html>
- [6] Margherita Grandini, Enrico Bagli, and Giorgio Visani. 2020. Metrics for Multi-Class Classification: an Overview. arXiv:2008.05756 [stat.ML]
- [7] Stefan Hegselmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David Sontag. 2023. TabLLM: Few-shot Classification of Tabular Data with Large Language Models. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 206)*, Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (Eds.). PMLR, 5549–5581. <https://proceedings.mlr.press/v206/hegselmann23a.html>
- [8] Timon M J Hruschka and Markus Appel. 2023. Learning about informal fallacies and the detection of fake news: An experimental intervention. *PLoS One* 18, 3 (March 2023), e0283238.
- [9] Zhijing Jin, Abhinav Lalwani, Tejas Vaidhya, Xiaoyu Shen, Yiwen Ding, Zhiheng Lyu, Mrinmaya Sachan, Rada Mihalcea, and Bernhard Schoelkopf. 2022. Logical Fallacy Detection. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 7180–7198. <https://doi.org/10.18653/v1/2022.findings-emnlp.532>
- [10] Bo Li, Peng Qi, Bo Liu, Shuai Di, Jingen Liu, Jiquan Pei, Jinfeng Yi, and Bowen Zhou. 2023. Trustworthy AI: From Principles to Practices. *ACM Comput. Surv.* 55, 9, Article 177 (jan 2023), 46 pages. <https://doi.org/10.1145/3555803>
- [11] Jiyi Li. 2024. A Comparative Study on Annotation Quality of Crowdsourcing and LLM via Label Aggregation. arXiv:2401.09760 [cs.CL]
- [12] Morgan Lundy. 2023. TikTok and COVID-19 Vaccine Misinformation: New Avenues for Misinformation Spread, Popular Infodemic Topics, and Dangerous Logical Fallacies. *International Journal of Communication* 17, 0 (2023). <https://ijoc.org/index.php/ijoc/article/view/19847>
- [13] Elena Musi and Chris Reed. 2022. From fallacies to semi-fake news: Improving the identification of misinformation triggers across digital media. *Discourse & Society* 33, 3 (2022), 349–370. <https://doi.org/10.1177/09579265221076609> arXiv:https://doi.org/10.1177/09579265221076609
- [14] Nikolaos Nikolaidis, Nicolas Stefanovitch, and Jakub Piskorski. 2023. On Experiments of Detecting Persuasion Techniques in Polish and Russian Online News: Preliminary Study. In *Proceedings of the 9th Workshop on Slavic Natural Language Processing 2023 (SlavicNLP 2023)*, Jakub Piskorski, Michał Marcińczuk, Preslav Nakov, Maciej Ogrodniczuk, Senja Pollak, Pavel Přibáň, Piotr Rybak, Josef Steinberger, and Roman Yangarber (Eds.). Association for Computational Linguistics, Dubrovnik, Croatia, 155–164. <https://doi.org/10.18653/v1/2023.bsnlp-1.18>
- [15] Didin Nuruddin Hidayat, Nurhalimah, Maya Defianty, Ummi Kultsum, Zulkifli, and Agus Sufyan. 2020. Logical Fallacies in Social Media: A Discourse Analysis in Political Debate. In *2020 8th International Conference on Cyber and IT Service Management (CITSM)*, 1–5. <https://doi.org/10.1109/CITSM50537.2020.9268821>
- [16] Sarah Pratt, Ian Covert, Rosanne Liu, and Ali Farhadi. 2023. What Does a Platypus Look Like? Generating Customized Prompts for Zero-Shot Image Classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 15691–15701.
- [17] Saumya Sahai, Oana Balalau, and Roxana Horincar. 2021. Breaking Down the Invisible Wall of Informal Fallacies in Online Discussions. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 644–657. <https://doi.org/10.18653/v1/2021.acl-long.53>
- [18] Zhivar Sourati, Vishnu Priya Prasanna Venkatesh, Darshan Deshpande, Himanshu Rawlani, Filip Ilievski, Hông-Ân Sandlin, and Alain Mermoud. 2023. Robust and explainable identification of logical fallacies in natural language arguments. *Know-Based Syst.* 266, C (apr 2023), 22 pages. <https://doi.org/10.1016/j.knosys.2023.110418>
- [19] Christian Schlereth Sven Beisecker and Sebastian Hein. 2024. Shades of fake news: how fallacies influence consumers’ perception. *European Journal of Information Systems* 33, 1 (2024), 41–60. <https://doi.org/10.1080/0960085X.2022.2110000> arXiv:https://doi.org/10.1080/0960085X.2022.2110000
- [20] Zhiqiang Wang, Yiran Pang, and Yanbin Lin. 2023. Large Language Models Are Zero-Shot Text Classifiers. arXiv:2312.01044 [cs.CL]
- [21] Wenxuan Zhang, Yue Deng, Bing Liu, Sinno Jialin Pan, and Lidong Bing. 2023. Sentiment Analysis in the Era of Large Language Models: A Reality Check. arXiv:2305.15005 [cs.CL]

