

Predicting LLM Correctness in Prosthodontics Using Metadata and Hallucination Signals

Lucky Susanto¹^a, Anasta Pranawiyana², Cortino Sukotjo³^b, Soni Prasad⁴^c, and Derry Wijaya¹^d

¹*Department of Data Science, Monash University Indonesia, Tangerang, Indonesia*

²*Independent Researcher*

³*Department of Prosthodontics, University of Pittsburgh, Pittsburgh, Pennsylvania*

⁴*Department of Restorative Sciences, University of North Carolina Adams School of Dentistry, Chapel Hill, North Carolina*
lucky.susanto@monash.edu, adhibaye@gmail.com, COS162@pitt.edu, prasado@unc.edu, derry.wijaya@monash.edu

Keywords: Large Language Models, Hallucination, Correctness Prediction, Model Reliability, Prompting Strategies

Abstract: Large language models (LLMs) are increasingly adopted in high-stakes domains such as healthcare and medical education, where the risk of generating factually incorrect (i.e., hallucinated) information is a major concern. While significant efforts have been made to detect and mitigate such hallucinations, predicting whether an LLM's response is correct remains a critical yet underexplored problem. This study investigates the feasibility of predicting correctness by analyzing a general-purpose model (GPT-4o) and a reasoning-centric model (OSS-120B) on a multiple-choice prosthodontics exam. We utilize metadata and hallucination signals across three distinct prompting strategies to build a correctness predictor for each (model, prompting) pair. Our findings demonstrate that this metadata-based approach can improve accuracy by up to +7.14% and achieve a precision of 83.12% over a baseline that assumes all answers are correct. We further show that while actual hallucination is a strong indicator of incorrectness, metadata signals alone are not reliable predictors of hallucination. Finally, we reveal that prompting strategies, despite not affecting overall accuracy, significantly alter the models' internal behaviors and the predictive utility of their metadata. These results present a promising direction for developing reliability signals in LLMs but also highlight that the methods explored in this paper are not yet robust enough for critical, high-stakes deployment.


1 INTRODUCTION


Large language models (LLMs) have seen an increase in adoption, even in high-stakes domains such as healthcare. Despite their capabilities, LLMs are still prone to generating hallucinated content. This is particularly dangerous, as these hallucinations often sound logical yet contain non-factual information (Huang et al., 2025). This issue is particularly concerning in medical education, where LLMs are increasingly used as learning aids or evaluators. In such settings, hallucinated yet plausible explanations can misinform students or reinforce misconceptions, indirectly affecting clinical competence (Abd-alrazaq et al., 2023).


Due to the danger of hallucinations, many efforts


have been made to detect and mitigate them (Manakul et al., 2023; Kadavath et al., 2022). These works detect hallucinations by utilizing information such as consistency, log probability, entropy distribution, and many more. However, to our knowledge, previous work has not directly attempted to predict model correctness. As previously mentioned, predicting correctness is vital in medical education, where inaccurate or misleading model responses could negatively influence learners' understanding or clinical reasoning.

In this study, we analyze two distinct LLMs on a specialized, non-reasoning task (i.e., a multiple-choice question answering problem) in the field of prosthodontics. We examine how metadata, specifically response consistency and answer token log probability, can be used to predict whether the model will answer the question correctly or not. Moreover, we then compare different prompting strategies and how they affect these two different models. Through this work, we bring the following contributions:

^a <https://orcid.org/0009-0003-3494-9249>

^b <https://orcid.org/0000-0002-2171-004X>

^c <https://orcid.org/0000-0002-4297-5012>

^d <https://orcid.org/0000-0002-0848-4703>

1. We demonstrate that it is feasible to predict model correctness by only looking at its consistency and log probabilities, achieving an increased accuracy of up to +7.14% when compared to simply assuming LLMs are always correct, alongside a maximum precision of 83.12%.
2. We show that under an idealized setup, hallucination signals are strong indicators of model correctness. However, this requires a powerful hallucination detection model.
3. In relation to the previous finding, we also show that consistency and token log probabilities are not predictive of hallucination signals.
4. Lastly, we found that even though differing prompting strategies have no significant effect on model accuracy, different prompting strategies affect GPT-4o’s behavior more than OSS-120B’s, as reflected in the predictiveness of metadata.

2 RELATED WORKS

2.1 Reasoning LLMs

Reasoning LLMs are a relatively new technology. Popularized by DeepSeek-R1 (DeepSeek-AI et al., 2025), these LLMs show incredible promise in tasks such as math, coding, and other reasoning-heavy tasks (Comanici, 2025; Yang et al., 2025). Perhaps due to their novelty, there is only a minimal amount of work analyzing the effect of reasoning on non-reasoning tasks. A recently released paper notes that we should induce reasoning only for specific tasks requiring reasoning, as they found that a general LLM would be more efficient otherwise (Boizard et al., 2025).

2.2 Prompting Strategies

After the rise of LLMs, a particular prompting strategy rose above the rest: Chain-of-Thought prompting. Shortened to CoT, this prompting strategy has been shown to induce reasoning capabilities in LLMs (Wei et al., 2023). In fact, reasoning LLMs such as DeepSeek-R1 are trained through post-training on reasoning-oriented datasets containing CoT exemplars, designed to incentivize step-by-step reasoning.

However, recent work has cast doubt on this approach, as it has been shown that CoT prompting can be unfaithful (Turpin et al., 2023) and that the reasoning induced by CoT is potentially a mirage (Zhao et al., 2025). Specifically, it seems that CoT is used by the model to match the pattern of reasoning, not

to perform reasoning itself. In contrast, Answer-then-Explain (post-hoc explanation) is less popular, though there exist works that utilize it to improve model performance (Krishna et al., 2023).

Lastly, it has been reported that despite differing prompting strategies having the same average faithfulness score, they can actually have a high degree of disagreement in the explanations they generate (Huang et al., 2023). This shows that different prompting strategies affect the behavior of how LLMs produce their answers. However, this work has only been done on general LLMs, leaving room to explore this for reasoning LLMs.

2.3 LLM-as-a-Judge

Due to the impressive capabilities of modern LLMs, a new paradigm has recently emerged. Known as “LLM-as-a-judge” (Li et al., 2025), this method utilizes an LLM to act as a judge that rates the output of another model based on specific metrics. Of course, it is possible that the judge itself is prone to error, thus choosing a reliable model is a must. Even when a reliable model is chosen, it is imperative to understand that LLMs are not perfect.

