Breaking the Chain: A Causal Analysis of LLM Faithfulness to Intermediate Structures

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models (LLMs) increasingly generate intermediate reasoning structures — rubrics, checklists, proof graphs — to make their decisions more interpretable. But are these structures causal mediators of the final answer, or decorative by-products? We introduce a causal evaluation protocol that tests LLM faithfulness via interventions to original prompt or corresponding intermediate structures. Across nine models and four benchmarks with annotated intermediates, the protocol reveals a systematic gap: while models rely on structures more than the original text (> 60% consistency under interventions to original prompt), they fail to update under logically significant structural edits more than 50% of the time. Surprisingly, models are more faithful to their self-generated structures than to gold ones, suggesting that the act of generation elicits reasoning more effectively than passive consumption. Our study provides the causal and systematic evidence that current LLMs treat intermediate structures as context rather than true mediators of decision making.

1 Introduction

What happens if we edit a model's reasoning steps? If an LLM is faithful to its intermediate structures — rubrics, checklists, proof graphs — then logically significant edits should change its final decision. If not, interventions lose their effect, undermining interpretability and human—AI collaboration. By human—AI collaboration we mean workflows where a human monitors, corrects, or guides an AI system: a teacher adjusting a rubric item in automated grading, a fact-checker flipping an evidence tag in claim verification, or a scientist editing a proof graph in multi-step reasoning. In each case, the human expects that the correction propagates to the final outcome. If the model ignores such edits, the collaboration becomes illusory — humans can view the reasoning, but cannot steer it.

Intermediate structures are crucial for enabling large language models (LLMs) to handle multi-step tasks, such as verifying factual claims, solving scientific questions, or supporting decision making. They provide an explicit reasoning pathway that humans can inspect and edit, with the expectation that such changes will be faithfully reflected in the final decision of the model (Sonkar et al., 2024; Cabrera et al.; Bussone et al., 2015).

Faithfulness means that an explanation or intermediate output reflects the true decision process (Jacovi & Goldberg, 2020). Prior work shows that LLM-generated reasoning is often unfaithful (Turpin et al., 2023; Lanham et al., 2023; Paul et al., 2024; Matton et al., 2025), raising doubts about when such intermediates can be trusted. Since model reasoning, or chain-of-thought, can be used for performance monitoring (Baker et al., 2025), faithfulness becomes a crucial issue.

In our work, we study the faithfulness of LLM to interventions on intermediate structures. Sometimes, intermediate structures can be used as input to deterministic systems such as verifiers or interpreters (Kirchner et al., 2024; Wang et al., 2024). But even if deterministic systems could replace mediators in narrowly defined tasks, the central research question is whether LLMs can be reliably steered by structured intermediates at all or are they ignored in favor of hidden shortcuts?

If models cannot be influenced by such structures, their reasoning remains a black box, limiting interpretability and undermining human-AI collaboration (Figure 1). If they can, then structured mediators provide actionable handles through which humans or external systems can intervene, cor-

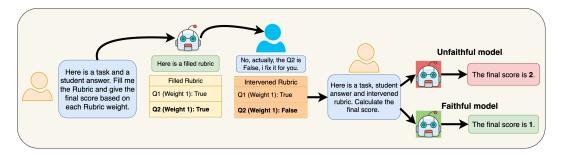


Figure 1: Illustration of causal intervention on intermediate structures. A model receives a task and produces a filled rubric as an explicit intermediate reasoning step, which is then used to generate a final score. By intervening on the rubric (e.g., correcting Q2 from True to False), we can test whether the final prediction is causally mediated by this intermediate structure. If the score changes consistently with the rubric edit, the model is faithful to the mediation; if the score ignores the correction, the model is unfaithful and relies on hidden shortcuts.

rect, and guide model behavior—an essential property for deploying LLMs in real-world decision-making workflows.

We introduce a causal evaluation protocol for testing whether LLMs faithfully rely on intermediate reasoning structures when making final decisions (Section 2.3). The protocol applies two kinds of interventions: (i) altering the intermediate structure while holding the input fixed, to test whether decisions causally depend on the structure, and (ii) varying the input while holding the structure fixed, to test whether models rely on the stated structure rather than surface cues. Faithfulness is then measured with three indicators — Hold Structure, Vary Text (HSVT), Local Edit Consistency (LEC), and Global Edit Consistency (GEC) (Section 2.2). Applying this framework to nine modern LLMs across four benchmarks with gold intermediates (rubrics, checklists, proof graphs, and verification queries), we find that while models rely on structured intermediates more than surface text (HSVT > 60%), they are not causally mediated by them (LEC/GEC = 40–60%) (Sections 5.1 and 5.2). Models are paradoxically more faithful to their own generated structures than to ground truth, global structural breaks are more disruptive than local edits, and soft input variations preserve faithfulness better than hard ones.

2 PROTOCOL FOR FAITHFULNESS EVALUATION OVER INTERMEDIATE STRUCTURES

2.1 PROBLEM FORMULATION

We consider a setting where an LLM receives an input X (e.g., a science question-answer, claim or hypothesis with supporting facts) and produces two outputs: an intermediate reasoning representation M (e.g., a checklist, rubric, or structured proof) and a final decision Y (e.g., sentiment label, correctness judgment, or entailment decision) which is based on M. Let U denote unobserved factors, such as the model's latent internal reasoning, that may influence both M and Y.

