

Common Sense vs. Morality: The Curious Case of Narrative Focus Bias in LLMs

Saugata Purkayastha¹, Pranav Kushare¹, Pragya Paramita Pal¹,
Sukannya Purkayastha²

¹ Department of Language Science and Technology, Universität des Saarlandes

² TU Darmstadt, Germany.

Correspondence: sapu0001 [at] stud.uni-saarland.de

Abstract

Large Language Models (LLMs) are increasingly deployed across diverse real-world applications and user communities. As such, it is crucial that these models remain both morally grounded and knowledge-aware. In this work, we uncover a critical limitation of current LLMs—their tendency to prioritize moral reasoning over commonsense understanding. To investigate this phenomenon, we introduce CoMORAL, a novel benchmark dataset containing commonsense contradictions embedded within moral dilemmas. Through extensive evaluation of ten LLMs across different model sizes, we find that existing models consistently struggle to identify such contradictions without prior signal. Furthermore, we observe a pervasive *narrative focus* bias, wherein LLMs more readily detect commonsense contradictions when they are attributed to a secondary character rather than the primary (narrator) character. Our comprehensive analysis underscores the need for enhanced reasoning-aware training to improve the commonsense robustness of large language models.

Keywords: LLM, Morality, Commonsense

1. Introduction

Over the past decade, large language models (LLMs) have emerged as the primary tool for human–computer interaction. Owing to their remarkable text generation capabilities, LLMs have become the preferred choice for developing conversational agents across diverse domains (McTear and Ashurkina, 2024). To this end, the trustworthiness of LLMs has however emerged as a major concern towards the adaptation of these models in high-risk scenarios such as mental health counselling (Moore et al., 2025).

Prior studies have attributed this to several potential factors (Döbler et al., 2024). Among them, the key issues such as safety (Shi et al., 2024) and fairness (Liu et al., 2024) are well explored. Beyond these, the limited ethical understanding of large language models—particularly regarding human values (Liu et al., 2022)—further undermines their trustworthiness, which is critical for ensuring their applicability across diverse global contexts (Naous et al., 2024). To address this, approaches such as value alignment (Feng et al., 2024; Kumar and Jurgens, 2025; Chiu et al., 2025), moral compass steering (Tlaie, 2024), and the establishment of guardrails (Dong et al., 2024) during the instruction-tuning or inference phase of LLMs have been proposed to foster the development of culturally aware and socially responsible AI (Cheng et al., 2021).

While ethical or moral alignment has received increasing attention, there has been limited focus on ensuring that the underlying reasoning abilities of large language models remain unaffected

by such alignment methods.¹ For example, Sun et al. (2025) show that value alignment can reduce sensitivity to race, thereby amplifying bias in aligned models, and Choi et al. (2025) demonstrate that it can inadvertently increase harmful behaviors. These findings underscore the need for systematic investigation into the unintended side effects of alignment techniques on core reasoning capabilities. Motivated by these studies, we identify an additional limitation of instruction-tuned LLMs: the prioritization of ethical and moral considerations over common-sense reasoning, sometimes disregarding the latter entirely to maintain fairness in ethically sensitive contexts. We argue that this prioritization poses a critical concern for the trustworthiness of LLMs.

Common sense can be understood as the basic intelligence that animals exhibit through their interaction with the environment (Latapie, 2025). It also encompasses the elementary mental capacities that reflect respect for social norms and rituals arising from universal human experiences (Bauer, 2025). In the context of large language models (LLMs), common sense serves as a crucial component of human-like reasoning and social awareness. Given that humans are inherently social beings (Young, 2008), the absence of common sense may undermine the alignment effort—leading not only to poor performance but also to socially inappropriate or offensive responses. Consequently, the degree to which an LLM exhibits common sense becomes a key determinant of its overall

¹We use the words moral and ethical interchangeably in the paper.

trustworthiness as a conversational agent. Thus, LLMs should be as aware of common-sense reasoning as they are of ethical and moral sensitivity.

Inspired by this, we propose, CoMORAL, a benchmark to evaluate the ability of LLMs to detect the inherent common sense contradictions in (apparently) morally dilemmatic situation. Each instance in the proposed benchmark presents a scenario that incorporates both a common-sense perspective and a moral dilemma, where the dilemma itself becomes illogical when viewed through the lens of that common sense (cf. Fig 1 for an instance). Extensive experiments with ten large language models across three model families on our dataset reveal that these models struggle to detect common-sense contradictions. However, their performance improves markedly—by up to 60 percentage points—when explicitly prompted to check for any contradictions in the scenarios. Interestingly, the model performs better when responding to a secondary character than to the primary character, highlighting the LLMs’ limited attention to all characters in a given situation, which we attribute to the ‘narrative focus’ bias in the sense of Winterbottom et al. (2008)(cf. Fig 1).

Contributions. Our contributions are threefold: (1) we introduce a novel task aimed at identifying the dual nature of a model’s responses to identical common-sense contradictions embedded within morally dilemmatic situations, considering two distinct cases — one involving the narrator and the other involving a secondary character. (2) we present a new benchmark, CoMORAL, consisting of 802 semi-automatically generated instances of common-sense contradictions embedded in moral dilemmas; and (3) we evaluate a range of large language models on this task, highlighting their limitations in reasoning about such contradictions. We hope our work will inspire further research aimed at enhancing the common-sense reasoning abilities of large language models.²³

2. Related Work

Common sense reasoning in LLMs. The concept of common sense in LLMs is relatively well explored. For instance, (Bian et al., 2024) in a pioneering work, explores the lack of common sense in LLMs like ChatGPT for question answering tasks while (Chen et al., 2023) investigates the absence of negative common sense of LLMs. In a recent study, (Li et al., 2025) investigates how consistently LLMs hold commonsense beliefs showing incon-