3 METHODOLOGY

3.1 Dataset

We use the 2025 National Prosthodontics Resident Exam questions in this work. The exam originally consists of 150 multiple-choice questions. We remove questions that use books as the source of reference due to the difficulty of extracting ground truth explanations. Furthermore, we exclude questions that include images. After filtering, we end up with a final dataset of 98 questions.

To obtain the ground truth explanation, we use GPT-4o to extract an explanatory paragraph with respect to a given question and answer. The extracted paragraph is then evaluated by two experts in the field of prosthodontics, ensuring its quality. Any incorrectly extracted paragraphs were then manually fixed by the authors.

3.2 Models

We evaluate two models in this work:

1. **GPT-4o** (OpenAI et al., 2024): A proprietary generalist LLM

2. **OSS-120B** (OpenAI et al., 2025): An open-source reasoning LLM

3.3 Prompting Strategy

We use three types of prompting in this work:

1. **Answer only (A):** Provided a question, the model answers directly.
2. **Answer then Explain (AE):** Provided a question, the model answers and then provides a justification.
3. **Explain then Answer (EA):** Commonly known as Chain-of-Thought; provided a question, the model reasons to reach an answer.

We use the following system prompt to enforce a consistent answer format for both models:

```
Instruction: You must answer this
question using the exact template
below
Response:
Answer: Only the selected option (e.g.,
A [without anything else])
Explanation: Your explanation
```

To change the prompting strategy, we modify the system prompt as follows:

1. Leave it unchanged for **Answer then Explain**.
2. Move the "Explanation" line above the "Answer" line for **Explain then Answer**.
3. Delete the "Explanation" line for the **Answer only** strategy.

3.4 Evaluation

For each question in the dataset, we prompt each model five times. We then take the majority answer as the model's final answer.

3.5 Hallucinations

We use **LLM-as-a-judge** (Li et al., 2025), with Gemini-2.5-flash (Comanici, 2025) as a judge. This model is chosen due to its capabilities, with a reported hallucination rate of 1.3% on the Vectara hallucination benchmark (Bao et al., 2024)¹. The prompt is available in Appendix 7. Note that the Vectara hallucination benchmark is not specific to the prosthodontics domain and thus this number may be an overestimate of the model's capabilities.

In this work, we focus on three types of hallucinations:

¹ As reported on the Vectara LLM Hallucination Leaderboard.

1. **Intrinsic Hallucination:** The explanation provided by the model directly contradicts the reference explanation.
2. **Extrinsic Hallucination:** The explanation provided by the model contains unverifiable information with respect to the reference explanation.
3. **Logical Hallucination:** The explanation provided by the model does not support its final answer.

We further refine the definition for **Extrinsic Hallucination**:

1. **Safe:** There is no unverifiable information from the model's generated explanation.
2. **Benign:** The unverifiable information generated by the model is superfluous and the core premise does not differ from the reference explanation.
3. **Harmful:** The unverifiable information generated by the model has a core premise that differs from the reference explanation.

Examples illustrating the difference between classes of extrinsic hallucination are provided in Appendix 7. In this work, we group extrinsic hallucinations into two categories: harmful and non-harmful (i.e., safe or benign). We made this choice because our setup limits the reference explanation to a single paragraph. However, this simplification introduces a limitation: a model's explanation may be factual yet unverifiable within the relatively short reference. This results in a higher occurrence of benign extrinsic hallucinations that would have been deemed safe if the reference explanation were more complete.

3.6 Metadata

This work chooses two types of metadata as the focus:

1. **Consistency:** How often each model outputs its most common answer across five attempts.
2. **Answer token log probability:** A proxy for the model's confidence when deciding which choice it should answer with (i.e., X in "Answer: [X]"). A log probability of 0 denotes absolute confidence.

4 EXPLORATORY ANALYSIS

4.1 Model Performance

Despite previous works reporting Reasoning LLMs' higher capabilities in domains such as math, coding, and other reasoning-heavy tasks, we found results indicating that this ability does not help in a




Figure 1: Different prompting strategies have a non-significant effect on model accuracy.

non-reasoning task, supporting results from previous works (Boizard et al., 2025).

4.2 Consistency




Figure 2: Different prompting strategies have a non-significant effect on model consistency.

We find that GPT-4o is relatively more consistent across the board when compared to OSS-120B. However, differences in prompting methods do not significantly affect consistency.

4.3 Disagreement

Despite having similar accuracy, each prompting strategy impacts what each model eventually responds with.

4.4 Correlation

The correlation plot shows that consistency has a positive correlation with the model’s correctness for both GPT-4o and OSS-120B. However, except in the case of *Explain then Answer*, answer token log probability has near-zero correlation for OSS-120B. We note that OSS-120B is extremely confident in most cases, with




Figure 3: Despite similar performance, each pair of prompting strategies shows some degree of disagreement.




Figure 4: Correlation of each metadata and hallucination rate with respect to each model’s correctness. We group safe and benign extrinsic hallucinations into one group (non-harmful extrinsic hallucination) and compare it to harmful extrinsic hallucination.

an answer token log probability near zero (indicating absolute confidence). Feature distribution visualization is available in Appendix 7, Figure 7.

We define hallucination rate as the average occurrence of each hallucination for a specific question when prompted five times. As previously mentioned, we group safe and benign extrinsic hallucinations into one group, treating them as if no extrinsic hallucination occurred, with harmful extrinsic hallucination being the other case.

5 RESULT: PREDICTING CORRECTNESS

This section addresses the main research question: “Can we predict correctness based on available meta-data?” To answer this question, we use leave-one-out cross validation (LOOCV)² due to our limited number of data points. Table 1 reports the results for all experimental setups.

²https://scikit-learn.org/stable/modules/generated/sklearn.model_selection.LeaveOneOut.html

Table 1: Performance comparison of GPT-4o and OSS-120B across prompting strategies. *Gray rows (Oracle)* denote upper-bound potential, not comparable results. Value inside brackets indicates improvement compared to baseline (assuming that all output made by an LLM is always correct). Under the baseline scenario, recall achieve a perfect score of 1. REG (regression) and CLS (classification) refers to the addition of predicted hallucination signals. Oracle refers to the addition of actual hallucination signals, where *Rate* denotes the average occurrence of each hallucination and *Maj* denotes whether or not each hallucination occurs at least three times when the question is prompted five times.