Because LLMs decode auto-regressively, we view this process as a two-stage generation: first M is produced from $p_{\theta}(M \mid X)$, and then Y is produced from

$$p_{\theta}(Y \mid X, M) = \prod_{t=1}^{|Y|} p_{\theta}(y_t \mid X, M, y_{< t}), \tag{1}$$

where y_t denotes the t-th token of the final decision. In the *faithful* case (Figure 2a), Y is functionally determined by M (given X), so that task-relevant edits to M alter the final decision. In the *unfaithful* case (Figure 2b), the model conditions on M structurally but relies semantically primarily on X or latent knowledge U; thus interventions on M fail to change Y. Here M may appear coherent while exerting little or no causal influence on the decision.

This corresponds to a front-door structure (Pearl, 2001), where the visible reasoning M is the sole conduit through which X affects Y during auto-regressive decoding. In the *unfaithful* case (Fig-

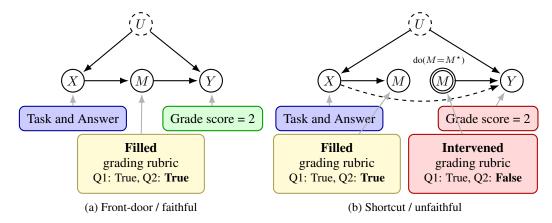


Figure 2: Causal framing of intervention on intermediate structure. (a) **Faithful mediation:** The input X influences the final answer Y primarily through the explicit intermediate structure M, with a possible confounder U. Edits to M therefore change Y, consistent with M being the causal mediator. (b) **Shortcut / unfaithful:** Although Y is formally conditioned on both X and M, the model may rely mainly on X or latent knowledge U, rendering M causally irrelevant. A *double-outlined node* indicates a variable set by intervention via the do() operator (here, M is replaced by M^*). The intervention tests whether altering M changes Y: if Y remains unchanged, the model is effectively ignoring the mediated pathway and defaulting to direct reliance on X or U.

ure 2b), a direct edge $X \to Y$ bypasses M, allowing the model to ignore the stated reasoning and rely on latent knowledge from U. In such cases, M may appear coherent while having little or no causal influence on Y.

We operationalize the faithfulness test via *interventions* on M (given X). Given a generated M, we construct an intervened version M^* using a dataset-specific transformation $\mathcal{I}(M)$ (e.g., flipping checklist items, altering rubric entries, or rewiring proof edges) that changes reasoning content without modifying X. The model is then re-prompted with M^* to produce Y^* . For a faithful model it is necessary that

$$M^* \neq M \Rightarrow Y^* \neq Y,$$
 (2)

under interventions that logically change the correct decision. Failure to update Y indicates that the model is not using M as a true mediator but instead defaulting to shortcuts from X or U e.g. from internal states.

2.2 FAITHFULNESS EVALUATION METRICS

We propose evaluation methodology from two complementary angles. First, we test whether model decisions genuinely depend on the surface properties of the input text. Second, we ask whether decisions are causally *mediated* by the structure itself. This decomposition separates two distinct failure modes: reliance on cues from original text versus ignoring the mediator entirely.

Hold Structure, Vary Text (HSVT). We hold the intermediate structure M fixed and perturb the input X with meaning-preserving edits (e.g., paraphrases, renamings, or style shifts) or totally replacing input (e.g., replacing with different text, the intermediate structure was originally based on). A faithful model should produce \hat{Y}^* as the predicted decision \hat{Y} before intervention. Formally,

$$HSVT = \frac{1}{N} \sum_{i=1}^{N} \mathbf{1} \left[\hat{Y}_{i}^{*} = \hat{Y}_{i} \middle| M \text{ fixed, } X \text{changed} \right], \tag{3}$$

where N is the number of samples. High HSVT indicates reliance on M rather than cues in X.

Local Edit Consistency (LEC). We flip individual elements of M (e.g., checklist entries, proof edges) while keeping X fixed. A faithful model should update its decision \hat{Y}_j^* after the edit j to

Algorithm 1 Inference with mediator intervention. In the Gold setting, we provide M_{gold} instead of generating \hat{M} .

Require: Dataset $\mathcal{D} = \{(x_i, M_i, y_i)\}_{i=1}^N$, model p_{θ}

- 1: **for** each $(x, M_{gold}, y) \in \mathcal{D}$ **do**
 - 2: Construct prompt from x to predict mediator \hat{M} and decision \hat{Y}
 - 3: Query p_{θ} to generate a completion
 - 4: Parse completion into (\hat{M}, \hat{Y})
 - 5: Apply intervention $\mathcal{I}(\hat{M}) \mapsto M^*$
 - 6: Compute counterfactual target $C(M^*) \mapsto \tilde{Y}$
 - 7: Form prompt (x, M^*) and query p_{θ} for decision \hat{Y}^*
 - 8: Evaluate HSVT with $(\hat{Y}^{\star}, \hat{Y})$ and LEC\GEC with $(\hat{Y}^{\star}, \tilde{Y})$
 - 9: end for

match expected counterfactual target Y. For each sample with K edits we calculate average over all edits LEC_{sample} and the calculate the dataset score LEC:

$$LEC_{\text{sample}} = \frac{1}{K} \sum_{i=1}^{K} \mathbf{1} \left[\hat{Y}_{j}^{*} = \tilde{Y} \right], LEC = \frac{1}{N} \sum_{i=1}^{N} LEC_{\text{sample},i}.$$
 (4)