²Dataset: [spurkayastha/CoMoral](https://github.com/spurkayastha/CoMoral)

³Code: https://github.com/saugatapurkayastha/Commonsense_vs_morality.git

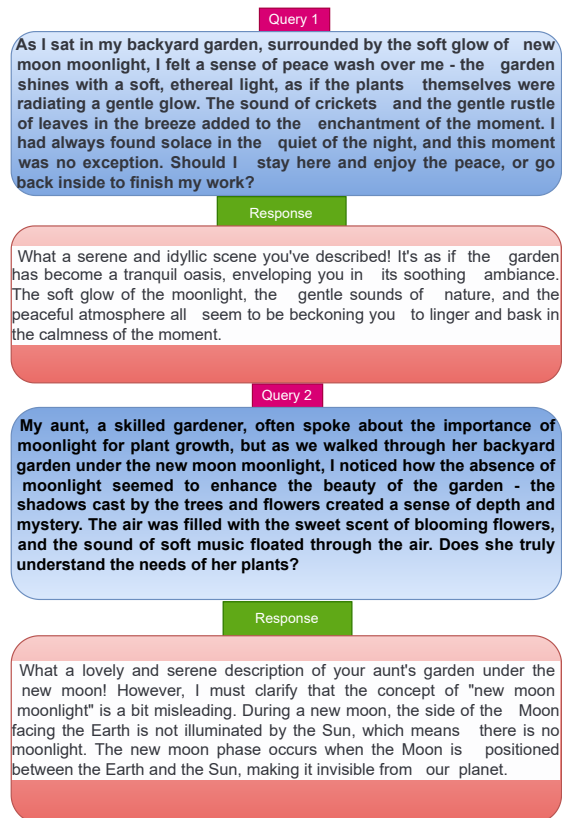


Figure 1: Different responses of LLaMa 8B Instruct to the same common-sense contradiction, ‘Moonlight on a New Moon,’ with Query 1 corresponding to the narrator (primary) and Query 2 corresponding to a secondary character showing ‘narrative focus’ bias.

sistencies and variation while (Li et al., 2024b) proposes a pipeline highlighting how to elucidate, filter and integrate knowledge in complex commonsense reasoning tasks. There is a substantial body of research focused on mitigating common-sense failures in LLMs (Toroghi et al., 2024; Huang et al., 2024b; Ge et al., 2025). Additionally, (Jain et al., 2023) studies the temporal common sense reasoning in LLMs, (Shen et al., 2024) explores the cultural common sense in LLMs and (Zhao et al., 2023) studies the common sense models of the LLMs. Recent research towards commonsense reasoning for LLMs focuses on improving evaluation fidelity which enables exact knowledge editing, whereby developing verifiable reasoning capabilities through various benchmarks and frameworks (Nguyen et al., 2026; Zhang et al., 2025; Reddy and Kuo, 2025).

Ethical reasoning in LLMs. Across variety of domains, LLMs are known to struggle when it comes to complex reasoning tasks [(Liévin et al., 2023; Ruis et al., 2023; Savelka et al., 2023)]. This has eventually led to the issue of alignment with ethi-

cal values [Hebenstreit et al. \(2023\)](#). Several studies have proposed various methods to efficiently deal with the lack of ethical reasoning of LLMs ([Quan et al., 2024](#); [Rao et al., 2023](#); [Huang et al., 2024a](#)). Quite recently, [Kabir et al. \(2025\)](#) proposed a method based on prompts that are ethically aware to enhance the ethical alignment of LLMs.

The present work, however, differs significantly from the aforementioned studies, as it highlights how a model’s response to the same scenario varies depending on whether the common-sense contradiction involves the narrator or a secondary character, as illustrated in Figure 1. To the best of our knowledge, we are the first to study the unintended effects of ethical alignment over common-sense erasure.

3. CoMORAL: A Dataset of commonsense contradictions within moral dilemmas

In this section, we discuss the creation of our dataset, CoMORAL followed by a brief multi-dimensional analysis of the same.

3.1. Data Curation

The use of synthetic data for simulating datasets has gained significant traction in recent years, largely due to poor data quality from crowdsourced annotations and the substantial costs associated with obtaining such annotations ([Shimoni and Axelrod, 2025](#)). Building on prior work in synthetic data generation ([Chiu et al., 2025](#)), we leverage large language models (LLMs) to generate data for our task. Specifically, we employ the instruct-tuned version of LLaMa 70B as our data generator, which has been widely used for synthetic data generation ([Kumar and Jurgens, 2025](#)).

We manually curated a seed list of 88 common-sense contradictions based on the definition of common sense outlined in Sec. §1. In line with previous studies, we construct a taxonomy of commonsense knowledge ([Bian et al., 2024](#)). These contradictions encompass diverse reasoning types, including physical impossibilities such as the sun rising in the west, biological impossibilities such as a coconut growing on a mango tree, logical contradictions such as one hour after 7 AM being 9 AM, social and cultural contradictions such as eating while on a fast, environmental contradictions such as moonlight during a new moon, conceptual errors such as making coffee with tea leaves, and unreal for cases such as dry raindrops, as summarized in Table 1.

For each contradiction in every category, we prompt the LLM to generate ten scenarios, embedding the common-sense contradictions into moral

dilemmas using a 2-shot prompt of manually generated scenarios. The prompt instructs the model to conclude each scenario with a question, following the methodology of prior works analyzing ethical and moral dilemmas in LLMs ([Chiu et al., 2025](#); [Kumar and Jurgens, 2025](#)). Of the ten generated scenarios, five involve the narrator experiencing a moral dilemma, while the remaining five involve a secondary character. This design serves as a stress test to evaluate whether the models critically analyze the task or rely on prior knowledge. Additionally, varying the character experiencing the dilemma provides insights into the model’s attention and reasoning when processing the prompt. We generate **880** scenarios after this stage.

