

Garbage In, Reasoning Out? Why Benchmark Scores are Unreliable and What to Do About It

Seyed Mahed Mousavi[†], Edoardo Cecchinato^{†*}, Lucia Horníková^{‡*}, Giuseppe Riccardi[†]

[‡] Masaryk University, Czech Republic

[†] Signals and Interactive Systems Lab, University of Trento, Italy

{mahed.mousavi, giuseppe.riccardi}@unitn.it

Abstract

We conduct a systematic audit of three widely used social reasoning benchmarks, SocialIQA, FauxPas-EAI, and ToMi, and uncover pervasive flaws in both benchmark items and evaluation methodology. Using five LLMs (GPT-3, 3.5, 4, o1, and LLaMA 3.1) as diagnostic tools, we identify structural, semantic, and pragmatic issues in benchmark design (e.g., duplicated items, ambiguous wording, and implausible answers), as well as scoring procedures that prioritize output form over the reasoning process. Through systematic human annotation and re-evaluation on cleaned benchmark subsets, we find that model scores often improve not due to erratic surface wording variations and not to improved reasoning. In fact, further analyses show that model performance is highly sensitive to minor input variations such as context availability and phrasing, revealing that high scores may reflect alignment with format-specific cues rather than consistent inference based on the input. These findings challenge the validity of current benchmark-based claims about social reasoning in LLMs, and highlight the need for evaluation protocols that assess reasoning as a process of drawing inference from available information, rather than as static output selection. We release audited data and evaluation tools to support more interpretable and diagnostic assessments of model reasoning¹.

1 Introduction

The state-of-the-art performance of Large language models (LLMs) across downstream tasks has led to a surge of interest in their potential cognitive and reasoning capabilities, including factuality (Mousavi et al., 2025), common sense reasoning (Sap et al., 2019; Joshi et al., 2025), and theory of mind (Sarangi et al., 2025). While several studies embrace LLMs’ strong performance on various

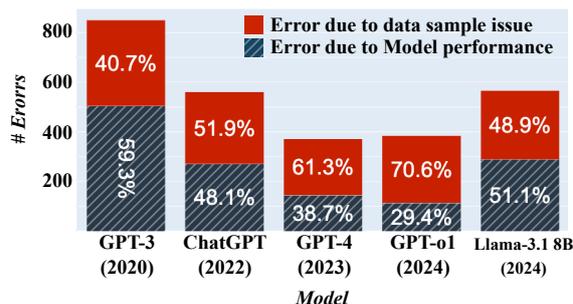


Figure 1: Breakdown of total errors for each model on SocialIQA dev set, divided into *reasoning errors* (errors due to model shortfalls, bottom segment, hatched blue) and *data errors* (errors due to benchmark flaws, top segment, solid red). A substantial portion of errors in each model is consistently due to faulty data.

reasoning benchmarks (Webb et al., 2023; Zhang et al., 2025; Liu et al., 2025), other works argue that such performance may be attributed to superficial pattern recognition (Shapira et al., 2024), dataset artifacts (Jiang et al., 2024), or prompt sensitivity (Mondorf and Plank, 2024). This divergence in findings underscores the need to reevaluate existing benchmarks, as it remains unclear whether they reflect the models’ ability to reason and draw inferences, or merely mask underlying limitations in how models process and generalize information.

Despite the growing use of benchmarks to evaluate the reasoning abilities of LLMs, closer examination reveals that many of these datasets and evaluation protocols suffer from critical flaws. Much of the current assessment of model reasoning relies on surface-level performance metrics that conflate output accuracy with reasoning capability (Mondorf and Plank, 2024). Models often produce correct outputs through inconsistent or fallacious reasoning processes, revealing a disconnect between prediction and explanation (Liu et al., 2024; Jiang et al., 2024). Recent studies (Shapira et al., 2024; Holliday et al., 2024; Jiang et al., 2024) further empha-

*Equal Contribution

¹Code and data available at [Repo Link](#).

size that benchmark success is often fragile, and highly sensitive to minor rewordings, input context, or task framing. These issues cast doubt on the validity of benchmark-derived claims about LLM reasoning, and they complicate efforts to draw principled conclusions about models’ cognitive capabilities. As a result, both the strengths and weaknesses of current LLMs may be misrepresented, hindering progress toward more robust models.

In this work, we shift the focus from model scores to the validity of the benchmarks and evaluation protocols themselves. Studies have identified isolated issues, such as the prompt sensitivity of ToM evaluations (Shapira et al., 2024), the mismatch between plausibility and gold labels in a subset ($\approx 6\%$) of SocialIQa (Palta et al., 2024), and self-contradictory reasoning across commonsense datasets (Liu et al., 2024). However, no prior work offers a systematic, human-verified audit of benchmark quality at scale.

Our approach complements analyses like Mondorf and Plank (2024), who argue for evaluation grounded in reasoning processes rather than outcomes, and Jiang et al. (2024), who highlight how token-level biases can distort perceived competence. We follow this perspective by defining reasoning as the process of drawing context-sensitive inferences from available information, rather than relying on superficial cues or memorized patterns. By systematically auditing datasets and evaluation protocols, we aim to expose not only where current benchmarks fail, but also how these failures obscure the reasoning shortfalls of models. Our findings advocate for a shift from static outcome-based metrics to evaluations that treat reasoning as a dynamic, explainable process, highlighting both the need for better data and fundamentally new evaluation paradigms.