Global Edit Consistency (GEC). We apply a *global break* to M that destroys all valid support. $\hat{Y}^*_{(gb)}$ is the expected prediction after intervention. The decision should flip as in expected counterfactual target \tilde{Y} :

$$GEC_{\text{sample}} = \mathbf{1} \left[\hat{Y}_{(gb)}^* = \tilde{Y} \right], \qquad GEC = \frac{1}{N} \sum_{i=1}^{N} GEC_{\text{sample},i}.$$
 (5)

Evaluation Settings. We distinguish two settings: (i) *Predicted*, where the model generates M and then predicts \hat{Y} , and (ii) *Gold*, where the ground-truth structure M_{gold} is provided and used to predict \hat{Y} . Applying interventions in both settings reveals whether failures stem from incorrect generation of \hat{M} or from ignoring even correct, externally supplied structures. This also helps us to understand the role of intermediate structures in idealized (gold) and real (predicted) setting.

2.3 IMPLEMENTATION PROTOCOL

We run two settings per example. In the *predicted* regime, the prompt contains only X (task + evidence + an empty structured template), and the model must first generate \hat{M} (e.g., a filled checklist) and then \hat{Y} . In the *gold* regime, we append the ground-truth mediator $M_{\rm gold}$ to the dialogue and ask the model to output only \hat{Y} . This isolates whether the model can use a provided structure without re-deriving the decision from X.

After the initial run, we construct M^* by applying dataset-specific edits to M while keeping X fixed: (i) **HSVT**: replace or paraphrase X (hard: swap in another answer or table row; soft: paraphrase) while holding M fixed; (ii) **Local**: flip individual slots or edges in M; (iii) **Global**: apply a full break (e.g., invert all rubric entries or cut all proof edges). For each intervention, we also compute the expected counterfactual target \tilde{Y} based on the altered mediator M^* . We then re-prompt the model with (X, M^*) and obtain \hat{Y}^* , comparing it against \hat{Y} to evaluate HSVT or against \tilde{Y} to evaluate LEC\GEC. The protocol is summarized in Algorithm 1.

Importantly, our prompt design never explicitly instructs the model to maintain consistency between M and Y. Instead, we describe the task objectives and provide few-shot examples without interventions. This ensures our evaluation measures the model's intrinsic reliance on the mediator rather than its ability to follow an instruction to "be faithful." Prompt examples are given in Appendix A.1.

To avoid stochastic variability, we disable sampling and decode deterministically (temperature = 0, greedy). We enforce a minimal output schema (e.g., "Final grade: <float>" or a categorical label) and extract \hat{Y} (and \hat{M} in the predicted regime) with simple pattern matching.

3 Datasets and Interventions

We evaluate our framework on four datasets that instantiate the general protocol across different reasoning formats and decision types. In all cases, the model first produces an intermediate structure M, which we intervene on before re-querying the model for the final decision Y. This allows us to test whether decisions are invariant to surface changes in X (HSVT) and whether they are causally mediated by M (local/global edits).

RiceChem (Rubric Intervention) (Sonkar et al., 2024). A chemistry grading task where each sample contains a question, a student answer, and a rubric of True/False sub-steps with real-valued weights. The dataset comprised on 4 tasks, with each task having 5 to 8 checklist items. The model fills M (the rubric), sums the weights of satisfied items, and outputs a score. For mediation, we flip rubric entries, which directly changes the implied score; a faithful model should update its prediction accordingly rather than re-deriving the score from latent knowledge. For HSVT, we replace the student answer text X with another answer to the same question while keeping the rubric fixed. A faithful model should output the same score, since M already determines it.

Averitec (QA–Rubric Intervention) (Schlichtkrull et al., 2023). A fact-checking dataset where each claim is paired with binary support questions. Each sample has has from 1 to 3 support questions. We treat the answers to these sub-questions as M, which mediates the final verdict. Interventions flip rubric entries (e.g., changing a sub-answer from "Yes" to "No"), requiring the model to update the claim label accordingly. For HSVT, we paraphrase or stylistically alter the claim text X while keeping the rubric unchanged. Faithful predictions should remain stable under these surface variations.

EntailmentBank (Proof Intervention) (Dalvi et al., 2021). A multi-hop entailment task where the model produces a structured proof M linking premises to a hypothesis, and outputs a final binary decision whether the hypothesis is correct or not. We intervene by corrupting the proof (removing or rewiring edges) while keeping textual content identical. Since all interventions yield invalid proofs, a faithful model should flip its entailment decision. For HSVT, we paraphrase premises or rename entities in X while leaving the proof fixed. Faithful models should preserve the entailment decision.

TabFact (Verification Query Intervention) (Wenhu Chen & Wang, 2020). A table-based fact verification dataset where each statement is grounded in a structured query M over table cells and operators. We intervene by altering the query (e.g., swapping columns or operators) while keeping the statement text fixed, which should flip the entailment label if the model is faithful. For HSVT, we substitute the original statement text X with a different statement from the dataset that corresponds to the same table while holding the query constant. Faithful models should maintain the same decision.