#	Model	Prompting	Accuracy	Precision	Recall	F1	ROC-AUC
1	GPT-4o	Answer Only	78.57% (+4.08%)	78.89% (+4.40%)	97.26% (-2.74%)	87.12% (+1.74%)	0.573
2	GPT-4o	Answer then Explain	74.49% (-2.04%)	76.60% (+0.07%)	96.00% (-4.00%)	85.21% (-1.50%)	0.484
3	GPT-4o	+ REG	77.55% (+1.02%)	78.49% (+1.96%)	97.33% (-2.67%)	86.90% (+0.19%)	0.763
4	GPT-4o	+ CLS	78.57% (+2.04%)	79.35% (+2.82%)	97.33% (-2.67%)	87.43% (+0.72%)	0.710
5	GPT-4o	+ (Oracle) Rate	85.71% (+9.18%)	89.61% (+13.08%)	92.00% (-8.00%)	90.79% (+4.08%)	0.897
6	GPT-4o	+ (Oracle) Maj	84.69% (+8.16%)	87.50% (+10.97%)	93.33% (-6.67%)	90.32% (+3.61%)	0.868
7	GPT-4o	Explain then Answer	79.59% (+5.10%)	80.46% (+5.97%)	95.89% (-4.11%)	87.50% (+2.12%)	0.822
8	GPT-4o	+ REG	80.61% (+6.12%)	81.40% (+6.91%)	95.89% (-4.11%)	88.05% (+2.67%)	0.815
9	GPT-4o	+ CLS	80.61% (+6.12%)	82.14% (+7.65%)	94.52% (-5.48%)	87.90% (+2.52%)	0.803
10	GPT-4o	+ (Oracle) Rate	78.57% (+4.08%)	83.33% (+8.84%)	89.04% (-10.96%)	86.09% (+0.71%)	0.889
11	GPT-4o	+ (Oracle) Maj	82.65% (+8.16%)	85.90% (+11.41%)	91.78% (-8.22%)	88.74% (+3.36%)	0.890
12	OSS-120B	Answer Only	77.55% (+4.08%)	81.25% (+7.78%)	90.28% (-9.72%)	85.53% (+0.82%)	0.643
13	OSS-120B	Answer then Explain	78.57% (+5.10%)	80.72% (+7.25%)	93.06% (-6.94%)	86.45% (+1.74%)	0.591
14	OSS-120B	+ REG	79.59% (+6.12%)	80.95% (+7.48%)	94.44% (-5.56%)	87.18% (+2.47%)	0.594
15	OSS-120B	+ CLS	79.59% (+6.12%)	80.95% (+7.48%)	94.44% (-5.56%)	87.18% (+2.47%)	0.613
16	OSS-120B	+ (Oracle) Rate	84.69% (+11.22%)	89.04% (+15.57%)	90.28% (-9.72%)	89.66% (+4.95%)	0.929
17	OSS-120B	+ (Oracle) Maj	89.80% (+16.33%)	91.89% (+18.42%)	94.44% (-5.56%)	93.15% (+8.44%)	0.939
18	OSS-120B	Explain then Answer	79.59% (+7.14%)	83.12% (+10.67%)	90.14% (-9.86%)	86.49% (+2.47%)	0.736
19	OSS-120B	+ REG	79.59% (+7.14%)	83.12% (+10.67%)	90.14% (-9.86%)	86.49% (+2.47%)	0.740
20	OSS-120B	+ CLS	79.59% (+7.14%)	83.12% (+10.67%)	90.14% (-9.86%)	86.49% (+2.47%)	0.733
21	OSS-120B	+ (Oracle) Rate	88.78% (+16.33%)	92.86% (+20.41%)	91.55% (-8.45%)	92.20% (+8.18%)	0.890
22	OSS-120B	+ (Oracle) Maj	89.80% (+17.35%)	92.96% (+20.51%)	92.96% (-7.04%)	92.96% (+8.94%)	0.894

5.1 Using Metadata

We initialize this experiment using only two metadata features: consistency and the log probability of the answer token. The results are shown in Table 1, specifically in rows 1, 2, 7, 12, 13, and 18. Although previous analyses indicate that prompting strategy has no statistically significant effect on model accuracy, the table shows that the *Explain then Answer* strategy produces metadata that is more predictive of correctness. For GPT-4o, *Explain then Answer* yields a +5.12% improvement in accuracy over the baseline. For OSS-120B, while *Explain then Answer* provides a +7.14% increase, other prompting strategies perform comparably, suggesting that prompting order has a smaller impact on metadata predictiveness for this model.

5.2 Adding Predicted Hallucination Labels

Obtaining labels for extrinsic and intrinsic hallucinations necessitates the availability of a golden reference, which is typically unattainable in practical scenarios. Consequently, we employ metadata as input to train two predictor models: (1) a classifier that

predicts a binary label (i.e., whether a hallucination occurs) and (2) a regressor that estimates the hallucination rate (i.e., the average proportion of hallucinated content across five attempts). Under the present configuration, the consistency and answer token log probability **are not predictive** of hallucination occurrence. The outcomes of these two models are provided in Appendix 7, Table 2.

Although Table 1 indicates that incorporating predicted hallucination labels marginally improves the correctness predictor model, we attribute this effect to random noise. In most cases, the inclusion of hallucination labels results in either negligible improvements (see rows 7–9 and 13–15) or no improvement at all (see rows 18–20). The apparent performance gains observed in rows 2–4 are likewise interpreted as artifacts of random fluctuations, compounded by the weak baseline performance reported in row 2.

5.3 Oracle Setting

The oracle setting represents the upper-bound performance of the correctness predictor when actual hallucination labels are provided. As shown in the gray rows of Table 1, there remains substantial room for improvement. In this idealized setup, GPT-4o with

the *Answer then Explain* strategy provides the highest predictability of correctness. For OSS-120B, while the potential for improvement remains large, the upper bounds for *Answer then Explain* and *Explain then Answer* are identical.




Figure 5: Feature importance weights of the correctness predictor when provided with metadata and oracle **hallucination majority** signal. **Cons.** = consistency, **LogP** = predicted answer token log probability, **Ext. H.** = extrinsic hallucination, **Int. H.** = intrinsic hallucination, **Log. H.** = logical hallucination.

Figure 5 and 6 showcase the correctness predictor’s feature weights. For GPT-4o, the *Answer then Explain* strategy makes harmful extrinsic and intrinsic hallucinations the strongest predictors of correctness, whereas under *Explain then Answer*, consistency becomes the most important feature.

For OSS-120B, both prompting strategies highlight consistency and hallucination signals as predictive, though the answer token log probability is only useful under *Explain then Answer*. This observation shows that prompting strategy has an existing effect on the predictiveness of correctness through metadata, despite the same upper-bound performance. We hypothesize that the choice of prompting strategy for OSS-120B still affects the model’s internal mechanism, as indicated by the disagreement between strategies, though we leave this question for future work.