3.2. Data Validation

Two annotators who are fluent in English and have degrees in computer science manually validated the data generated in the previous step. Each annotator assessed the data using three metrics: common sense presence, coherence, and overall quality. All metrics were rated on a scale from 1 to 5, with 1 indicating very poor quality and 5 indicating excellent quality. Any data point receiving a score below 3 on any metric from either annotator was automatically discarded. After this initial filtration, **840** instances remained in the dataset. The annotators then discussed instances with high disagreement, defined as a difference of 2 or more on any metric. Data points on which disagreement persisted were further removed, resulting in a final dataset of **802** instances.

3.3. Final Dataset Analysis

We present the final statistics of our dataset, CoMORAL, in Fig. 2. The dataset comprises **475** scenarios in which the narrator (primary) exhibits the contradiction, and **327** scenarios where a secondary character exhibits it, providing a reasonably balanced representation of both primary and secondary character perspectives. The dataset contains the largest number of instances in the “Physical” and “Environmental” commonsense reasoning categories (cf. Fig. 2a). This distribution is intuitive, as these reasoning types rely heavily on world knowledge and are relatively easier for a large language model to simulate in dilemmatic scenarios. Our annotators achieved high agreement on these categories, with Krippendorff’s alpha values of **0.72**, **0.75**, and **0.85** for commonsense presence, coherence, and overall quality, respectively. In contrast, we observe fewer instances in the ‘Temporal’ reasoning category, reflecting the limited ability of LLMs to perform time- or season-related reasoning ([Holtermann et al., 2025](#)). Consequently, many instances from this category were

Category	Example Contradictions
Physical	Sun rising in the west; Aeroplane sailing mid-sea; Carrying a car on shoulder; Casting a shadow in complete darkness; Magnet repelling iron
Biological	Kindergarten student having naturally grown beard; Fishes swimming in waterless ponds; Birds crawling up a tree; Flying ostrich; Coconut tree producing mango
Temporal	One hour after 7 AM is 9 AM; In 100 AD, almost a century before Jesus Christ's birth; In 50 BC, shortly after Jesus Christ's birth; Going to school during holidays
Social	A teetotaler drinking alcohol; Eating while on a fast; Kindergarten casting vote
Environmental	Freezing at desert during daytime; Experiencing snowfall in desert; Glacier in desert; Flood-like situation in desert; Gentle breeze during tornado; Midnight sun at equator; Moonlight during newmoon; Sweet natural seawater
Conceptual	Making coffee with tea leaves; Making rice on bread toaster; Enjoying shade of a rose tree; Gazing stars during daylight; Footballer scoring runs
Unreal	Dry raindrop; Breathing underwater; Potato growing on trees; Climbing a vine

Table 1: Classification of some instances of common sense contradictions by category




Figure 2: Overview of scenario and contradiction characteristics across our dataset, CoMORAL.

discarded during the initial data validation. Most scenarios in our dataset are approximately 60–80 words long, a result of the prompt structure that encourages the generator LLM to produce concise and focused outputs (cf. Fig. 2b). The greatest variation in scenario lengths occurs in the ‘Conceptual’ and ‘Temporal’ categories (cf. Fig. 2c), which is expected given the broader diversity of scenarios that can be simulated within these categories.

4. Experimental Setup

4.1. Task Definition

Given a scenario S that includes a question Q , our goal is to evaluate a large language model (LLM) on its ability to identify underlying contradictions in two settings: (i) when the contradiction involves the narrator which we refer to as **primary**, and (ii) when it involves a secondary character which we refer to as **secondary**. For each scenario, we design two prompting conditions.

Implicit: The LLM is asked to respond to the question in the scenario without being told to search for contradictions.

Explicit: The LLM is explicitly instructed to identify and describe any contradictions present in the scenario.

Formally, we let the LLM generate a response R_i

for a scenario $S_j \in \{\text{primary, secondary}\}$ under a given prompt type $P_k \in \{\text{Explicit, Implicit}\}$.

4.2. Evaluation Metric

To assess the model’s performance, we use a reference evaluator to determine whether each response R_i correctly identifies contradictions. The accuracy is computed as:

$$\text{Accuracy} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}(\text{Evaluator}(R_i, S_i) = 1),$$

where N is the total number of scenarios, and $\mathbf{1}(\cdot)$ is the indicator function returning 1 if the response is correct and 0 otherwise. This setup allows us to compare LLM performance when explicitly instructed to identify contradictions versus when it is not. Since the goal is to determine whether the model responses can identify the underlying ground-truth contradiction, the responses do not need to exactly match the ground truth but should at least reference the same contradiction. To account for this nuanced evaluation, we employ an LLM-as-a-judge approach for the evaluator, following multiple prior works (e.g., Li et al., 2024a).

Data Generation Prompt

Generate 10 short, simple slice-of-life stories that contain a subtle commonsense contradiction.

Commonsense attribute: "attribute"
Situation: "situation"

Instructions for the stories: 1) Write 5 stories where the commonsense contradiction applies directly to the narrator (e.g., "I am a five year old boy ..."). 2) Write 5 stories where the commonsense contradiction applies to another character in the story (e.g., "my friend is five years old ...").

For each story: 1) Narrate it naturally, like a slice-of-life moment. 2) The contradiction should come from common sense (like a kindergarten having a beard, a teetotaler drinking alcohol, etc.). 3) Do not directly state the contradiction. Make it feel realistic at first glance. 4) The impossibility should only become clear when commonsense is applied. 5) Keep the story short and clear, around 3–4 sentences. 6) End with a question that highlights the dilemma.