Counterfactual Targets. For each dataset, we define a counterfactual target Y that reflects the expected decision after intervention on M: (i) recomputed scores from edited rubrics in RiceChem, (ii) flipped verdicts from inverted sub-answers in Averitec, (iii) not correct labels from broken proofs in EntailmentBank, (iv) refuted labels from corrupted queries in TabFact. Evaluating against these targets allows us to measure whether models' predictions are causally mediated by the structures they generate or consume.

4 RELATED WORK

Intermediate structures. Intermediate structures are widely used in decision-making systems to decompose complex tasks, support monitoring, and enable verification. They appear in domains such as education (rubric-based essay grading (Sonkar et al., 2024)), semantic parsing tasks (like

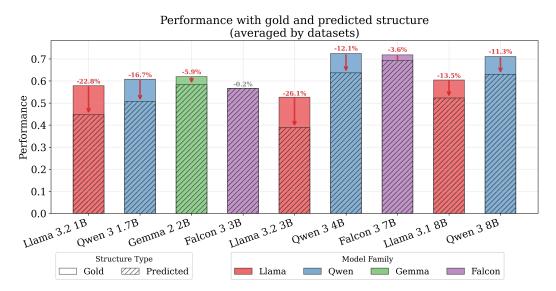


Figure 3: With gold structure, the model is provided with ground-truth intermediate structure M. With predicted structure, the model has to generate intermediate structure M from scratch. We report metrics averaged over 4 datasets considered in the paper.

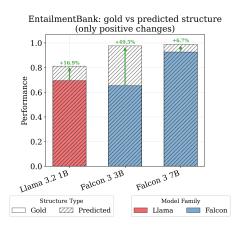
decomposition of the task into consecutive elements)(Lee et al., 2022; Yu et al.), contract review (clause extraction) (Amazou et al., 2025), and clinical decision support, where transparency improves professional trust (Bussone et al., 2015). In NLP and AI planning, intermediate structures take the form of schema linking in text-to-query (Yu et al.; Li et al., 2023) or formal action sequences in PDDL (Silver et al., 2022), both of which mediate between input and final execution. In this work, we evaluate such settings through the lens of faithfulness: whether LLMs align their decisions with structured intermediates and respond appropriately to interventions.

Reasoning faithfulness. Recent work has examined whether model-generated rationales faithfully mediate predictions. (Turpin et al., 2023; Chua & Evans, 2025; Chen et al., 2024; Korbak et al., 2025) test reasoning faithfulness by truncating, paraphrasing, or injecting mistakes, finding task-dependent variation and showing that rationales are often helpful but not reliably causal. Also the models verbalization was inspected of the use of injected hints (e.g., sycophancy, grader hacking, metadata) and report very low "reveal rates", concluding that test-time reasoning monitoring cannot be relied upon to expose shortcut use (Lyu et al., 2023; Lanham et al., 2023). While faithfulness and causality are closely related, most studies do not treat the rationale as an explicit causal mediator, however it is an emergent topic recently – Paul et al. (2024) formalizes reasoning as a mediator and propose FRODO method, which separates rationale generation and reasoning to improve mediated effects and OOD robustness. Similarly, Tutek et al. (2025) quantify parametric faithfulness in a causal framework instead of contextual faithfullness which was done in other works and what we do in this work.

However, prior studies intervene only on free-form reasoning text, making edits hard to control and benchmark targets difficult to align. In contrast, we select tasks with gold structured mediators—rubrics, checklists, proof graphs, and verifier queries—that serve as explicit causal handles. By intervening on these mediators, we measure model faithfulness under both idealized (gold mediator) and realized (predicted mediator) settings, allowing us to disentangle failures from incorrect structure generation versus failures from ignoring the structure altogether.

5 RESULTS

We evaluate our protocol on four benchmarks using models from four LLM families: Qwen 3, Falcon 3, Gemma 2, and LLaMA 3. Our experiments cover a range of model sizes, including Qwen (1.7B, 4B, 8B) (Yang et al., 2025), Falcon (Team, 2024) (3B, 7B), LLaMA Grattafiori et al. (2024)



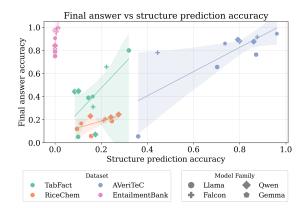


Figure 4: Left: with gold structure, the model is provided with ground-truth intermediate structure M. With predicted structure, the model has to generate intermediate structure M from scratch. Surprisingly, for some models the latter approach works better. Right: We measure correlation between accuracies of predicting intermediate structure M and final answer Y. As expected, it is positive for all datasets.

Table 1: R² scores along with confidence intervals for final answer / structure prediction accuracy.

Dataset	R^2	95% CI
TabFact	0.408	(0.003, 0.867)
RiceChem	0.371	(0.040, 0.910)
AVeriTeC	0.600	(0.066, 0.964)
EntailmentBank	0.347	(0.175, 0.780)

(3.1 8B, 3.2 1B, 3.2 3B), and Gemma Team et al. (2024) (2B). All models are instruction-tuned. For Qwen 3, we disable the built-in reasoning mode to ensure fair comparison across models. In Section 5.1 we compare performance metrics between gold and predicted evaluation settings, and in Section 5.2 we analyze faithfulness across evaluation settings and interventions.