We examine the reliability of three widely used reasoning benchmarks, SocialIQa (Sap et al., 2019), FauxPas-EAI (Le et al., 2019), and ToMi (Shapira et al., 2023), by evaluating how five language models (GPT-3, ChatGPT (GPT-3.5), GPT-4, GPT-o1, and LLaMA 3.1 8B Instruct) perform under standard prompting and evaluation settings used in prior work (Shapira et al., 2024). Rather than benchmarking the models themselves, we treat them as diagnostic tools to reveal flaws in the datasets and evaluation procedures. We then conduct a structured human audit to systematically identify both data-level and methodological issues, categorize them, and assess their impact on model

outputs. Furthermore, we observe that model performance is highly sensitive to surface-level variations, i.e. small changes in phrasing or minor data errors can cause large shifts in performance, exposing a brittle reliance on surface cues rather than reasoning process grounded in context. Our findings reveal that current evaluation setups often the reasoning ability, equating lexical alignment with cognitive competence. Based on these insights, we advocate for a shift toward process-oriented evaluation paradigms that better capture inferential competence. We release audited and annotated data, along with evaluation tools, to support more transparent and interpretable assessments of model reasoning in future research. Our contributions are:

- We perform a diagnostic assessment of three widely used social reasoning benchmarks, using five LLMs as tools to uncover data artifacts and methodological inconsistencies.
- We conduct a large-scale, systematic human audit to identify and categorize dataset-level and evaluation-level flaws, and re-evaluate model performance on curated benchmark subsets to quantify their impact.
- We expose fundamental limitations of current evaluation practices and argue for a paradigm shift from static output-based metrics toward reasoning-as-process evaluation frameworks.
- We release the audited data, the annotation materials, and the evaluation tools to support interpretable and diagnostic assessments of reasoning in LLMs.

2 Literature Review

Webb et al. (2023) report that GPT-3 and GPT-4 outperform humans on structured analogical reasoning tasks inspired by cognitive assessments. Shapira et al. (2024) show that model performance on social reasoning tasks is fragile, driven by shallow heuristics, and correlates with dataset age, suggesting data exposure during pretraining, rather than genuine reasoning ability. Holliday et al. (2024) find 1) even top-performing models make basic logical errors and inconsistencies; 2) performance correlates with model size, but no model surpasses 90% accuracy; 3) chain-of-thought prompting improves scores in some settings but fails to eliminate inconsistencies, and in certain cases, introduces new logical failures. Liu et al. (2024) and Jiang

et al. (2024) demonstrate that LLMs frequently produce correct answers using flawed or logically inconsistent reasoning, often relying on superficial cues, and token bias. While the mentioned studies focus mostly on model performance, (Palta et al., 2024) examine the quality of the benchmarks used for such evaluations. The authors sampled 250 questions ($\approx 6\%$) from SocialIQa (Sap et al., 2019) and CommonsenseQA (Talmor et al., 2019) as prominent commonsense reasoning benchmarks, and find that over 22% of samples are problematic. A systematic survey presented by Mondorf and Plank (2024) argues that evaluation practices often conflate correct answers with sound reasoning, and that data leakage, prompt sensitivity, and shallow metrics inflate perceptions of model competence.

3 Garbage-In Model Evaluation

We begin by evaluating model performance under standard conditions commonly adopted in prior work (Shapira et al., 2024). Our aim is not to assess the models themselves, but to use their outputs as diagnostic signals. By combining model outputs with human annotation, we identify how flaws in data construction and evaluation protocols distort reported performance and obscure what is actually being measured.

SocialIQa (Sap et al., 2019) is a crowd-sourced multiple-choice benchmark for evaluating social and emotional commonsense reasoning. Each example consists of a short context, a social inference question, and three answer options, only one of which is correct. As the test set is not publicly released to prevent overfitting and data leakage, we conduct all evaluations on the development set, which contains 1,954 questions. Reported human performance on this subset is 87%.

ToMi (Le et al., 2019) is a synthetic benchmark designed to evaluate Theory of Mind (ToM) reasoning in language models. Each story involves multiple agents and object relocations, followed by six questions probing memory, reality, first-order beliefs, and second-order beliefs. ToMi includes randomized story generation, distractor agents, and varied event sequences to reduce exploitable patterns and test belief reasoning more robustly. We evaluate models on a batch of 100 generated stories (600 questions)².

FauxPas-EAI (Shapira et al., 2023) includes 44

²The set is generated using the official code from Le et al. (2019).

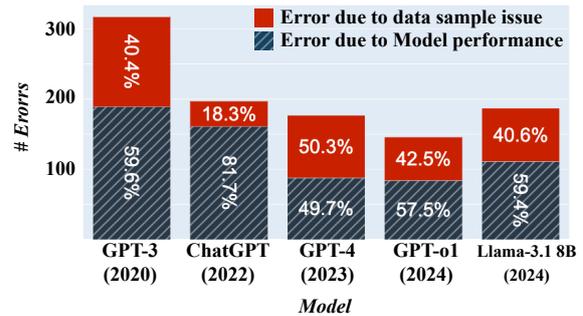


Figure 2: Breakdown of total errors for each model on ToMi, divided into *reasoning errors* (errors due to model shortfalls, bottom segment, hatched blue) and *data errors* (errors due to benchmark flaws, top segment, solid red). Data errors include *Incorrect Ground Truth* and *Answer Granularity Mismatch*.

short narrative stories, 22 containing a social faux pas and 22 matched control stories, each followed by four multiple-choice questions: a detection question, a statement identification question, a comprehension check, and a false-belief probe. The benchmark is adapted from a clinical Theory of Mind test (Baron-Cohen et al., 1999) designed to assess social reasoning and false-belief understanding. As noted by Shapira et al. (2023), FauxPas-EAI can be viewed as a compositional task that combines SocialIQa and ToMi.