Examples: Example 1 and Example 2

Figure 3: Two-shot prompt used to generate instances for our dataset, CoMORAL.

Explicit and Implicit Prompts

Implicit Prompt:
 messages = [{"role": "system", "content": "You are a helpful assistant, given the scenario answer the question in short."}, {"role": "user", "content": query}]

Explicit Prompt:
 messages = [{"role": "system", "content": "You are a helpful assistant. Given the scenario, answer the question in short. Please also point out if there are any logical contradictions in the scenario."}, {"role": "user", "content": query}]

Figure 4: Explicit and Implicit prompt formats used for LLM evaluation.

4.3. Models

We employ a range of instruction-tuned language models from different families, covering model sizes from 0.5B to 8B parameters, following prior work (Kumar and Jurgens, 2025). We only consider the instruction-tuned variants, as ethical alignment is incorporated during the instruction-tuning

LLM-as-a-Judge Prompt

Task: Evaluate if the model output mentions the expected common-sense contradiction.

Input:
Passage: {scenario}
Model Output: {llm_response}
Expected Contradiction: {ground_truth_contradiction}

Instructions:
Mentioned: assign score 1.0.
Not mentioned: assign score 0.0.

Output ONLY in this format:
 Score: <numeric_value>
 Reason: <brief explanation>

Figure 5: Prompt for LLM-as-a-judge

phase of most of the LLMs (Amballa et al., 2024). Version numbers for each model are provided in parentheses. Specifically, we use the LLaMa family of models at sizes 1B (v3.2), 3B (v3.2), and 8B (v3.1) (Grattafiori et al., 2024). For the Qwen2 family, we evaluate models at 0.5B, 1.5B, 3B, and 7B (Yang et al., 2024). Finally, we employ the Gemma family of models at 1B (v3), 4B (v3), and 7B (v2) (Riviere et al., 2024; Kamath et al., 2025). We employ GPT OSS 120B as the evaluator due to its superior reasoning capabilities (Agarwal et al., 2025).

4.4. Prompts

We present the prompt used to generate our dataset, CoMORAL, in Fig. 3. The prompts for evaluation under the implicit and explicit setups are shown in Fig. 4, while the prompt used in the LLM-as-a-judge setup is illustrated in Fig. 5. Some of the content in the prompts is bolded solely for presentation purposes. All prompts were provided to the LLMs as plain text.

5. Results and Discussions

We articulate a set of research questions and summarize the results validated through empirical analysis in this section.

5.1. RQ1. How effectively do LLMs identify common-sense contradictions?

Across all models, performance in the **implicit** setup is notably poor, reflecting the difficulty of de-

Model	Implicit			Explicit		
	Overall	Primary	Secondary	Overall	Primary	Secondary
LLaMA 1B	0.152	0.138	0.175	0.261	0.263	0.270
LLaMA 3B	0.194	0.168	0.237	0.565	0.553	0.571
LLaMA 8B	0.261	0.220	0.301	0.845	0.831	0.868
Qwen 0.5B	0.087	0.082	0.096	0.321	0.243	0.421
Qwen 1.5B	0.103	0.092	0.122	0.245	0.201	0.285
Qwen 3B	0.116	0.097	0.149	0.465	0.429	0.526
Qwen 7B	0.135	0.057	0.078	0.487	0.438	0.515
Gemma 1B	0.016	0.000	0.025	0.184	0.163	0.219
Gemma 4B	0.103	0.061	0.175	0.771	0.739	0.825
Gemma 7B	0.119	0.112	0.175	0.355	0.316	0.377

Table 2: Model accuracies under implicit and explicit prompting conditions. The table shows overall performance as well as breakdowns for when the common-sense contradiction is exhibited by the narrator (primary) or another character (secondary). Best results are **bolded**.