5.1 Performance in Gold and predicted settings

Gold structures do not always improve task performance. We compare performance in two settings: $gold\ structure$, where the model is provided with ground-truth intermediate structure M, and $predicted\ structure$, where the model has to first generate M from scratch. On average across four datasets, models perform better when gold structures are available (Figure 3). However, this is not uniformly true: on EntailmentBank, some models achieve higher accuracy when generating their own structures (Figure 4, left). Across all datasets, we observe such improvements in 7 out of 36 cases, with gains up to 49.5% (see Figure 7 in Appendix A.3). We hypothesize that this occurs when the gold structure M is highly unlikely under the model's predictive distribution. Conditioning on such unlikely prefixes pushes the model into an out-of-distribution regime where its capabilities are weaker, while unconstrained generation allows it to stay closer to training distribution and leverage Chain-of-Thought reasoning (Wei et al., 2022).

Accurate structure prediction correlates with final task accuracy. We also measure the correlation between accuracy in predicting intermediate structures and accuracy on the final answer. By design, we expect models to solve tasks more reliably when they correctly reconstruct the intermediate structure. Indeed, on Figure 4, right side, we observe positive correlations across all datasets, confirming that structural accuracy and decision accuracy are closely linked. For a quantitative evaluation, we report R^2 with 95% confidence intervals obtained via bootstrapping in Table 5.1.

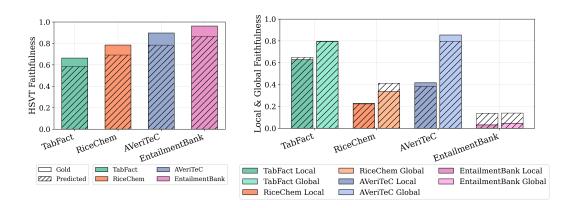


Figure 5: Faithfulness of LLMs to intermediate structures under different interventions. Results are reported for 4 datasets and averaged across all models. Left: reliance on structures (HSVT). Right: causal faithfulness under local vs. global interventions, shown separately for gold and predicted structures (solid vs. striped bars).

5.2 FAITHFULNESS ANALYSIS

To determine whether models rely on their intermediate representations or treat them as mere decorations, we apply our intervention protocol (described in Section 2.3) to both "gold" and "predicted" settings. This allows us to distinguish between failures to generate accurate representations and failures to utilize even accurate ones. The results are summarized in Figure 5, and key findings are discussed below.

Models rely on intermediate structures, but they are not the dominant factor in their decisions. High average HSVT scores across all datasets (with a mean exceeding 60% and reaching over 75% for most individual models; see Appendix A.2) indicate that models rely primarily on the provided structure for information. When the surface text is changed (e.g., paraphrased or replaced with semantically similar inputs) or even totally replaced, the model's decision remains consistent, suggesting that it does not use input text cues. However, low LEC and GEC scores reveal a disconnect: LEC is below 40% on 3 out of 4 datasets and GEC falls below 40% on 2 out of 3 datasets, dropping to under 15% on EntailmentBank. This suggests that reliance on structure does not lead to causal mediation. Even when the structure is altered in ways that contradict the correct answer logically (e.g., flipping rubric items or breaking logical chains), the model often makes the same final decision. This suggests that the model bypasses the explicit reasoning process and instead relies on indirect inferences from the input or its internal knowledge (U in the causal graph in Figure 2).

Global structural breaks are more disruptive than local edits. As shown in Figure 5, global interventions (such as inverting all rubric entries in RiceChem or removing all edges supporting a hypothesis in EntailmentBank) consistently elicit stronger reactions from models than local edits (e.g., flipping a single checklist item). In two of three datasets, global faithfulness exceeds local by 20–40%. This suggests that models tolerate minor inconsistencies within the structure, treating them as noise or irrelevant details, while radical, system-wide changes are more likely to disrupt the model's internal state and trigger a re-evaluation.

Self-generated intermediate structures lead to higher faithfulness despite lower accuracy. While models rely more heavily on gold structures (higher HSVT), they exhibit a greater degree of causal faithfulness to the predicted ones: LEC and GEC scores are consistently higher when intervening on self-generated mediators. On average, the gap is 10–15%, but it reaches 20–25% in many cases—and exceeds 60% for specific models (see Appendix A.2). This indicates that although models consider gold structures to be more authoritative anchors, they are also more responsive to logical changes in structures they have generated themselves. This trend contrasts with overall task

performance, as shown in Section 5.1: gold structures typically lead to higher final-answer accuracy, whereas self-generated structures can sometimes harm performance. However, paradoxically, it is in the predicted setting – where accuracy is often lower – that models demonstrate a stronger sensitivity to edits. We hypothesize that generating a structure is an integral part of a model's reasoning process, as it embeds the model's understanding into the generated structure, making subsequent edits more effective. In contrast, passively consuming a predefined "correct" structure may lead a model to treat it as context, ignoring the logical content, even if that structure improves accuracy on the final task. This challenges the assumption that providing a perfect CoT (contextualized output) or rubric (set of rules) will ensure faithful reasoning, as correctness does not necessarily imply causal mediation.