We use five LLMs as diagnostic tools in our benchmark analysis. These include GPT-3³ (Brown et al., 2020), an early autoregressive model trained on a broad range of internet text; ChatGPT (GPT-3.5)⁴, a fine-tuned variant optimized for dialogue and instruction-following; GPT-4⁵, a more advanced model with improved capabilities; GPT-o1⁶, a recent optimized release of GPT-4 with enhanced performance; and, LLaMA 3.1 8B Instruct⁷, an open-source instruction-tuned model trained to align with human feedback. Further implementation details are presented in §A.1 and §A.2.

Results As shown in Table 2, model performance under standard evaluation settings reproduces known trends across all three benchmarks. Newer models (GPT-4 and GPT-o1) outperform earlier versions such as GPT-3.5 and GPT-3. On SocialIQa, accuracy ranges from 56% (GPT-3) to 81% (GPT-4), with intermediate scores from ChatGPT and LLaMA 3.1. A similar pattern holds for

³davinci-002

⁴gpt-3.5-turbo-1106

⁵gpt-4-1106-preview

⁶o1-preview

⁷LLaMA-3.1

ToMi, where performance increases from 47% to 76%. On FauxPas-EAI, GPT-4 achieves the highest scores on both question-level and story-level metrics (82%|39%), while other models show greater variability. These results align with previous findings (Shapira et al., 2024) and confirm that our evaluation setup reproduces previously observed performance trends across models and benchmarks. Full implementation details and prompt templates are provided in our repository.

4 Benchmark Audit

We audit benchmark integrity by using LLMs as diagnostic tools; when model predictions appear inconsistent, unstable, or penalized despite valid reasoning process, we investigate whether these behaviors point to underlying flaws in benchmark construction or evaluation. We group these issues into two categories: (1) *Dataset-level issues*, which concern the quality and validity of benchmark items; and (2) *Evaluation-level issues*, which arise from scoring procedures and presentation formats.

4.1 Data-Related Issues

We conducted a manual evaluation of all items across the three datasets. Adapting the human evaluation framework proposed by Mousavi et al. (2022), two researchers reviewed each item. The process was supervised by a senior researcher who developed the annotation guidelines and oversaw the process to ensure consistency and annotation quality. Faulty items were identified and categorized into three major dimensions: 1) structural flaws (e.g., duplicated questions, malformed answer options); 2) semantic flaws (e.g., ambiguous wording, internally inconsistent logic); and, 3) pragmatic flaws (e.g., unrealistic scenarios, culturally biased assumptions).

SocialIQA

Our audit of the SocialIQA development set revealed that more than 28% of the items contain data-related flaws that compromise the benchmark’s reliability for evaluating reasoning:

I. Structural Flaws ($\approx 4\%$) include surface-level or syntactic issues, such as duplicated entries, repetitive answer options, incomplete prompts, broken syntax, and grammatical errors in the context or answer choices.

II. Semantic Flaws ($\approx 18\%$) involve inconsistencies or ambiguities in meaning, such as mismatches in semantic roles (e.g., assigning actions

Error Category	SocialIQA	ToMi
Structural Flaws	85 (4.5%)	24% (50 Qs)
Semantic Flaws	361 (18.5%)	21% (126 Qs)
Pragmatic Flaws	130 (6.7%)	45% (270 Qs)
Total	576 (29.5%)	-

Table 1: Distribution of dataset-level flaws (structural, semantic, and pragmatic) across SocialIQA and ToMi identified through human audit. For SocialIQA, counts and percentages reflect item-level errors across 1,954 development set questions. For ToMi (100 stories, equal to 600 questions), percentages reflect the proportion of stories (and questions) affected within the evaluated batch. The flaws in both benchmarks raise concerns about their reliability for reasoning evaluation.

to the wrong participant), violations of temporal logic, referential confusion, and under-specified or contextually vague questions that allow multiple answer options to be equally plausible.

III. Pragmatic Flaws ($\approx 7\%$) include answer options that are implausible, socially incoherent, or misaligned with real-world expectations. These include responses that fail to reflect reasonable human motivations, emotional reactions, or common social behavior, as well as answers that are awkwardly phrased or pragmatically unnatural given the context.

ToMi

While ToMi benefits from the controllability of synthetic generation, this does not guarantee immunity from annotation or construction errors. Our audit of the ToMi benchmark revealed that a significant portion of items contain data-related flaws:

I. Structural Flaws ($\approx 0.8\%$) include incorrect ground truth labels that contradict the logical implications of the story, leading to valid model responses being incorrectly marked as wrong, (found in 50 questions affecting 24 stories).

II. Semantic Flaws ($\approx 21\%$) involve misclassified story types; for example, labeling a false belief as true belief, which breaks the alignment between the intended reasoning task and the question’s logical demands.