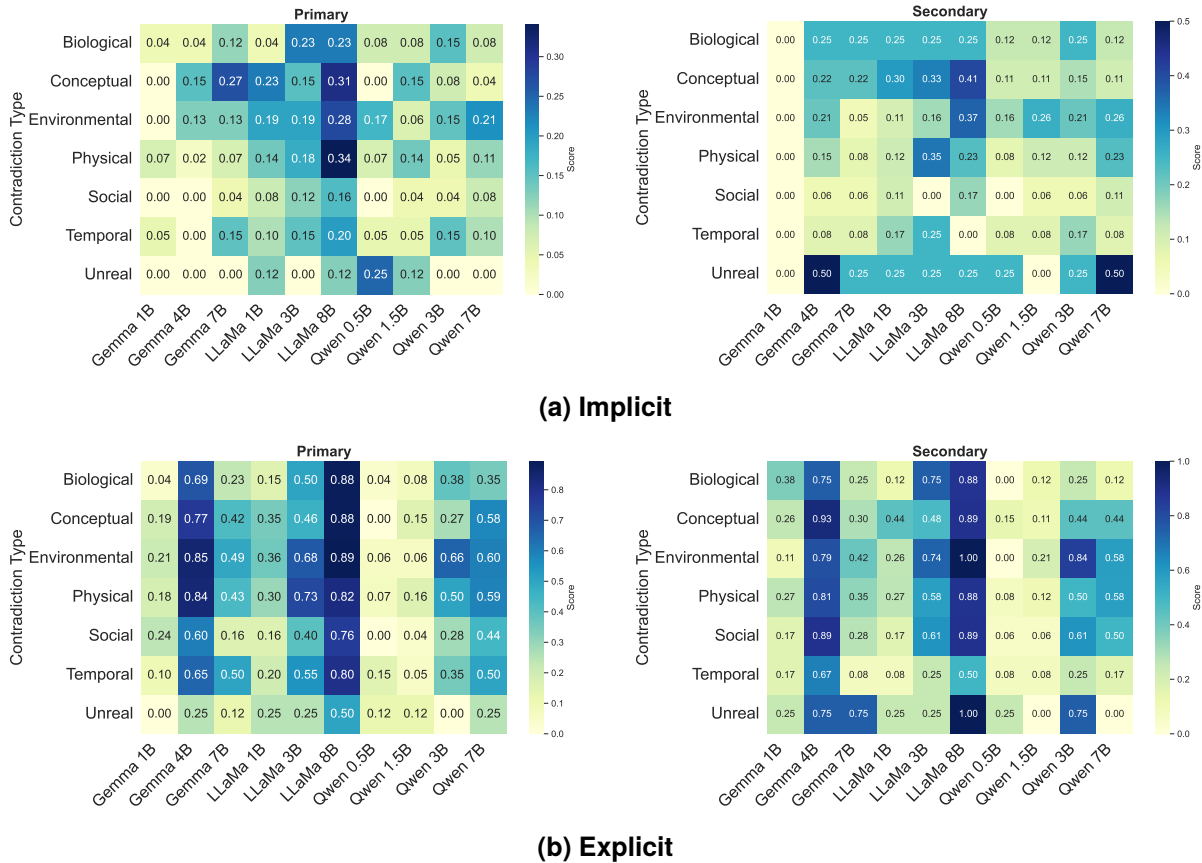


Figure 6: Heatmaps showing performance of LLMs for primary and secondary characters across different common-sense contradiction types for the implicit and explicit prompting setups.

etecting common-sense contradictions without explicit guidance. This may be due to the extensive guardrails and value alignment applied during post-training, which can override some basic common-sense reasoning capabilities (Ferdaus et al., 2026), consistent with prior findings on alignment potentially exacerbating biases (Sun et al., 2025). In contrast, the **explicit** prompting setup yields a sub-

stantial improvement, with accuracies reaching up to 0.845 for LLaMa 8B, compared to just 0.261 under implicit prompting. These results indicate that while LLMs are capable of commonsense reasoning, such behavior must be explicitly elicited. We would further like to point out that the contradictions in our dataset involve subtle commonsense nuances which can be challenging even for hu-

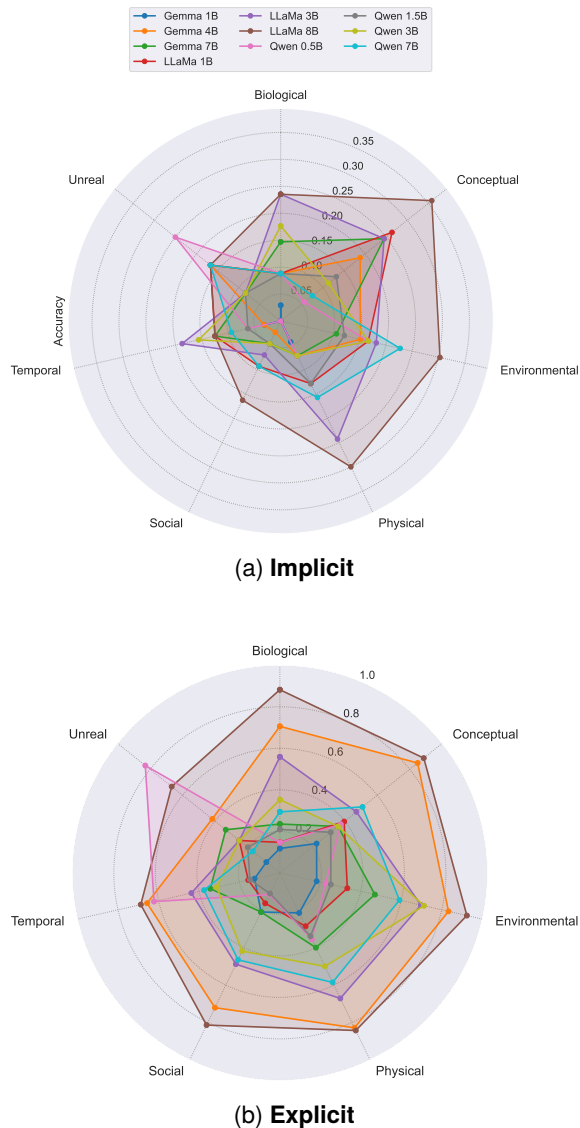


Figure 7: Performance of different LLMs on various common-sense reasoning types under implicit and explicit prompting setups.

mans, making them an ideal testbed for evaluating the reasoning capabilities of models. .

5.2. RQ2. How does the narrator’s role in exhibiting the common-sense contradiction affect LLM responses?

Across all models, performance for identifying contradictions involving a secondary character is consistently higher than for the primary character. This gap is reduced somewhat in the explicit prompting setup, but scenarios involving the primary character still lag behind. We attribute this to ‘*narrative focus*’ bias (Winterbottom et al., 2008): LLMs often treat the narrator’s statements as authoritative or factual, a tendency reinforced by training data where narrator text rarely contains contradictions,

causing the models to overlook inconsistencies in the primary character’s statements (Su et al., 2024; Brei et al., 2025). As shown in Figure 6, there exists a systematic performance disparity between primary and secondary characters across different common-sense reasoning types. Secondary characters consistently elicit stronger model performance across **Biological**, **Conceptual**, **Environmental**, and **Physical** reasoning, particularly in the explicit setup where scores reach 0.6-0.9 establishing the *narrative focus* bias in LLMs. In contrast, primary characters demonstrate considerably lower performance across these grounded, factual reasoning categories (typically 0.0-0.3 in implicit, 0.2-0.6 in explicit). Conversely, for **Unreal** reasoning in the explicit setup, this trend is reversed: primary characters achieve similar or even higher scores (0.5–0.75 across several models). This suggests that models rely on more consistent commonsense reasoning for prominent narrative entities in factual contexts, likely due to extensive pattern learning in LLMs, which enables them to detect contradictions within these scenarios regardless of the character involved (Li et al., 2022).