Soft interventions preserve faithfulness more than the hard ones. Our experiments show that "soft" interventions (such as paraphrasing or entity renaming) maintain faithfulness (i.e., high HSVT and moderate LEC/GEC) more than "hard" interventions (like swapping entire answers or table rows). Notably, HSVT under soft interventions is 20–30% higher on AVeriTeC and EntailmentBank compared to RiceChem and TabFact. This suggests that while models are structurally anchored, they still have some sensitivity to surface features of the input. Hard interventions can create a mismatch between the input and the model's structure, which the model may not be able to resolve without resorting to latent shortcuts.

Figure 5 presents aggregated trends, but these patterns hold consistently at the model level. Notably, model size has little effect on faithfulness — small and large models exhibit similarly varied behavior. For details, see Appendix A.2.

In summary, our causal intervention analysis reveals a systematic gap between structural reliance and causal fidelity: models rely on intermediate structures when surface cues vary, but largely ignore their logical content when it is altered. Faithfulness is not a binary property - it depends on the type of intervention (global > local), the source of the structure (predicted > gold), and the nature of the change (soft > hard). These findings emphasize that intermediate structures are not always automatic causal mediators, even if they appear central to the output of a model.

6 Conclusion

We presented a causal protocol for evaluating the faithfulness of LLMs to their intermediate structures. By intervening on structured reasoning outputs such as rubrics, checklists, proof graphs, and table queries, we assessed whether model decisions are causally mediated by these structures or bypassed through latent shortcuts. Our experiments across four datasets and nine modern LLMs reveal that while models rely on intermediate structures when surface cues are altered (high HSVT scores), they often fail to update their predictions under logically significant local or global edits, indicating weak causal dependence. Surprisingly, models are more faithful to their self-generated structures than to gold ones, suggesting that the act of generation embeds reasoning into the structure in ways that passive consumption does not. These results challenge the assumption that providing perfect reasoning paths guarantees faithful use and highlight the need for new training and prompting methods that enforce causal reliance on intermediate structures. More broadly, our results suggest that intermediate structures could serve as actionable mediators for human—AI collaboration through structured reasoning interventions, but current LLMs do not reliably support this use.

REPRODUCIBILITY STATEMENT

We have taken several steps to ensure the reproducibility of our work. All four datasets used in this study (RiceChem, Averitec, EntailmentBank, TabFact) are publicly available. Our evaluation protocol is described in Section 2.3, with implementation details and deterministic decoding settings. We release prompts used for both gold and predicted mediator settings in Appendix A.1, ensuring that intervention strategies can be replicated. The large language models we evaluate (Qwen 3, LLaMA 3, Falcon 3, Gemma 2) are publicly accessible in instruct-tuned versions. Finally, our source code for running interventions, computing counterfactual targets, and reproducing all metrics and figures is provided in the supplementary material to facilitate replication of results.

USE OF LARGE LANGUAGE MODELS

Large Language Models (LLMs) were used in this work as an assistive tool for polishing the text, improving clarity, and suggesting alternative phrasings. They were not used for research ideation, experimental design, analysis, or result generation. All scientific contributions, experiments, and conclusions are the responsibility of the authors.

REFERENCES

- Youssra Amazou, Faouzi Tayalati, Houssam Mensouri, Abdellah Azmani, and Monir Azmani. Accurate ai assistance in contract law using retrieval-augmented generation to advance legal technology. *International Journal of Advanced Computer Science and Applications*, 16, 02 2025. doi: 10.14569/IJACSA.2025.01602113.
- Bowen Baker, Joost Huizinga, Leo Gao, Zehao Dou, Melody Y Guan, Aleksander Madry, Wojciech Zaremba, Jakub Pachocki, and David Farhi. Monitoring reasoning models for misbehavior and the risks of promoting obfuscation. *arXiv preprint arXiv:2503.11926*, 2025.
- Adrian Bussone, Simone Stumpf, and Dympna O'Sullivan. The role of explanations on trust and reliance in clinical decision support systems. In 2015 international conference on healthcare informatics, pp. 160–169. IEEE, 2015.
- Ángel Alexander Cabrera, Adam Perer, and Jason I Hong. Improving human-ai collaboration with descriptions of ai behavior. *Proceedings of the ACM on Human-Computer Interaction*, 7 (CSCW1). doi: 10.1145/3579612. URL https://par.nsf.gov/biblio/10602196.
- Yanda Chen, Ruiqi Zhong, Narutatsu Ri, Chen Zhao, He He, Jacob Steinhardt, Zhou Yu, and Kathleen McKeown. Do models explain themselves? counterfactual simulatability of natural language explanations. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 7880–7904, 2024.
- James Chua and Owain Evans. Are deepseek r1 and other reasoning models more faithful? *arXiv* preprint arXiv:2501.08156, 2025.
- Bhavana Dalvi, Peter Jansen, Oyvind Tafjord, Zhengnan Xie, Hannah Smith, Leighanna Pipatanangkura, and Peter Clark. Explaining answers with entailment trees. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Alon Jacovi and Yoav Goldberg. Towards faithfully interpretable nlp systems: How should we define and evaluate faithfulness? In *58th Annual Meeting of the Association for Computational Linguistics*, *ACL 2020*, pp. 4198–4205. Association for Computational Linguistics (ACL), 2020.
- Jan Hendrik Kirchner, Yining Chen, Harri Edwards, Jan Leike, Nat McAleese, and Yuri Burda. Prover-verifier games improve legibility of llm outputs. *arXiv preprint arXiv:2407.13692*, 2024.
- Tomek Korbak, Mikita Balesni, Elizabeth Barnes, Yoshua Bengio, Joe Benton, Joseph Bloom, Mark Chen, Alan Cooney, Allan Dafoe, Anca Dragan, et al. Chain of thought monitorability: A new and fragile opportunity for ai safety. *arXiv preprint arXiv:2507.11473*, 2025.
- Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, et al. Measuring faithfulness in chain-of-thought reasoning. *arXiv preprint arXiv:2307.13702*, 2023.
- Gyubok Lee, Hyeonji Hwang, Seongsu Bae, Yeonsu Kwon, Woncheol Shin, Seongjun Yang, Minjoon Seo, Jong-Yeup Kim, and Edward Choi. Ehrsql: A practical text-to-sql benchmark for electronic health records. *Advances in Neural Information Processing Systems*, 35:15589–15601, 2022.