III. Pragmatic Flaws ($\approx 45\%$) include misaligned or inconsistently assigned agent roles, which interfere with perspective-taking tasks (The benchmark defines Agent 0 as the actor who moves the object, and Agent 1 as the other participants). This is particularly damaging in belief and desire questions, where agent identity is central to infer-

ring mental states.

FauxPas-EAI

We identified data-related issues in 4 (10%) of the stories. These include: I) *Structural flaws*, such as incorrect labeling of question–answer pairs (e.g., misattributing the gift giver) and awkward or grammatically incorrect phrasing (e.g., “at the kitchen”); II) *Pragmatic flaws*, involving scenarios that fail to clearly constitute a faux pas, as well as questions that hinge on speculative knowledge (e.g., assuming whether a teacher knows a student’s parent), introducing ambiguity that is difficult to resolve in a multiple-choice format. These flawed items span all three story sources; the original clinical set from [Baron-Cohen et al. \(1999\)](#), the human-authored additions, and the LLM-generated stories by [Shapira et al. \(2023\)](#).

4.2 Evaluation-Level Issues

In addition to dataset-level flaws, we identified several evaluation-level issues that affect how benchmark performance is evaluated and the scores are interpreted. These concern not the content of individual items, but the procedures and assumptions involved in applying the benchmarks and scoring model outputs.

ToMi

Beyond data-level flaws, our analysis of ToMi reveals several evaluation-level issues that further undermine its reliability as a reasoning benchmark:

I. Context Fragmentation In the literature, the ToMi benchmark is typically applied by presenting each question to the model in isolation, without access to prior questions or answers from the same story. In contrast, human participants process the full narrative and associated questions sequentially, accumulating context throughout. This discrepancy limits the model’s ability to condition its output on earlier events, particularly in second-order belief and desire tasks, where reasoning depends on prior mental state inferences. The fragmented prompting setup thus distorts evaluation by disrupting essential context needed for multi-step inference.

II. Semantics-Agnostic Scoring In ToMi, model outputs are evaluated using strict string matching against gold labels, with no tolerance for lexical variation. This scoring method fails to recognize semantically valid responses. As a result, models may be scored as incorrect despite semantic validity. Conversely, answers that share surface

lexical overlap with the gold label may be scored as correct even when they fail to convey the intended meaning.

III. Granularity Mismatch In $\approx 30\%$ of responses marked as incorrect, we observed a granularity mismatch between the expected answer and the model’s semantically valid but coarser output. This issue arises from a lack of specificity in the question prompts, which fail to clearly indicate the level of detail required in the answer. For example, when asked to identify the location of an object, expecting the object container as the valid response, models would sometimes respond with the correct room in which the container (and thus the object) is located. While such answers are semantically valid, they are nonetheless penalized.

FauxPas-EAI

Besides minor data-level issues, we identified several evaluation-level problems that affect how performance is measured on FauxPas-EAI:

I. Context Fragmentation As with ToMi, models are prompted one question at a time in FauxPas-EAI, without access to prior questions or previous answers. In contrast, human evaluators process the full story and all four associated questions as a single, coherent unit. This discrepancy creates a context asymmetry, particularly on identification questions, where reasoning depends on earlier responses. The lack of contextual continuity limits the model’s ability to build and track social inference over multiple steps.

II. Semantics-Agnostic Scoring FauxPas-EAI uses Levenshtein distance ([Levenshtein, 1966](#)) to score model outputs based on string similarity. This scoring approach rewards lexical alignment over interpretive depth, weakening the benchmark’s diagnostic value. It results in paraphrased or logically correct responses being marked incorrect, while answers that share surface-form overlap with the gold label may be scored as correct even when they fail to reflect valid reasoning.

III. Heterogeneous Story Composition FauxPas-EAI combines three qualitatively distinct story sources; 20 clinical samples from [Baron-Cohen et al. \(1999\)](#), 12 human-authored stories, and 12 LLM-generated stories validated post hoc ([Shapira et al., 2023](#)). While treated as a unified benchmark in the literature, these subsets differ in narrative complexity and reasoning demands. The original clinical stories rely on implicit social knowledge and belief modeling, while the newer

Benchmark	#Samples	GPT-3	ChatGPT	GPT-4	GPT-o1	LLaMA 3.1 _T
SocialIQA						
Shapira et al. (2024)	400	19	67	79	-	-
<i>Our Results</i> (section 3)	1954	56	71	81	80	71
Clean Subset (section 5)	1378	63	80	90	92	79
ToMi						
Shapira et al. (2024)	400	39	70	70	-	-
<i>Our Results</i> (section 3)	100	47	67	70	76	68
Clean Set (section 5)	100	69	73	85	86	81
+ Context		66	84	82	91	79
FauxPas-EAI (QIS)						
Shapira et al. (2024)	44	63 14	73 25	74 27	-	-
<i>Our Results</i> (section 3)	44	60 16	75 23	82 39	70 36	61 16
Clean Subset (section 5)	40	59 13	75 23	84 40	71 40	60 15
Source BreakDown: (section 5)						
<i>Baron-Cohen et al. (1999)</i>	18	60 11	79 22	85 39	67 33	65 22
+ Context		71 33	79 39	92 67	74 44	67 22
<i>Shapira et al. (2023) (H.)</i>	11	54 09	77 27	82 36	75 45	57 09
+ Context		59 09	77 18	84 55	80 36	57 27
<i>Shapira et al. (2023) (AI+H.)</i>	11	61 18	66 18	82 45	75 45	57 09
+ Context		55 09	73 18	93 82	77 36	66 36

Table 2: Model performance across the original versions of three benchmarks: SocialIQA, FauxPas-EAI, and ToMi, under different evaluation settings. For each benchmark, we report accuracy (%) on the original dataset (our results and prior work), the cleaned subset, and versions enhanced with additional context. For **FauxPas-EAI**, we report question-level (Q) and story-level (S) accuracy, where the latter requires all four questions for a story to be answered correctly. Performance tends to increase with model recency and improves further with cleaned data and contextual prompting.

stories are more explicit and structurally simplified (Shapira et al., 2023). Merging these stories into a single evaluation set can inflate model scores in misleading ways.