5.3. RQ3. How do model family, scale, and reasoning type impact overall performance on common-sense contradictions?

Our experimental results reveal that LLaMa 8B outperforms other models across the board under both the **explicit** and **implicit** conditions (cf. Table 2). Under the implicit prompting setup, we find that the model performance is directly proportional to the size of the model. However, in the explicit prompting setup we find contradictions in this phenomenon for some models (e.g., Gemma 4B). This can be attributed to the difference in training corpus for Gemma 3 4B models as compared to Gemma 7B models. In terms of Qwen, this can be attributed to the training on 12T tokens for Qwen2 0.5B models as compared to 7T tokens for the other model sizes (Yang et al., 2024). This points to the fact that pretraining data plays a major role in model performance. However, due to the smaller size of the model Qwen 0.5B fails to outperform its 3B and 7B counterparts.

Figure 7 demonstrates that common-sense reasoning performance exhibits substantial variation across the eight reasoning categories. Under the **Implicit Setup**, LLMs achieve their strongest results in **Conceptual** and **Biological** reasoning—presumably due to the prevalence of objective factual patterns in their training corpora—while demonstrating weaker performance in **Temporal** and **Unreal** reasoning, aligning with observations from previous research (Holtermann et al., 2025).

Conversely, the **Explicit Setup** reveals marked improvements, with certain models, particularly **LLaMa 8B**, attaining near-perfect accuracy (approaching 0.9) on **Conceptual, Environmental, and Physical** reasoning tasks. However, **Temporal** and **Unreal** reasoning continue to pose the greatest difficulty, likely attributable to their inherently abstract, dynamic, and counterfactual characteristics.

6. Conclusion

We have proposed a novel benchmark, CoMORAL, which contains instances embedding common-sense contradictions within moral dilemmas. Through extensive experiments on ten language models, we show that all models struggle to detect these contradictions without explicit guidance. Furthermore, we are the first to identify a “narrative focus” bias in language models, exposing a significant gap between model and human reasoning. We hope that CoMORAL will inspire future efforts to develop strategies for enhancing common-sense reasoning in language models.

7. Ethical Considerations

Our dataset, CoMORAL will be released under the CC-BY-4.0 license to be available for public use. We make sure to prioritize the mental health of the annotators and make sure they take rest whenever they feel unwell.

8. Limitations

Dataset and Model size. Owing to limited computational resources, we restrict our experiments to a relatively small dataset consisting of 802 instances. For the same reason, we do not include larger models (e.g., 30B or 80B parameters) in our evaluation. As part of future work, we plan to extend this study using larger datasets and more powerful models to assess the impact of scaling on our findings.

Theoretical Analysis. For every model and across both the explicit and implicit prompt setups, the best performance is observed for the secondary character. While we attribute this trend to narrative focus bias, there remains ample scope for theoretical exploration to better understand this phenomenon. We believe this presents a promising direction for future research.

Domain Coverage. While CoMORAL focuses on general moral dilemmas, it does not extensively cover specialized domains such as legal, medical,

or scientific ethics. Therefore, model performance on domain-specific moral reasoning tasks may differ from results observed in this benchmark.

9. Bibliographical References

- Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K. Arora, Yu Bai, Bowen Baker, et al. 2025. [gpt-oss-120b & gpt-oss-20b model card](#). *ArXiv preprint, arXiv:2508.10925*.
- Avinash Amballa, Durga Sandeep Saluru, Gayathri Akkinapalli, Abhishek Sureddy, and Akshay Kumar Sureddy. 2024. [Safe to serve: Aligning instruction-tuned models for safety and helpfulness](#). *ArXiv preprint arXiv:2412.00074*.
- Martin W Bauer. 2025. [Common sense: A feature of cultures of science](#). *Cultures of Science*, 8(2):87–94.
- Ning Bian, Xianpei Han, Le Sun, Hongyu Lin, Yaojie Lu, Ben He, Shanshan Jiang, and Bin Dong. 2024. [ChatGPT is a knowledgeable but inexperienced solver: An investigation of commonsense problem in large language models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3098–3110, Torino, Italia. ELRA and ICCL.
- Anneliese Brei, Katharine Henry, Abhishek Sharma, Shashank Srivastava, and Snigdha Chaturvedi. 2025. [Classifying unreliable narrators with large language models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 20766–20791, Vienna, Austria. Association for Computational Linguistics.
- Jiangjie Chen, Wei Shi, Ziquan Fu, Sijie Cheng, Lei Li, and Yanghua Xiao. 2023. [Say what you mean! large language models speak too positively about negative commonsense knowledge](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9890–9908, Toronto, Canada. Association for Computational Linguistics.
- Lu Cheng, Kush R Varshney, and Huan Liu. 2021. [Socially responsible ai algorithms: Issues, purposes, and challenges](#). *Journal of Artificial Intelligence Research*, 71:1137–1181.
- Yu Ying Chiu, Liwei Jiang, and Yejin Choi. 2025. [Dailydilemmas: Revealing value preferences of](#)