- Haoyang Li, Jing Zhang, Cuiping Li, and Hong Chen. Resdsql: Decoupling schema linking and skeleton parsing for text-to-sql. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 13067–13075, 2023.
 - Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. Faithful chain-of-thought reasoning. In *The 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (IJCNLP-AACL 2023)*, 2023.
 - Katie Matton, Robert Ness, John Guttag, and Emre Kiciman. Walk the talk? measuring the faithfulness of large language model explanations. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=4ub9gpx9xw.
 - Debjit Paul, Robert West, Antoine Bosselut, and Boi Faltings. Making reasoning matter: Measuring and improving faithfulness of chain-of-thought reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 15012–15032, 2024.
 - Judea Pearl. Direct and indirect effects. *Probabilistic and Causal Inference*, 2001. URL https://api.semanticscholar.org/CorpusID:5947965.
 - Michael Schlichtkrull, Zhijiang Guo, and Andreas Vlachos. Averitec: A dataset for real-world claim verification with evidence from the web. *Advances in Neural Information Processing Systems*, 36: 65128–65167, 2023.
 - Tom Silver, Varun Hariprasad, Reece S Shuttleworth, Nishanth Kumar, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. Pddl planning with pretrained large language models. In *NeurIPS 2022 foundation models for decision making workshop*, 2022.
 - Shashank Sonkar, Kangqi Ni, Lesa Tran Lu, Kristi Kincaid, John S Hutchinson, and Richard G Baraniuk. Automated long answer grading with ricechem dataset. In *International Conference on Artificial Intelligence in Education*, pp. 163–176. Springer, 2024.
 - Falcon-LLM Team. The falcon 3 family of open models, December 2024. URL https://huggingface.co/blog/falcon3.
 - Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. Gemma 2: Improving open language models at a practical size. *arXiv* preprint arXiv:2408.00118, 2024.
 - Miles Turpin, Julian Michael, Ethan Perez, and Samuel Bowman. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. *Advances in Neural Information Processing Systems*, 36:74952–74965, 2023.
 - Martin Tutek, Fateme Hashemi Chaleshtori, Ana Marasović, and Yonatan Belinkov. Measuring chain of thought faithfulness by unlearning reasoning steps. *arXiv preprint arXiv:2502.14829*, 2025.
 - Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6):186345, 2024.
 - Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
 - Jianshu Chen Yunkai Zhang Hong Wang Shiyang Li Xiyou Zhou Wenhu Chen, Hongmin Wang and William Yang Wang. Tabfact: A large-scale dataset for table-based fact verification. In *International Conference on Learning Representations (ICLR)*, Addis Ababa, Ethiopia, April 2020.
 - An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

Donghan Yu, Sheng Zhang, Patrick Ng, Henghui Zhu, Alexander Hanbo Li, Jun Wang, Yiqun Hu, William Yang Wang, Zhiguo Wang, and Bing Xiang. Decaf: Joint decoding of answers and logical forms for question answering over knowledge bases. In *The Eleventh International Conference on Learning Representations*.

A APPENDIX

648

649 650

651 652

653

654 655

656

657

658

659

660

661

662

663

664

665

666

669

670

671

673

674

675

676

677

678

679

680

681

682

685

686

688

689 690

691 692

693

694

695

696

697

698

699

700

701

A.1 PROMPTS

Here we present prompts for RiceChem and Tabfact with predicted mediator. In Gold mediator we pass the model the gold mediator as generated one and only ask the model to output its final target.

RiceChem prompt (predicted mediator).

Role: You are an automated grader for a college-level chemistry class. First fill the *checklist* strictly from the student's answer, then compute the final grade.