4.3 Analysis and Implications

Our assessment reveals pervasive flaws across all three benchmarks, spanning both dataset- and evaluation-levels. Structural, semantic, and pragmatic flaws (Table 1) compromise the integrity of benchmark items, while semantic-agnostic scoring protocols and context-fragmented prompting fail to reflect the reasoning capability they aim to measure. These errors inject noise into evaluation and distort conclusions. As a result, models may be penalized for appropriate reasoning or rewarded for token overlap. Figures 1 and 2 illustrate that, in both SocialIQA and ToMi, a large share of model “errors” are attributable not to model limitations, but to flaws in benchmark design.

Benchmarks relying on multiple-choice formats (SocialIQA) or pre-defined task schemas (ToMi, FauxPas-EAI) reduce reasoning to surface-level selection. When tasks are framed as choosing from a fixed set of options without requiring explanation or justification, correct reasoning becomes indis-

tinguishable from shallow token overlap. In such settings, correct answers can arise from token-level heuristics rather than conceptual understanding. These setups tend to reward surface-level alignment over valid reasoning processes and equate output correctness with reasoning competence. More broadly, they reflect a deeper limitation of current benchmark design: **benchmarks based on pre-defined schemas and token-matching metrics fail to capture the flexible, contextual nature of reasoning**, and risk conflating format adherence with cognitive ability.

5 High Scores Conceal Reasoning Gaps

To better understand what current benchmark scores reflect, we examine model performance under modified evaluation settings that address flaws identified in our audit. Our goal is to assess how reported performance is shaped by data quality, context availability, input phrasing, and scoring methodology. We analyze four dimensions: cleaned datasets (A), context-aware prompting (B), semantically equivalent rephrasings (C), and meaning-aware scoring using LLMs (D). This analysis probes the stability of model predictions and the robustness of benchmark scores to input-

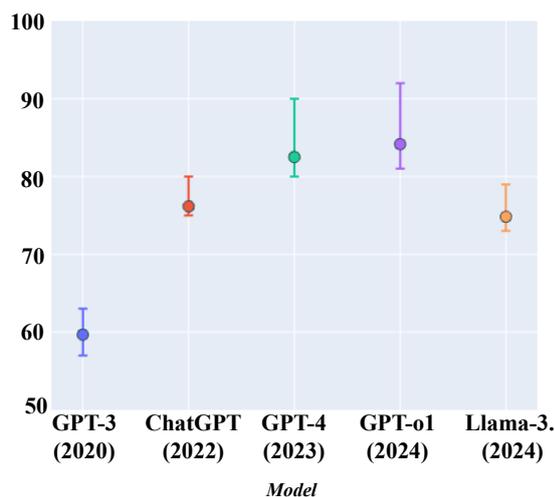


Figure 3: Performance on SocialIQa (clean subset) and five semantically equivalent rephrasings per item. Despite no change in content, all models exhibit notable score fluctuations, revealing sensitivity to surface-level variation and the limitation of the benchmark to assess robust reasoning.

level noise, including annotation inconsistencies, missing context, and superficial lexical variation.

A. Evaluation on Cleaned Data We investigate how performance scores vary when input-noise is minimized. We filtered out flawed items from SocialIQA and FauxPas-EAI, corrected labeling and agent-role inconsistencies in ToMi, adjusted scoring procedures to account for granularity mismatches, and manually reviewed all model responses. As shown in Table 2, all models exhibit improved accuracy on the cleaned datasets. For example, GPT-o1 achieves 92% on the cleaned SocialIQA subset (up from 80%), and GPT-4 improves to 85% on ToMi. FauxPas-EAI also shows modest gains after data cleaning, though performance remains variable. Nevertheless, rather than enhanced reasoning ability, **these gains highlight how sensitive model performance is to input-noise, annotation quality, and benchmark design.** We also examine performance across the three story sources in FauxPas-EAI (clinical, human-authored, and LLM-generated), and observe notable variation across subsets. This variation suggests that differences in narrative structure and complexity, rather than reasoning competence, may account for observed model performance.