Feature Importance by Model and Prompting Strategy Using Oracle Hallucination Rate




Figure 6: Feature importance weights of the correctness predictor when provided with metadata and oracle **hallucination rate** signal, identical abbreviations as Figure 5.

6 DISCUSSION

The use of LLMs is still too risky in high-stakes domains Both GPT-4o and OSS-120B achieve only around 75% and 73% accuracy on the exam dataset, respectively. This is excessively low for high-stakes domains that require absolute precision. Our work shows that both consistency and answer token log probability are useful for flagging potentially erroneous responses, achieving a maximum precision of 83.12%, which is still unreliable for high-stakes applications.

The effect of prompting strategy is still unclear While we find that different prompting strategies have no effect on model performance for this specific task, they still alter the predictive power of metadata features, which we attribute to differences in the model’s internal mechanisms. This is especially surprising for OSS-120B, a reasoning model designed to innately reason. Specifically, we found that *Explain then Answer* may result in different final answers compared to *Answer then Explain*. This is particularly intriguing since the task itself does not require reasoning and is primarily factual.

How different prompting strategies affect LLMs and reasoning LLMs is still understudied

Our results show that even on factual, non-reasoning tasks, prompting strategies can change which features are most predictive of correctness, and can even lead to different final answers for the same model. While GPT-4o shows more pronounced differences between *Answer then Explain* and *Explain then Answer*, OSS-120B exhibits subtler changes in metadata predictiveness despite identical upper-bound performance. These findings suggest that prompting can influence internal decision processes in ways not captured by accuracy alone, highlighting the need for further research into how prompts interact with reasoning models' internal mechanisms across domains and tasks.

Implications for medical education The exam dataset we use reflects real-world medical knowledge assessments, where accuracy is critical. Our results indicate that even strong LLMs like GPT-4o and OSS-120B achieve only 75–79% accuracy, corresponding roughly to a B– to B grade—insufficient for high-stakes decisions. While metadata such as consistency and answer log probability can help flag potentially incorrect responses, the maximum precision of 83% still leaves substantial risk. Even under the oracle setting, accuracy reaches only 89% (B+), highlighting the gap between current LLM capabilities and safe deployment in medical education contexts. This underscores the importance of cautious, supervised use of LLMs as educational aids rather than as unsupervised decision-makers.

7 LIMITATIONS AND FUTURE WORKS

Small dataset The primary limitation of this work lies in the small scope (Prosthodontics) and limited dataset size (n=98), which restrict the generalizability and statistical power of our results. Marginal performance increases should therefore be interpreted with caution. Despite this, the key insight of our work comes from the oracle setting, which demonstrates that correctness is predictable with the help of accurate hallucination detection. We urge future work to explore more sophisticated and less restrictive hallucination detection methods (i.e., those not requiring a reference explanation). We also suggest expanding the dataset and evaluating the correctness predictor across broader domains.

Classification of extrinsic hallucination Our work classifies extrinsic hallucination using a single reference explanation per question. While this simplifies the task, it may misclassify valid alternative explanations as "benign extrinsic hallucinations" when they deviate from the reference. We mitigate this issue by standardizing the annotation into a binary classification: harmful versus non-harmful (i.e., safe or benign) extrinsic hallucinations. This allows us to focus on misleading and factually incorrect explanations.

ACKNOWLEDGEMENTS

This work utilizes large language models to help proof-read the writing of the authors in many sections of this work. The author utilizes Gemini-2.5-Pro to provide codes for the visualizations shown in this work. Other acknowledgements are pending due to double-blind.

REFERENCES

- Abd-alrazaq, A., AlSaad, R., Alhuwail, D., Ahmed, A., Healy, P. M., Latifi, S., Aziz, S., Damsch, R., Alabed Alrazak, S., and Sheikh, J. (2023). Large language models in medical education: Opportunities, challenges, and future directions. *JMIR Med Educ*, 9:e48291.
- Bao, F., Li, M., Luo, R., and Mendelevitch, O. (2024). HHEM-2.1-Open.
- Boizard, N., Gisserot-Boukhlef, H., El-Haddad, K., Hudelot, C., and Colombo, P. (2025). When does reasoning matter? a controlled study of reasoning's contribution to model performance.
- Comanici, G. e. a. (2025). Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities.
- DeepSeek-AI, Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., Zhang, X., Yu, X., Wu, Y., Wu, Z. F., Gou, Z., Shao, Z., Li, Z., Gao, Z., Liu, A., Xue, B., Wang, B., Wu, B., Feng, B., Lu, C., Zhao, C., Deng, C., Zhang, C., Ruan, C., Dai, D., Chen, D., Ji, D., Li, E., Lin, F., Dai, F., Luo, F., Hao, G., Chen, G., Li, G., Zhang, H., Bao, H., Xu, H., Wang, H., Ding, H., Xin, H., Gao, H., Qu, H., Li, H., Guo, J., Li, J., Wang, J., Chen, J., Yuan, J., Qiu, J., Li, J., Cai, J. L., Ni, J., Liang, J., Chen, J., Dong, K., Hu, K., Gao, K., Guan, K., Huang, K., Yu, K., Wang, L., Zhang, L., Zhao, L., Wang, L., Zhang, L., Xu, L., Xia, L., Zhang, M., Zhang, M., Tang, M., Li, M., Wang, M., Li, M., Tian, N., Huang, P., Zhang, P., Wang, Q., Chen, Q., Du, Q., Ge, R., Zhang, R., Pan, R., Wang, R., Chen, R. J., Jin, R. L., Chen, R., Lu, S., Zhou, S., Chen, S., Ye, S., Wang, S., Yu, S., Zhou, S., Pan, S., Li, S. S., Zhou, S., Wu, S., Ye, S.,