Instructions

- Use only the given question and student's answer.
- Keep the checklist exactly as provided (same order, wording, and weights).
- Replace each trailing <True/False> with True or False.
- Mark an item True only if it is explicitly satisfied in the answer.
- The final grade equals the sum of weights for items marked True.
- Output only:
 - 1. Checklist: (the filled checklist, line-for-line)
 - 2. Final grade (0--8): <float>

Few-shot examples

Example 1 Question: Why does removing successive electrons from silicon require increasing energy? Answer: Each removal decreases repulsion and increases attraction to the nucleus; the big jump occurs after the 4th electron since the 5th is from a lower shell. Checklist: cites decreased electron–electron repulsion (weight: 1.0) (True/False): True links repulsion decrease to stronger nuclear attraction (weight: 1.0) (True/False): True explains 3rd/4th electrons same shell/core charge (weight: 1.0) (True/False): True explains 5th electron from lower shell (weight: 1.0) (True/False): True Final grade (0–8): 4.0

Example 2 Question: Why can atoms absorb only certain frequencies, but eject electrons with any frequency above a threshold? Answer: Absorption requires exact energy differences between levels; ejection only requires exceeding the threshold. Checklist: states energy levels are quantized (weight: 1.5) (True/False): True explains matching frequency to level gap (weight: 2.0) (True/False): True notes threshold energy for ejection (weight: 1.0) (True/False): True additional energy becomes kinetic energy (weight: 1.0) (True/False): True Final grade (0–8): 5.5

Question

```
{QUESTION_TEXT}
```

Student Answer

{STUDENT_ANSWER}

Checklist

```
{ITEM_1} (weight: {w_1}) (True/False): <True/False>
{ITEM_2} (weight: {w_2}) (True/False): <True/False>
{ITEM_3} (weight: {w_3}) (True/False): <True/False>
...
```

TabFact prompt (predicted mediator).

Role: You are an expert table fact-checking system. First construct a *Verifier Query* in the provided DSL that encodes the reasoning steps, then return the execution result as the final verdict.

Instructions

- Use only the given claim and table.
- Construct a syntactically valid DSL expression that ends with =True or =False.
- The query must encode the logical verification of the claim (e.g., compare values, filter rows, aggregate).
- The final verdict is the boolean result of executing the query: True or False.
- Output only:
 - 1. Verifier Query: ¡DSL expression;

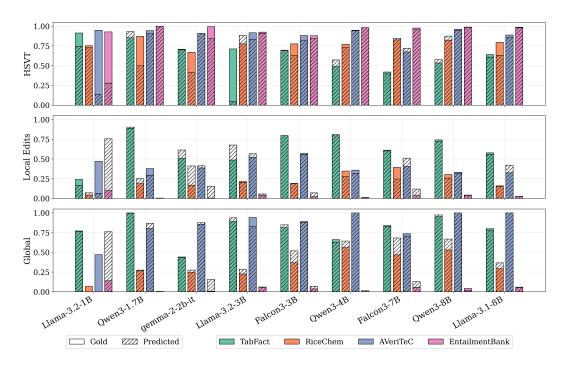


Figure 6: Faithfulness scores (HSVT, Local Edits, Global Edits) for each individual model across all four datasets. Results are shown separately for gold (solid) and predicted (striped) structures.

```
2. Execution Result:
                              <True/False>
Few-shot examples
Example 1 Table:
rank#athlete#nation#gold
1#Usain Bolt#Jamaica#2
2#Shawn Crawford#United States#1
Claim: Usain Bolt won more gold medals than Shawn Crawford.
Verifier
           Query:
                          greater{hop{filter_eq{all_rows; athlete;
Usain Bolt}; gold}; hop{filter_eq{all_rows; athlete; Shawn
Crawford}; gold}}=True
Execution Result: True
Example 2 Table:
player#team#goals
Messi#PSG#30
Ronaldo#AlNassr#25
Claim: Ronaldo scored more goals than Messi.
Verifier Query: greater{hop{filter_eq{all_rows; player; Ronaldo};
goals}; hop{filter_eq{all_rows; player; Messi}; goals}}=True
Execution Result: False
Table
{TABLE_CONTENT}
Claim
```

A.2 FAITHFULNESS

{STATEMENT}

Figure 6 presents per-model faithfulness results without aggregation. While the main trends are already visible in the averaged Figure 5, the non-aggregated view reveals additional insights:

 No clear scaling trend: The size of the model does not consistently correlate with its faithfulness. Larger models do not always outperform smaller ones in terms of LEC or

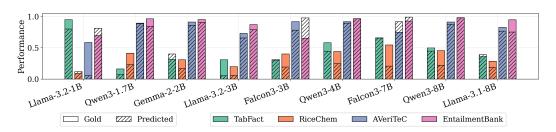


Figure 7: Performance scores (final answer prediction accuracy) for each individual model across all four datasets. Results are shown separately for gold (solid) and predicted (striped) structures.

GEC. The behavior of the models varies significantly even within the same family, making it difficult to predict which model will perform best.

• **Predicted is better than Gold holds per model.** Despite this lack of consistency, the paradox of higher causal faithfulness to self-generated structures is evident at the individual model level. Most models show higher LEC and GEC under predicted mediators, even though the gold scores may be lower.

A.3 PERFORMANCE

Figure 7 presents final answer prediction accuracy for each model without aggregation over datasets.

Gold structure is not always beneficial. In 7 out of 36 cases models perform better without gold structure. This happens at least once on every dataset, suggesting that it could be a general phenomenon, not related to individual dataset idiosyncrasies.