B. Context-Aware Evaluation We evaluate models in a context-aware setting, where they are exposed to the full sequence of story questions. This allows to accumulate context across turns, par-

ticularly important for tasks involving belief tracking or multi-step inference. As shown in Table 2, performance improves further under this setup for most models. Nevertheless, the performance of a few models drop slightly in across benchmarks, suggesting that not all models benefit equally from additional context. These results confirm that isolated question prompting underestimates model performance on tasks requiring sequential social reasoning. However, **higher scores in context-aware settings do not signal improved reasoning, but rather expose the inadequacy of fragmented evaluation setups.**

C. Format Sensitivity To investigate the impact of input format on model performance, we generated five alternative phrasings for each item in the cleaned SocialIQA subset using GPT-4o. A random 15% of the generated items were manually validated to ensure semantic equivalence and linguistic correctness. We then re-evaluated model performance on these rephrasings. Despite no change in content, model scores varied considerably, Figure 3. These findings reveal that what performance scores on static benchmarks represent is not reasoning process, but a fragile reliance on specific lexical patterns and token-level cues. The observed volatility underscores that even minimal input variation can destabilize model performance, exposing a brittleness inconsistent with inferential capabilities. **Rather than assessing models’ ability to draw inferences from the input context, current static benchmarks reflect sensitivity to surface-level cues and variations in phrasing.**

D. LLM-as-a-Judge In this evaluation, we relied on human annotations to mitigate the limitations of semantics-agnostic scoring. Nevertheless, we explored the use of LLM-as-a-judge (Ho et al., 2025; Bavaresco et al., 2024) as a scalable cost-effective surrogate for meaning-aware evaluation. We assessed GPT-4o in evaluating the semantic validity of model predictions given the question and gold reference. On FauxPas-EAI items, GPT-4o achieved near-perfect agreement with human annotations (Cohen’s $\kappa = 0.89$), outperforming Levenshtein distance ($\kappa = 0.77$). These findings suggest that LLMs can provide an alternative to rigid string-matching metrics. However, we emphasize that **LLM-as-a-judge should be treated as a diagnostic aid, not a substitute for rigorous human evaluation.**

6 Discussion

As demonstrated in Section 5, high benchmark scores often collapse under input variation or when superficial cues are removed, revealing a persistent gap between what benchmark scores reflect and the actual reasoning processes exhibited by models. Once flawed data is corrected and phrasing is controlled, models struggle to produce stable, contextually grounded inferences. This suggests that current benchmarks do not reliably assess reasoning as inference from input content, but instead reward lexical conformity and alignment with format-specific cues.

Static Benchmarks Conflate Accuracy with Reasoning. Current static benchmarks reduce reasoning to static tasks evaluated through string matching or token-level overlap, failing to reflect reasoning as a dynamic, contextual, and inferential process. Pre-defined formats and static schema encourage pattern matching and fluency over correct reasoning process. Our audit and prior work (Jiang et al., 2024; Liu et al., 2024) show that models often produce correct answers while offering incoherent or contradictory rationales. Even GPT-o1, which scores 92% on the cleaned SocialIQa subset surpassing human accuracy, exhibits brittle behavior when confronted with minor rephrasings or alternative phrasings of the same question. Such instability (Section 5) reveals that what benchmarks reward is not inference, but format sensitivity. *A model can achieve perfect accuracy and still fail to reason.*

Benchmarks Create the Illusion of Progress. Observed high scores are frequently driven by overfitting to benchmark artifacts, not gains in reasoning ability. Many widely-used datasets suffer from contamination, outdated task designs, or structural inconsistencies. This inflates performance through memorization or exposure, rather than generalization. Moreover, models that cannot maintain logical consistency and frequently contradict their own outputs when asked to explain themselves lack the required coherence for reasoner models. Their success is not evidence of reasoning but of matching familiar patterns and memorization. Without contextual robustness or internal consistency, high scores serve as a misleading proxy for cognitive competence.

Toward Process-Aware Evaluation. The shortcomings we identify reflect not just model limitations, but deeper flaws in how reasoning is currently

evaluated. Current benchmarks reward surface-level alignment over inferential capabilities, emphasizing outcomes rather than the reasoning processes that produce them. More meaningful alternatives include structured rationale evaluation, contradiction detection, counterfactual testing, and interactive formats such as adaptive diagnostics (Mondorf and Plank, 2024). These approaches prioritize reasoning as a process rather than an outcome, and offer a more diagnostic view of model competence. By using current benchmarks, we risk mistaking statistical fluency for understanding, and building models that appear intelligent, but are not.

7 Conclusion

We conducted a detailed audit of widely used social reasoning benchmarks and evaluation practices to assess what current methods truly measure in LLMs. Our findings reveal that benchmark performance is frequently shaped by flawed data, fragile evaluation protocols, and a reliance on surface-level cues. Even when models achieve high scores, their outputs often reflect surface lexical alignment rather than input grounded reasoning process. In particular, we find that performance can degrade substantially under semantically equivalent rephrasings or minor data imperfections, highlighting a deep sensitivity to input noise that undermines reliability.

We argue that meaningful progress in developing and evaluating reasoning capabilities requires methods that treat reasoning as a contextual, coherent, and explainable process. To move forward, we must shift away from static, form-driven metrics toward process-level evaluation, tracking reasoning steps, identifying contradictions, and testing robustness across rephrasings and contexts. Without such a shift, benchmark success will continue to misrepresent what LLMs can actually do, and claims of reasoning competence will remain on shaky ground.

Limitations

Our study focuses only social reasoning and theory-of-mind tasks, leaving logical and mathematical reasoning benchmark unexplored. Our rephrasing experiments rely on a single backbone model, which may limit the generality and diversity of our format sensitivity findings. While our human audit is systematic, its scale is resource-intensive. Extending this diagnostic approach to more diverse

reasoning tasks, larger datasets, and multilingual settings remains important future work.