- Yun, T., Pei, T., Sun, T., Wang, T., Zeng, W., Zhao, W., Liu, W., Liang, W., Gao, W., Yu, W., Zhang, W., Xiao, W. L., An, W., Liu, X., Wang, X., Chen, X., Nie, X., Cheng, X., Liu, X., Xie, X., Liu, X., Yang, X., Li, X., Su, X., Lin, X., Li, X. Q., Jin, X., Shen, X., Chen, X., Sun, X., Wang, X., Song, X., Zhou, X., Wang, X., Shan, X., Li, Y. K., Wang, Y. Q., Wei, Y. X., Zhang, Y., Xu, Y., Li, Y., Zhao, Y., Sun, Y., Wang, Y., Yu, Y., Zhang, Y., Shi, Y., Xiong, Y., He, Y., Piao, Y., Wang, Y., Tan, Y., Ma, Y., Liu, Y., Guo, Y., Ou, Y., Wang, Y., Gong, Y., Zou, Y., He, Y., Xiong, Y., Luo, Y., You, Y., Liu, Y., Zhou, Y., Zhu, Y. X., Xu, Y., Huang, Y., Li, Y., Zheng, Y., Zhu, Y., Ma, Y., Tang, Y., Zha, Y., Yan, Y., Ren, Z. Z., Ren, Z., Sha, Z., Fu, Z., Xu, Z., Xie, Z., Zhang, Z., Hao, Z., Ma, Z., Yan, Z., Wu, Z., Gu, Z., Zhu, Z., Liu, Z., Li, Z., Xie, Z., Song, Z., Pan, Z., Huang, Z., Xu, Z., Zhang, Z., and Zhang, Z. (2025). Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning.
- Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., and Liu, T. (2025). A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2):1–55.
- Huang, S., Mamidanna, S., Jangam, S., Zhou, Y., and Gilpin, L. H. (2023). Can large language models explain themselves? a study of llm-generated self-explanations.
- Kadavath, S., Conerly, T., Askell, A., Henighan, T., Drain, D., Perez, E., Schiefer, N., Hatfield-Dodds, Z., Das-Sarma, N., Tran-Johnson, E., Johnston, S., El-Showk, S., Jones, A., Elhage, N., Hume, T., Chen, A., Bai, Y., Bowman, S., Fort, S., Ganguli, D., Hernandez, D., Jacobson, J., Kernion, J., Kravec, S., Lovitt, L., Ndousse, K., Olsson, C., Ringer, S., Amodei, D., Brown, T., Clark, J., Joseph, N., Mann, B., McCandlish, S., Olah, C., and Kaplan, J. (2022). Language models (mostly) know what they know.
- Krishna, S., Ma, J., Slack, D., Ghandeharioun, A., Singh, S., and Lakkaraju, H. (2023). Post hoc explanations of language models can improve language models.
- Li, D., Jiang, B., Huang, L., Beigi, A., Zhao, C., Tan, Z., Bhattacharjee, A., Jiang, Y., Chen, C., Wu, T., Shu, K., Cheng, L., and Liu, H. (2025). From generation to judgment: Opportunities and challenges of llm-as-a-judge.
- Manakul, P., Liusie, A., and Gales, M. J. F. (2023). Self-checkgpt: Zero-resource black-box hallucination detection for generative large language models.
- OpenAI, :, Agarwal, S., Ahmad, L., Ai, J., Altman, S., Applebaum, A., Arbus, E., Arora, R. K., Bai, Y., Baker, B., Bao, H., Barak, B., Bennett, A., Bertao, T., Brett, N., Brevdo, E., Brockman, G., Bubeck, S., Chang, C., Chen, K., Chen, M., Cheung, E., Clark, A., Cook, D., Dukhan, M., Dvorak, C., Fives, K., Fomenko, V., Garipov, T., Georgiev, K., Glaese, M., Gogineni, T., Goucher, A., Gross, L., Guzman, K. G., Hallman, J., Hehir, J., Heidecke, J., Helyar, A., Hu, H., Huet, R., Huh, J., Jain, S., Johnson, Z., Koch, C., Kofman, I., Kundel, D., Kwon, J., Kyrylov, V., Le, E. Y., Leclerc, G., Lennon, J. P., Lessans, S., Lezcano-Casado, M., Li, Y., Li, Z., Lin, J., Liss, J., Lily, Liu, Liu, J., Lu, K., Lu, C., Martinovic, Z., McCallum, L., McGrath, J., McKinney, S., McLaughlin, A., Mei, S., Mostovoy, S., Mu, T., Myles, G., Neitz, A., Nichol, A., Pachocki, J., Paino, A., Palmie, D., Pantuliano, A., Parascandolo, G., Park, J., Pathak, L., Paz, C., Peran, L., Pimenov, D., Pokrass, M., Proehl, E., Qiu, H., Raila, G., Raso, F., Ren, H., Richardson, K., Robinson, D., Rotsted, B., Salman, H., Sanjeev, S., Schwarzer, M., Sculley, D., Sikchi, H., Simon, K., Singhal, K., Song, Y., Stuckey, D., Sun, Z., Tillet, P., Toizer, S., Tsimpouras, F., Vyas, N., Wallace, E., Wang, X., Wang, M., Watkins, O., Weil, K., Wendling, A., Whinnery, K., Whitney, C., Wong, H., Wang, L., Yang, Y., Yasunaga, M., Ying, K., Zaremba, W., Zhan, W., Zhang, C., Zhang, B., Zhang, E., and Zhao, S. (2025). gpt-oss-120b gpt-oss-20b model card.
- OpenAI, :, Hurst, A., Lerer, A., Goucher, A. P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., Madry, A., Baker-Whitcomb, A., Beutel, A., Borzunov, A., Carney, A., Chow, A., Kirillov, A., Nichol, A., Paino, A., Renzin, A., Passos, A. T., Kirillov, A., Christakis, A., Conneau, A., Kamali, A., Jabri, A., Moyer, A., Tam, A., Crookes, A., Tootoochian, A., Tootoonchian, A., Kumar, A., Vallone, A., Karpathy, A., Braunstein, A., Cann, A., Codispoti, A., Galu, A., Kondrich, A., Tulloch, A., Mishchenko, A., Baek, A., Jiang, A., Pelisse, A., Woodford, A., Gosalia, A., Dhar, A., Pantuliano, A., Nayak, A., Oliver, A., Zoph, B., Ghorbani, B., Leimberger, B., Rossen, B., Sokolowsky, B., Wang, B., Zweig, B., Hoover, B., Samic, B., McGrew, B., Spero, B., Gierler, B., Cheng, B., Lightcap, B., Walkin, B., Quinn, B., Guarraci, B., Hsu, B., Kellogg, B., Eastman, B., Lugaresi, C., Wainwright, C., Bassin, C., Hudson, C., Chu, C., Nelson, C., Li, C., Shern, C. J., Conger, C., Barette, C., Voss, C., Ding, C., Lu, C., Zhang, C., Beaumont, C., Hallacy, C., Koch, C., Gibson, C., Kim, C., Choi, C., McLeavey, C., Hesse, C., Fischer, C., Winter, C., Czarnecki, C., Jarvis, C., Wei, C., Koumouzelis, C., Sherburn, D., Kappler, D., Levin, D., Levy, D., Carr, D., Farhi, D., Mely, D., Robinson, D., Sasaki, D., Jin, D., Valladares, D., Tsipras, D., Li, D., Nguyen, D. P., Findlay, D., Oiwoh, E., Wong, E., Asdar, E., Proehl, E., Yang, E., Antonow, E., Kramer, E., Peterson, E., Sigler, E., Wallace, E., Brevdo, E., Mays, E., Khorasani, F., Such, F. P., Raso, F., Zhang, F., von Lohmann, F., Sulit, F., Goh, G., Oden, G., Salmon, G., Starace, G., Brockman, G., Salman, H., Bao, H., Hu, H., Wong, H., Wang, H., Schmidt, H., Whitney, H., Jun, H., Kirchner, H., de Oliveira Pinto, H. P., Ren, H., Chang, H., Chung, H. W., Kivlichan, I., O’Connell, I., O’Connell, I., Osband, I., Silber, I., Sohl, I., Okuyucu, I., Lan, I., Kostrikov, I., Sutskever, I., Kanitscheider, I., Gulrajani, I., Coxon, J., Menick, J., Pachocki, J., Aung, J., Betker, J., Crooks, J., Lennon, J., Kiro, J., Leike, J., Park, J., Kwon, J., Phang, J., Teplitz, J., Wei, J., Wolfe, J., Chen, J., Harris, J., Varavva, J., Lee, J. G., Shieh, J., Lin, J., Yu, J., Weng, J., Tang, J., Yu, J., Jang, J., Candela, J. Q., Beutler, J., Landers, J., Parish, J.,