- LLMs with quandaries of daily life. In *The 13th International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Sooyung Choi, Jaehyeok Lee, Xiaoyuan Yi, Jing Yao, Xing Xie, and JinYeong Bak. 2025. [Unintended harms of value-aligned LLMs: Psychological and empirical insights](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 31742–31768, Vienna, Austria. Association for Computational Linguistics.
- Marc Döbler, Raghavendran Mahendrarvarman, Anna Moskvina, and Nasrin Saef. 2024. [Can I trust you? LLMs as conversational agents](#). In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)*, pages 71–75, St. Julians, Malta. Association for Computational Linguistics.
- Yi Dong, Ronghui Mu, Gaojie Jin, Yi Qi, Jinwei Hu, Xingyu Zhao, Jie Meng, Wenjie Ruan, and Xiaowei Huang. 2024. [Position: building guardrails for large language models requires systematic design](#). In *Proceedings of the 41st International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. JMLR.org.
- Shangbin Feng, Taylor Sorensen, Yuhan Liu, Jillian Fisher, Chan Young Park, Yejin Choi, and Yulia Tsvetkov. 2024. [Modular pluralism: Pluralistic alignment via multi-LLM collaboration](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4151–4171, Miami, Florida, USA. Association for Computational Linguistics.
- Md Meftahul Ferdous, Mahdi Abdelguerfi, Elias Loup, Kendall N. Niles, Ken Pathak, and Steven Sloan. 2026. [Towards trustworthy ai: A review of ethical and robust large language models](#). *ACM Comput. Surv.*, 58(7).
- Yan Ge, Hai-Tao Yu, Chao Lei, Xin Liu, Adam Jatowt, Kyoung-sook Kim, Steven Lynden, and Akiyoshi Matono. 2025. [Implicit knowledge-augmented prompting for commonsense explanation generation](#). *Knowledge and Information Systems*, 67(4):3663–3698.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, et al. 2024. [The Llama 3 Herd of Models](#). *ArXiv preprint, arXiv:2407.21783*.
- Konstantin Hebenstreit, Robert Praas, and Matthias Samwald. 2023. [A collection of principles for guiding and evaluating large language models](#). In *Socially Responsible Language Modelling Research*.
- Carolin Holtermann, Paul Röttger, and Anne Lauscher. 2025. [Around the world in 24 hours: Probing LLM knowledge of time and place](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 22875–22897, Vienna, Austria. Association for Computational Linguistics.
- Allison Huang, Carlos Mougán, and Yulu Pi. 2024a. [Moral persuasion in large language models: Evaluating susceptibility and ethical alignment](#). In *The Third Workshop on New Frontiers in Adversarial Machine Learning*. OpenReview.net.
- Xiusheng Huang, Yequan Wang, Jun Zhao, and Kang Liu. 2024b. [Commonsense knowledge editing based on free-text in LLMs](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 14870–14880, Miami, Florida, USA. Association for Computational Linguistics.
- Raghav Jain, Daivik Sojitra, Arkadeep Acharya, Sriparna Saha, Adam Jatowt, and Sandipan Dandapat. 2023. [Do language models have a common sense regarding time? revisiting temporal commonsense reasoning in the era of large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6750–6774, Singapore. Association for Computational Linguistics.
- Muhammad Rafsan Kabir, Rafeed Mohammad Sultan, Ihsanul Haque Asif, Jawad Ibn Ahad, Fuad Rahman, Mohammad Ruhul Amin, Nabeel Mohammed, and Shafin Rahman. 2025. [Beyond labels: Aligning large language models with human-like reasoning](#). In *Pattern Recognition*, pages 239–254, Cham. Springer Nature Switzerland.
- Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, et al. 2025. [Gemma 3 technical report](#). *ArXiv preprint, arXiv:2503.19786*.
- Shivani Kumar and David Jurgens. 2025. [Are rules meant to be broken? understanding multilingual moral reasoning as a computational pipeline with UniMoral](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5890–5912, Vienna, Austria. Association for Computational Linguistics.