References

- Simon Baron-Cohen, Michelle O’riordan, Valerie Stone, Rosie Jones, and Kate Plaisted. 1999. Recognition of faux pas by normally developing children and children with asperger syndrome or high-functioning autism. *Journal of autism and developmental disorders*, 29:407–418.
- Anna Bavaresco, Raffaella Bernardi, Leonardo Bertolazzi, Desmond Elliott, Raquel Fernández, Albert Gatt, Esam Ghaleb, Mario Giulianelli, Michael Hanna, Alexander Koller, et al. 2024. Llms instead of human judges? a large scale empirical study across 20 nlp evaluation tasks. *arXiv preprint arXiv:2406.18403*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Xanh Ho, Jiahao Huang, Florian Boudin, and Akiko Aizawa. 2025. Llm-as-a-judge: Reassessing the performance of llms in extractive qa. *arXiv preprint arXiv:2504.11972*.
- Wesley H. Holliday, Matthew Mandelkern, and Cedegao E. Zhang. 2024. [Conditional and modal reasoning in large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 3800–3821, Miami, Florida, USA. Association for Computational Linguistics.
- Bowen Jiang, Yangxinyu Xie, Zhuoqun Hao, Xiaomeng Wang, Tanwi Mallick, Weijie J Su, Camillo Jose Taylor, and Dan Roth. 2024. [A peek into token bias: Large language models are not yet genuine reasoners](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4722–4756, Miami, Florida, USA. Association for Computational Linguistics.
- Abhinav Joshi, Areeb Ahmad, Divyaksh Shukla, and Ashutosh Modi. 2025. [Towards quantifying commonsense reasoning with mechanistic insights](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 9633–9660, Albuquerque, New Mexico. Association for Computational Linguistics.
- Matthew Le, Y-Lan Boureau, and Maximilian Nickel. 2019. [Revisiting the evaluation of theory of mind through question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5872–5877, Hong Kong, China. Association for Computational Linguistics.
- Vladimir I. Levenshtein. 1966. Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Doklady*, 10(8):707–710.
- Wei Liu, Ruochen Zhou, Yiyun Deng, Yuzhen Huang, Junteng Liu, Yuntian Deng, Yizhe Zhang, and Junxian He. 2025. [Learn to reason efficiently with adaptive length-based reward shaping](#). *arXiv preprint arXiv:2505.15612*.
- Ziyi Liu, Soumya Sanyal, Isabelle Lee, Yongkang Du, Rahul Gupta, Yang Liu, and Jieyu Zhao. 2024. [Self-contradictory reasoning evaluation and detection](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3725–3742, Miami, Florida, USA. Association for Computational Linguistics.
- Philipp Mondorf and Barbara Plank. 2024. [Beyond accuracy: Evaluating the reasoning behavior of large language models - a survey](#). In *First Conference on Language Modeling*.
- Seyed Mahed Mousavi, Simone Alghisi, and Giuseppe Riccardi. 2025. Llms as repositories of factual knowledge: Limitations and solutions. *arXiv preprint arXiv:2501.12774*.
- Seyed Mahed Mousavi, Gabriel Roccabruna, Michela Lorandi, Simone Caldarella, and Giuseppe Riccardi. 2022. [Evaluation of response generation models: Shouldn’t it be shareable and replicable?](#) In *Proceedings of the 2nd Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 136–147, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Shramay Palta, Nishant Balepur, Peter Rankel, Sarah Wiegrefe, Marine Carpuat, and Rachel Rudinger. 2024. [Plausibly problematic questions in multiple-choice benchmarks for commonsense reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3451–3473, Miami, Florida, USA. Association for Computational Linguistics.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. [Social IQa: Commonsense reasoning about social interactions](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4463–4473, Hong Kong, China. Association for Computational Linguistics.
- Sneheel Sarangi, Maha Elgarf, and Hanan Salam. 2025. [Decompose-ToM: Enhancing theory of mind reasoning in large language models through simulation and task decomposition](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 10228–10241, Abu Dhabi, UAE. Association for Computational Linguistics.

- Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuhui Zhou, Yejin Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. 2024. [Clever hans or neural theory of mind? stress testing social reasoning in large language models](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2257–2273, St. Julian’s, Malta. Association for Computational Linguistics.
- Natalie Shapira, Guy Zwirn, and Yoav Goldberg. 2023. [How well do large language models perform on faux pas tests?](#) In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10438–10451, Toronto, Canada. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Taylor Webb, Keith J Holyoak, and Hongjing Lu. 2023. Emergent analogical reasoning in large language models. *Nature Human Behaviour*, 7(9):1526–1541.
- Zhen Zhang, Xuehai He, Weixiang Yan, Ao Shen, Chenyang Zhao, Shuohang Wang, Yelong Shen, and Xin Eric Wang. 2025. Soft thinking: Unlocking the reasoning potential of llms in continuous concept space. *arXiv preprint arXiv:2505.15778*.

A Appendix

A.1 Implementation Details

For every benchmark, we ran the experiments using five different models: GPT-3⁸, ChatGPT (GPT-3.5)⁹, GPT-4¹⁰, GPT-o1¹¹, and LLaMA 3.1 8B Instruct¹². We used the official OpenAI APIs¹³ for the first four and the HuggingFace APIs for LLaMA 3.1¹⁴. We set the temperature of all the models to 0 for reproducibility reasons, with the only exception of GPT-o1 for which is set to 1 by default. We queried every model sequentially for every data entry, and following we present the number of queries for each benchmark and related tasks.