Heidecke, J., Schulman, J., Lachman, J., McKay, J., Uesato, J., Ward, J., Kim, J. W., Huizinga, J., Sitkin, J., Kraaijeveld, J., Gross, J., Kaplan, J., Snyder, J., Achiam, J., Jiao, J., Lee, J., Zhuang, J., Harriman, J., Fricke, K., Hayashi, K., Singhal, K., Shi, K., Karthik, K., Wood, K., Rimbach, K., Hsu, K., Nguyen, K., Gu-Lemberg, K., Button, K., Liu, K., Howe, K., Muthukumar, K., Luther, K., Ahmad, L., Kai, L., Itow, L., Workman, L., Pathak, L., Chen, L., Jing, L., Guy, L., Fedus, L., Zhou, L., Mamitsuka, L., Weng, L., McCallum, L., Held, L., Ouyang, L., Feuvrier, L., Zhang, L., Kondraciuk, L., Kaiser, L., Hewitt, L., Metz, L., Doshi, L., Afak, M., Simens, M., Boyd, M., Thompson, M., Dukhan, M., Chen, M., Gray, M., Hudnall, M., Zhang, M., Aljube, M., Litwin, M., Zeng, M., Johnson, M., Shetty, M., Gupta, M., Shah, M., Yatbaz, M., Yang, M. J., Zhong, M., Glaese, M., Chen, M., Janner, M., Lampe, M., Petrov, M., Wu, M., Wang, M., Fradin, M., Pokrass, M., Castro, M., de Castro, M. O. T., Pavlov, M., Brundage, M., Wang, M., Khan, M., Murati, M., Bavarian, M., Lin, M., Yesildal, M., Soto, N., Gimbelshein, N., Cone, N., Staudacher, N., Summers, N., LaFontaine, N., Chowdhury, N., Ryder, N., Stathas, N., Turley, N., Tezak, N., Felix, N., Kudige, N., Keskar, N., Deutsch, N., Bundick, N., Puckett, N., Nachum, O., Okelola, O., Boiko, O., Murk, O., Jaffe, O., Watkins, O., Godement, O., Campbell-Moore, O., Chao, P., McMillan, P., Belov, P., Su, P., Bak, P., Bakkum, P., Deng, P., Dolan, P., Hoeschele, P., Welinder, P., Tillet, P., Pronin, P., Tillet, P., Dhariwal, P., Yuan, Q., Dias, R., Lim, R., Arora, R., Troll, R., Lin, R., Lopes, R. G., Puri, R., Miyara, R., Leike, R., Gaubert, R., Zamani, R., Wang, R., Donnelly, R., Honsby, R., Smith, R., Sahai, R., Ramchandani, R., Huet, R., Carmichael, R., Zellers, R., Chen, R., Chen, R., Nigmatullin, R., Cheu, R., Jain, S., Altman, S., Schoenholz, S., Toizer, S., Miserendino, S., Agarwal, S., Culver, S., Ethersmith, S., Gray, S., Grove, S., Metzger, S., Hermani, S., Jain, S., Zhao, S., Wu, S., Jomoto, S., Wu, S., Shuaiqi, Xia, Phene, S., Papay, S., Narayanan, S., Coffey, S., Lee, S., Hall, S., Balaji, S., Broda, T., Stramer, T., Xu, T., Gogineni, T., Christianson, T., Sanders, T., Patwardhan, T., Cunningham, T., Degry, T., Dimson, T., Raoux, T., Shadwell, T., Zheng, T., Underwood, T., Markov, T., Sherbakov, T., Rubin, T., Stasi, T., Kaftan, T., Heywood, T., Peterson, T., Walters, T., Eloundou, T., Qi, V., Moeller, V., Monaco, V., Kuo, V., Fomenko, V., Chang, W., Zheng, W., Zhou, W., Manassra, W., Sheu, W., Zaremba, W., Patil, Y., Qian, Y., Kim, Y., Cheng, Y., Zhang, Y., He, Y., Zhang, Y., Jin, Y., Dai, Y., and Malkov, Y. (2024). Gpt-4o system card.

- Turpin, M., Michael, J., Perez, E., and Bowman, S. R. (2023). Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., and Zhou, D. (2023). Chain-of-thought prompting elicits reasoning in large language models.
- Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., Zheng, C., Liu, D.,

Zhou, F., Huang, F., Hu, F., Ge, H., Wei, H., Lin, H., Tang, J., Yang, J., Tu, J., Zhang, J., Yang, J., Yang, J., Zhou, J., Zhou, J., Lin, J., Dang, K., Bao, K., Yang, K., Yu, L., Deng, L., Li, M., Xue, M., Li, M., Zhang, P., Wang, P., Zhu, Q., Men, R., Gao, R., Liu, S., Luo, S., Li, T., Tang, T., Yin, W., Ren, X., Wang, X., Zhang, X., Ren, X., Fan, Y., Su, Y., Zhang, Y., Zhang, Y., Wan, Y., Liu, Y., Wang, Z., Cui, Z., Zhang, Z., Zhou, Z., and Qiu, Z. (2025). Qwen3 technical report.