A PROMPT FOR IDENTIFYING LOGICAL FALLACIES

There are seven fallacies. Each fallacy, its definition, and examples of texts with the fallacy are provided below.

Fallacy: hasty generalization

Definition: Drawing a conclusion about all or many instances of a phenomenon on the basis of one or a few instances of that phenomenon.

Examples:

- 1) "Four out of five dentists recommend Happy Glossy toothpaste. Therefore, it must be great."
- 2) No one in your family surfs? But I thought you said you lived in California before this.
- 3) All chess players are geniuses.

Fallacy: questionable cause

Definition: Concluding that one thing caused another simply because they are regularly associated.

Examples:

- 1) Children who play violent video games act more violently than those who don't.
- 2) "President Kumail raised taxes, and then the rate of violent crime went up. Kumail is responsible for the rise in crime."
- 3) People who eat yogurt have healthy guts. If I eat yogurt I will never get sick.

Fallacy: ad populum

Definition: Affirming that something is real or better because the majority in general or of a particular group thinks so.

Examples:

- 1) Everyone seems to support the changes in the vacation policy, and if everyone likes them, they must be good.
- 2) "Since 88% of people polled believe in UFOs, they must exist."
- 3) Don't be the only one not wearing Nike!

Fallacy: ad hominem

Definition: Attacking the person or some aspect of the person making the argument instead of addressing the argument directly.

Examples:

- 1) Only a selfish, non-caring person would believe that this is ok.
- 2) You can't believe Jack when he says there is a God because he doesn't even have a job.
- 3) The reason our company never makes any money is because we have a buffoon running it!

Fallacy: appeal to emotion

Definition: Manipulating the reader's emotions in order to win the argument in place of a valid reason.

Examples:

- 1) Regime vs. government Pro-death vs. pro-choice

- 2) "It is an outrage that the school wants to remove the vending machines. This is taking our freedom away!"
- 3) Television Advertisement: Get all of your stains out by using new and improved Ultra Suds and wash your blues away!

Fallacy: red herring

Definition: Attempting to divert the reader's attention from the original argument by offering a different point.

Examples:

- 1) I shouldn't be punished for staying out past my curfew because I did the dishes earlier today.
- 2) "Yes, we have safety issues in our factory. But we work really hard to make a good product. They sell so well!"
- 3) Mining may destroy the environment, but how about the people whose jobs will be affected? Should we ban it?

Fallacy: appeal to authority

Definition: Using an alleged authority who is not really an authority on the facts relevant to the argument as evidence.

Examples:

- 1) My professor, who has a Ph.D. in Astronomy, once told me that ghosts are real. Therefore, ghosts are real.
- 2) There is definitely a link between dementia and drinking energy drinks. I read about it on Wikipedia.
- 3) Coke is not as healthy for you as Pepsi. Besides, Britney Spears drinks Pepsi, so it must be healthier than Coke

For the text below, identify if there are any fallacies {hasty generalization, questionable cause, ad populum, ad hominem, appeal to emotion, red herring, appeal to authority}.

For each part of the text with a fallacy, use this template.
Part: {quote verbatim the part of the text with the fallacy}
Fallacy: {fallacy}

If the text does not contain any fallacy, use this template.
Part: none
Fallacy: none

The text:
--The text goes here--