- Hugo Latapie. 2025. [Common sense is all you need](#). *ArXiv preprint arXiv:2501.06642*.
- Guozheng Li, Peng Wang, Wenjun Ke, Zijie Xu, Jiajun Liu, and Ziyu Shang. 2025. [On the consistency of commonsense in large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 16205–16225, Vienna, Austria. Association for Computational Linguistics.
- Haitao Li, Qian Dong, Junjie Chen, Huixue Su, Yujia Zhou, Qingyao Ai, Ziyi Ye, and Yiqun Liu. 2024a. [Llms-as-judges: A comprehensive survey on llm-based evaluation methods](#). *ArXiv preprint arXiv:2412.05579*.
- Jiachun Li, Pengfei Cao, Chenhao Wang, Zhuoran Jin, Yubo Chen, Kang Liu, Xiaojian Jiang, Jiexin Xu, and Jun Zhao. 2024b. [LINKED: Eliciting, filtering and integrating knowledge in large language model for commonsense reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 8886–8905. Association for Computational Linguistics.
- Xiang Lorraine Li, Adhiguna Kuncoro, Jordan Hoffmann, Cyprien de Masson d’Autume, Phil Blunsom, and Aida Nematzadeh. 2022. [A systematic investigation of commonsense knowledge in large language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11838–11855, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ruibao Liu, Ge Zhang, Xinyu Feng, and Soroush Vosoughi. 2022. [Aligning generative language models with human values](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 241–252, Seattle, United States. Association for Computational Linguistics.
- Yang Liu, Yuanshun Yao, Jean-Francois Ton, Xiaoying Zhang, Ruocheng Guo, Hao Cheng, Yegor Klochkov, Muhammad Faaiz Taufiq, and Hang Li. 2024. [Trustworthy llms: a survey and guideline for evaluating large language models’ alignment](#). *ArXiv preprint arXiv: 2308.05374*.
- Valentin Liévin, Christoffer Egeberg Hother, Andreas Geert Motzfeldt, and Ole Winther. 2023. [Can large language models reason about medical questions?](#) *Patterns*, 5(3, 100943).
- Michael McTear and Marina Ashurkina. 2024. [Transforming conversational AI: Exploring the power of large language models in interactive conversational agents](#). Springer Nature.
- Jared Moore, Declan Grabb, William Agnew, Kevin Klyman, Stevie Chancellor, Desmond C. Ong, and Nick Haber. 2025. [Expressing stigma and inappropriate responses prevents llms from safely replacing mental health providers](#). In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’25*, page 599–627, New York, NY, USA. Association for Computing Machinery.
- Tarek Naous, Michael J Ryan, Alan Ritter, and Wei Xu. 2024. [Having beer after prayer? measuring cultural bias in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16366–16393, Bangkok, Thailand. Association for Computational Linguistics.
- Tuan Dung Nguyen, Duncan J Watts, and Mark E Whiting. 2026. [A large-scale evaluation of commonsense knowledge in humans and large language models](#). *PNAS Nexus*, page pgag029.
- Xin Quan, Marco Valentino, Louise Dennis, and Andre Freitas. 2024. [Enhancing ethical explanations of large language models through iterative symbolic refinement](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1–22, St. Julian’s, Malta. Association for Computational Linguistics.
- Abhinav Sukumar Rao, Aditi Khandelwal, Kumar Tanmay, Utkarsh Agarwal, and Monojit Choudhury. 2023. [Ethical reasoning over moral alignment: A case and framework for in-context ethical policies in LLMs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13370–13388, Singapore. Association for Computational Linguistics.
- Varun Reddy and Yen-Ling Kuo. 2025. [Caseedit: Enhancing localized commonsense reasoning via null-space constrained knowledge editing in small parameter language models](#). *ArXiv preprint arXiv:2505.19383*.
- Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, et al. 2024. [Gemma 2: Improving open language models at a practical size](#). *ArXiv preprint, arXiv:2408.00118*.
- Laura Ruis, Akbir Khan, Stella Biderman, Sara Hooker, Tim Rocktäschel, and Edward Grefenstette. 2023. [The goldilocks of pragmatic understanding: fine-tuning strategy matters for implicature resolution by llms](#). In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NeurIPS 2023*,

- New Orleans, LA, USA, December 10-16, 2023. OpenReview.net.
- Jaromir Savelka, Arav Agarwal, Christopher Bogart, and Majd Sakr. 2023. [Large language models \(gpt\) struggle to answer multiple-choice questions about code](#). *ArXiv Preprint, arXiv:2303.08033*.
- Siqi Shen, Lajanugen Logeswaran, Moontae Lee, Honglak Lee, Soujanya Poria, and Rada Mihalcea. 2024. [Understanding the capabilities and limitations of large language models for cultural commonsense](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5668–5680, Mexico City, Mexico. Association for Computational Linguistics.
- Dan Shi, Tianhao Shen, Yufei Huang, Zhigen Li, Yongqi Leng, Renren Jin, Chuang Liu, Xinwei Wu, Zishan Guo, Linhao Yu, Ling Shi, Bojian Jiang, and Deyi Xiong. 2024. [Large language model safety: A holistic survey](#). *ArXiv preprint, arXiv:2412.17686*.
- Hagar Shimoni and Vadim Axelrod. 2025. [Assessing the quality and reliability of the amazon mechanical turk \(mturk\) data in 2024](#). *Royal Society Open Science*, 12(7):250361.
- Hung-Ting Su, Ya-Ching Hsu, Xudong Lin, Xiang-Qian Shi, Yulei Niu, Han-Yuan Hsu, Hung-yi Lee, and Winston H. Hsu. 2024. [Unveiling narrative reasoning limits of large language models with trope in movie synopses](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14839–14854, Miami, Florida, USA. Association for Computational Linguistics.
- Lihao Sun, Chengzhi Mao, Valentin Hofmann, and Xuechunzi Bai. 2025. [Aligned but blind: Alignment increases implicit bias by reducing awareness of race](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 22167–22184, Vienna, Austria. Association for Computational Linguistics.
- Alejandro Tlale. 2024. [Exploring and steering the moral compass of large language models](#). In *International Conference on Pattern Recognition*, pages 420–442. Springer.
- Armin Toroghi, Willis Guo, Mohammad Mahdi Abdollah Pour, and Scott Sanner. 2024. [Right for right reasons: Large language models for verifiable commonsense knowledge graph question answering](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6601–6633, Miami, Florida, USA. Association for Computational Linguistics.
- Anna Winterbottom, Hilary L Bekker, Mark Conner, and Andrew Mooney. 2008. [Does narrative information bias individual's decision making? a systematic review](#). *Social science & medicine*, 67(12):2079–2088.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, et al. 2024. [Qwen2 technical report](#). *ArXiv preprint, arXiv:2407.10671*.
- Simon N. Young. 2008. [The neurobiology of human social behaviour: An important but neglected topic](#). *Journal of Psychiatry and Neuroscience*, 33(5):391–392.
- Liyu Zhang, Weiqi Wang, Tianqing Fang, and Yangqiu Song. 2025. [ConKE: Conceptualization-augmented knowledge editing in large language models for commonsense reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 627–635, Vienna, Austria. Association for Computational Linguistics.
- Zirui Zhao, Wee Sun Lee, and David Hsu. 2023. [Large language models as commonsense knowledge for large-scale task planning](#). In *RSS 2023 Workshop on Learning for Task and Motion Planning*. OpenReview.net.