Zhao, C., Tan, Z., Ma, P., Li, D., Jiang, B., Wang, Y., Yang, Y., and Liu, H. (2025). Is chain-of-thought reasoning of llms a mirage? a data distribution lens.

APPENDIX

Hallucination Predictor Model Performance

This appendix shows the performance of the hallucination predictor model, which is available in Table 2. The metrics used for the regressor are the following:

1. **R²**: The coefficient of the determinant. A value of 0 indicates that the model explains **none** of the variance in the target variable (the **Target** column)
2. **MAE**: Mean absolute error. A value of 0 indicates perfect model performance.
3. **RMSE**: Root mean squared error. A value of 0 indicates perfect model performance.

Hallucination Extraction Prompt

We use the following prompt to extract the judgement from Gemini-2.5-flash.

```
You are a meticulous, expert evaluator
in the field of prosthodontics.
Your task is to perform a three-
part evaluation of the "Candidate
Explanation" using the "Ground
Truth Context" and the detailed "
Evaluation Rubric" below.

**Evaluation Rubric:**

1.  **Intrinsic Hallucination:** Does
the Candidate Explanation directly
CONTRADICT a fact in the Reference
Explanation?
- "Yes" or "No".

2.  **Extrinsic Hallucination:** Does
the Candidate Explanation FABRICATE
new information not present in the
Reference Explanation? Classify
any fabrication based on its **harm
```

Table 2: Performance of hallucination predictors using metadata. Left: classification of hallucination class. Right: regression of hallucination rate. Note that both the classifier and the regressor perform close to the baseline, indicating that metadata is not usable to predict hallucination in our setup.

Classifier							Regressor		
Model	Prompting	Target	Baseline Acc	Accuracy	Precision	F1	R ²	MAE	RMSE
GPT-4o	A→E	extrinsic	0.81	0.81	0.00	0.00	-0.01	0.282	0.361
GPT-4o	A→E	intrinsic	0.70	0.69	0.00	0.00	-0.03	0.342	0.392
GPT-4o	A→E	logical	0.98	0.98	0.00	0.00	-0.02	0.051	0.117
GPT-4o	E→A	extrinsic	0.83	0.84	0.60	0.27	0.07	0.210	0.304
GPT-4o	E→A	intrinsic	0.74	0.72	0.00	0.00	0.01	0.299	0.364
GPT-4o	E→A	logical	0.90	0.90	0.00	0.00	-0.00	0.178	0.236
OSS-120B	A→E	extrinsic	0.78	0.78	0.00	0.00	0.05	0.295	0.367
OSS-120B	A→E	intrinsic	0.64	0.65	0.53	0.32	0.01	0.342	0.388
OSS-120B	A→E	logical	0.98	0.98	0.00	0.00	0.05	0.071	0.115
OSS-120B	E→A	extrinsic	0.69	0.77	0.67	0.55	0.16	0.274	0.353
OSS-120B	E→A	intrinsic	0.54	0.67	0.72	0.57	0.02	0.352	0.399
OSS-120B	E→A	logical	0.98	0.98	0.00	0.00	-0.05	0.080	0.124

```

to the explanation's core argument
**:
- **"Harmful":** The fabricated
  information is an **essential
  detail** to support the final
  answer. It introduces a new
  cause or reason that is
  essential to the logic,
  deviating from the core premise
  of the ground truth.
- **"Benign":** The fabricated
  information is a **superfluous
  detail**. It does not deviate
  from the core premise of the
  ground truth.
- **"None":** The explanation
  contains no new information.

3. **Logical Hallucination:** Does the
  reasoning in the Candidate
  Explanation logically support the
  Candidate's Final Answer?
  - "Aligned" or "Misaligned".

**Your Task:**
Provide your evaluation ONLY in the
following JSON format:
{{
  "intrinsic_hallucination": "Yes or No
  ",
  "extrinsic_hallucination": "None,
  Benign, or Harmful",
  "logical_rating": "Aligned or
  Misaligned",
  "reasoning": "An explanation for your
  evaluation, referencing the
  rubric."
}}
---
**Ground Truth Context:**

```

```

**Question:** {question_text}
**Reference Explanation:**
{ground_truth_explanation}
---
**Candidate to Evaluate:**
**Candidate's Final Answer:** {
  prediction_answer}
**Candidate Explanation:**
{prediction_explanation}
---
**JSON Output:**

```

Example of Extrinsic Hallucinations

Example #1: Benign Extrinsic Hallucination

Question: Which of the following statements accurately describes Pound's Lingualized Occlusion for denture teeth arrangement?

- The lingual cusps of the mandibular posterior teeth function in the central fossae of the maxillary posterior teeth.
- The lingual cusps of the maxillary posterior teeth function in the central fossae of the mandibular posterior teeth. (Correct answer and model answer)
- (Truncated)
- (Truncated)

Reference Explanation: The maxillary posterior teeth are set to aid in developing a lingualized occlusion, wherein the lingual cusps of the maxillary posterior teeth function in the central fossae of the