SocialIQA

We ran models evaluation for all the development set (1954 items), that for all the five models implies 9770 calls to the APIs. After this step, we generated five rephrasings for every data entry of the cleaned version of the benchmark (1378 items). To do so we queried GPT-4o¹⁵ 6890 times, using a temperature value of 1. Finally, we ran again all the initial models on the rephrasings (6890 per model), adding a number of 34,450 API calls. At the end of the experiment, for SocialIQA we obtained a total 51,110 calls.

ToMi

For ToMi we generated 100 stories from the original repository¹⁶, each containing 6 questions. For all the models, we initially evaluated every question sequentially (3000 API calls), and then providing the conversation history (3000 API calls). The conversation history consists of adding the previously answered questions for the same story, in addition to the context and the current question. In the end, the total number of queries for this benchmark was 6000.

FauxPas-EAI

FauxPas-EAI consists of 44 stories, each containing 4 questions (176 questions in total). Like we did for ToMi, we evaluated each model sequentially (880 API calls) and with the conversation history (880 API calls). Then, on the cleaned dataset, containing 40 stories (160 questions), we queried GPT-4o (temperature set to 0) as LLM-as-a-Judge evaluator instead of using Levenshtein distance. Doing this for every model predictions we executed 800 additional API calls. These LLM annotations were then compared with the humans ones using the Cohen’s κ value, implemented in this repository¹⁷. In the end, the final total number of queries for FauxPas-EAI was 2560.

A.2 Prompt Examples

Table 3 and 4 contain all the prompts used to run our experiments. The first one contains the benchmarks evaluation prompts, while the second one contains the prompts used for the SocialIQA rephrasing task and the FauxPas-EAI LLM-as-a-Judge evaluation.

⁸[davinci-002](#)

⁹[gpt-3.5-turbo-1106](#)

¹⁰[gpt-4-1106-preview](#)

¹¹[o1-preview](#)

¹²[LLaMA-3.1](#)

¹³[OpenAI developer platform](#)

¹⁴[meta-llama/Llama-3.1-8B-Instruct](#)

¹⁵[gpt-4o-2024-08-06](#)

¹⁶[facebookresearch/ToMi](#)

¹⁷[sislab-unitn/Human-Evaluation-Protocol](#)

Benchmark	Prompt
SocialIQA Sequential	<i>Context: {context}</i> <i>Question: {question}</i> <i>Based on the context, which option is the best answer to the question?</i> <i>Option 1: {option1}</i> <i>Option 2: {option2}</i> <i>Option 3: {option3}</i> <i>Answer with "1", "2" or "3", without explanations. In case of doubt, answer according to the most probable answer.</i>
ToMi Sequential	<i>Context: {context}</i> <i>Question: {question}</i> <i>Answer with one word only, without explanations. Answer:</i>
Conv-history	<pre> [{"role": "user", "content": {prompt1}} {"role": "assistant", "content": {model_answer1}}, {"role": "user", "content": {prompt2}}, {"role": "assistant", "content": {model_answer2}}, ... {"role": "user", "content": {prompt6}} </pre>
FauxPas-EAI Sequential	<i>{story}</i> <i>{question}</i> <i>Answer:</i>
Conv-history	<pre> [{"role": "user", "content": {prompt1}} {"role": "assistant", "content": {model_answer1}}, {"role": "user", "content": {prompt2}}, {"role": "assistant", "content": {model_answer2}}, ... {"role": "user", "content": {prompt4}} </pre>

Table 3: Prompts used for the benchmarks evaluation, both sequentially and with conversation history. In SocialIQA, for GPT-3 and LLaMA 3.1 we slightly changed the prompt structure, adding some guideline for the specific output formatting.

Task	Prompt
SocialQA rephrasing	<p>Your job is to generate 5 different rephrasings of the given item. For each rephrasing, rephrase the context, question and possible answers while maintaining their semantics and structure.</p> <p>## Response Format: Each rephrasing should be separated by "<i>—REPHRASING—</i>" and follow this exact format:</p> <p><i>REPHRASING 1:</i> Context: [rephrased context] Question: [rephrased question] Answer A: [rephrased answer A] Answer B: [rephrased answer B] Answer C: [rephrased answer C] <i>—REPHRASING—</i> <i>REPHRASING 2:</i> [same format as above] And so on for all 5 rephrasings.</p> <p>## Original Item to Rephrase: Context: {context} Question: {question} Answer A: {answerA} Answer B: {answerB} Answer C: {answerC}</p>
LLM-as-a-Judge (FauxPas-EAI)	<p>Your job is to evaluate a predicted answer by comparing it against the gold answer and the given question. You may refer to the provided context if needed.</p> <p>## Grading Criteria: * 1: The predicted answer matches the gold answer or is a valid alternative (e.g., different but correct ways of writing a name). * 0: The predicted answer is wrong or does not align with the gold answer. * In some ambiguous cases, where it is unclear whether the predicted answer is correct or not, please refer to the provided context and use it as the final source for making your judgment.</p> <p>## Response Format: Please answer with <i>ONLY '1' or '0'</i>.</p> <p>## Here is your task: Context: {context} Question: {question} Gold Answer: {original_answer} Predicted Answer: {predicted_answer}</p>

Table 4: Prompts used for additional evaluation tasks. The rephrasing prompt was used to generate five rephrased versions of every item in the cleaned version of SocialQA, while the LLM-as-a-judge one to run automatic evaluation for FauxPas-EAI, instead of the originally used Levenshtein distance.