<p>mandibular posterior teeth. This centralizing effect can be achieved by placing the upper posterior teeth in a manner that allows their lingual cusps to engage with the central fossae of the lower posterior teeth, thereby enhancing stability and occlusal function.</p> <p>Model Explanation: Pound’s Lingualized Occlusion involves the lingual cusps of the maxillary posterior teeth functioning in the central fossae of the mandibular posterior teeth to [create a more balanced and efficient occlusal contact, enhancing the stability and function of the dentures.]</p> <p>Gemini Reasoning: (Truncated)... [The phrase ‘to create a more balanced and efficient occlusal contact’ is an additional detail not explicitly in the reference, but it is consistent with the stated purpose of ‘enhancing stability and occlusal function’ and does not introduce a new, essential argument or contradict the core premise, classifying it as a benign extrinsic hallucination.]...</p> <p>=====</p> <p>Example #2: Harmful Extrinsic Hallucination</p> <p>Question: In a comparison of two surface treatments for highly translucent 5Y-TZP zirconia: 1. Non-thermal oxygen plasma (NTP) treatment shows significantly higher bond strength of zirconia to composite resin compared to airborne particle abrasion (APA) treatment. 2. NTP treatment shows significantly lower biaxial flexural strength of 5Y-TSP zirconia compared to APA.</p> <p>a. Both statements are true. b. The first statement is true, but the second statement is false. (Model Prediction) c. Both statements are false. (Correct Answer) d. The first statement is false, but the second statement is true.</p> <p>Reference Explanation: Based on the obtained results, the null hypotheses were rejected, since there was a significant difference among different surface treatments. Regarding the difference in DTP00,</p>	<p>the APA group showed a significantly higher difference in the DTP00 than other groups, followed by the TS, and the lowest value was for the NTP. This difference could be explained by the fact that airborne-particle abrasion resulted in increased surface irregularities causing diffuse light reflection and an increase in light scattering from the tested specimens. However, surface treatment with the NTP caused surface activation with minimal alteration in its surface topography, consistent with previous studies. Nevertheless, the NTP treatment showed significantly lower bond strength values; however, it improved strength.</p> <p>Model Explanation: Non-thermal oxygen plasma treatment [enhances the surface energy and wettability] of highly translucent 5Y-TZP zirconia, resulting in a significantly higher bond strength to composite resin compared with conventional airborne particle abrasion. (Truncated) ...</p> <p>Gemini Reasoning: (Truncated) ... [Extrinsic hallucination is present and classified as ‘Harmful.’] The Candidate fabricates essential details not found in the Ground Truth, such as ‘enhances the surface energy and wettability’ and ‘plasma does not mechanically roughen or induce surface flaws, it does not compromise the material’s bulk integrity.’ These fabricated details are crucial to supporting the Candidate’s incorrect claim about higher bond strength (which intrinsically hallucinates against the ground truth) and provide the logical basis for its overall argument, thereby deviating from the core premise of the ground truth ...</p>
<p>Feature Distribution</p> <p>Feature distribution of metadata and hallucination rate with respect to model’s correctness, as shown in Figure 7.</p>	

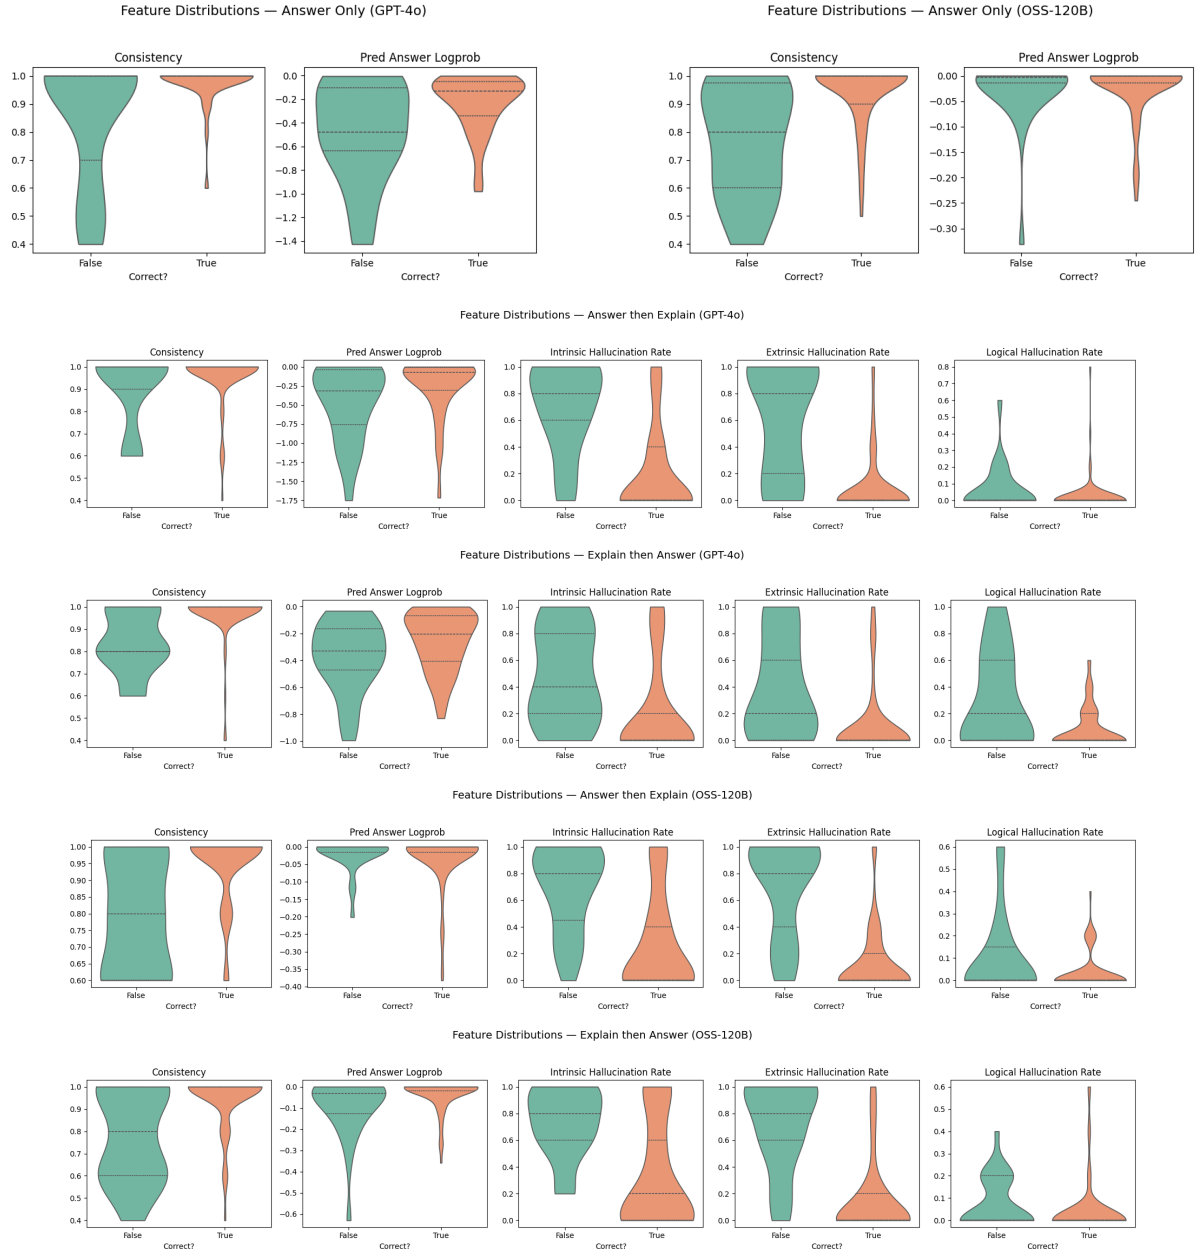


Figure 7: Distribution of metadata and hallucination features with respect to model’s correctness.