

Mary, the Cheeseburger-Eating Vegetarian: Do LLMs Recognize Incoherence in Narratives?

Karin de Langis¹, Püren Öncel², Ryan Peters¹, Andrew Elfenbein⁴,
Laura Kristen Allen², Andreas Schramm³, & Dongyeop Kang¹

¹Department of Computer Science and Engineering, University of Minnesota

²Department of Educational Psychology, University of Minnesota

³Department of Linguistics, Hamline University

⁴Department of English and Center for Cognitive Sciences, University of Minnesota
dento019@umn.edu

Abstract

Leveraging a dataset of paired narratives, we investigate the extent to which large language models (LLMs) can reliably separate incoherent and coherent stories. A probing study finds that LLMs’ internal representations can reliably identify incoherent events in narratives. However, this separation disappears by the narrative’s end, and weakens when the differences between coherent and incoherent stories are more subtle. When asked to rate overall coherence of narratives after reading, LLMs generate responses that fail to satisfactorily separate the coherent and incoherent narratives. Reasoning models tested do not eliminate these deficits, indicating that thought strings may not be able to fully address the discrepancy between model internal state and behavior. Additionally, we find that LLMs appear to be more sensitive to incoherence resulting from an event that violates the setting (e.g., a rainy day in the desert) than to incoherence arising from a character violating an established trait (e.g., Mary, a vegetarian, later orders a cheeseburger), suggesting that LLMs may rely more on prototypical world knowledge than building coherence through a meaning-based world model of the narrative setting. Together, our results indicate that LLMs lack robustness in their ability to recognize incoherence in narratives.¹

1 Introduction

Contemporary models of (human) reading comprehension characterize comprehension as a dynamic process in which the reader continually builds and updates representations to maintain coherence and integrate new information with prior knowledge, e.g., (Kintsch, 1998; Myers and O’Brien, 1998; Van den Broek et al., 1999). Specifically, narratives evoke mental simulations of events (Kintsch, 1998), and when processing narratives, readers focus on dimensions such as causality, time, space,

¹Our data is available at <https://github.com/minnesotanlp/narrative-coherence>.

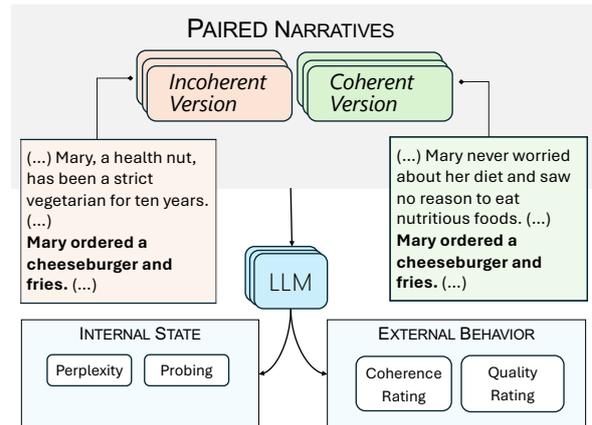


Figure 1: We use paired coherent and incoherent narratives to investigate the extent to which LLMs maintain coherence when tracking fictional entities. Internal measures differentiate coherent from incoherent stories at the incoherent location, yet LLMs’ explicit coherent ratings at the end of the story fail to do so.

motivations and protagonist to maintain coherence (Magliano et al., 1999; McNamara and Magliano, 2009; Zwaan and Radvansky, 1998).

For LLMs, claims of “narrative understanding” rely on whether an analogue of this coherence-monitoring process is present. Standard evaluations rarely isolate incoherence itself, confounding it with topic, difficulty, or style; as a result, they cannot adjudicate whether models register contradictions or rely on surface heuristics. This raises the question: *Are LLMs able to reliably recognize coherence in narratives?* To explore this question, we leverage a longstanding paradigm in human reading studies: paired narratives that are identical except for a critical inconsistency manipulation (Albrecht and O’Brien, 1993). Specifically, we adapt and augment materials from Rizzella and O’Brien (2024) to study LLMs, enabling precise comparison of LLM responses to coherent versus incoherent narratives (Figure 1).

We first examine *internal* LLM representations

of the narratives using token-level perplexity and a probing study of hidden states. Both measures indicate that LLMs can internally distinguish between the coherent and incoherent narratives, at least at the location where incoherence occurs (or does not occur). However, LLMs appear **more sensitive to event-setting incoherence** (violations of commonsense expectations about the scene, like being given chopsticks and then being served pizza) than to **trait-behavior incoherence** (violations of a character’s established personality or traits, like a vegetarian ordering a cheeseburger). Human readers show no such asymmetry, suggesting that LLMs may be relying more on background world knowledge to identify incoherence.

Secondly, we examine *external* LLM responses by asking LLMs to assess the coherence and quality of each story. Surprisingly, both instruction-tuned and reasoning models fail to reliably distinguish coherent from incoherent narratives on both Likert ratings and forced choice true or false questions. Examination of reasoning traces reveals that LLMs often apply muddled logic – for instance, by identifying innocuous details as incoherent – revealing **a discrepancy between their internal sensitivity to narrative structure and their explicit understanding of it**. This dissociation between internal and external indicators of LLM narrative understanding highlights a potential vulnerability: models may internally register indicators of narrative incoherence without being able to reliably recognize or report it when instructed to do so.

To further explore what might drive the dissociation between internal and external indicators of LLMs’ recognition of narrative coherence, we conduct an additional probing study to examine how well LLM internal representations separate *incoherent* from *unexpected* events. We find that probes are substantially more accurate for incoherent and coherent events versus incoherent and unexpected events, suggesting that LLM narrative understanding may be driven by statistically learned expectations of narratives rather than building a situation model (or world model) of the narrative setting.

As LLMs are increasingly deployed in real-world applications, and especially in educational settings focused on literature and creative writing, understanding how they construct or fail to construct narrative coherence is essential for evaluating their trustworthiness and alignment with human narrative reasoning. We believe that our findings, together with the paired narrative dataset in-

troduced here, provide a valuable foundation for future research on the mechanisms of narrative comprehension in LLMs.

2 Related Work

Our work intersects with previous research in comprehension processes in both humans and language models. We cover each of these separately in this section.

2.1 Comprehension processes in humans

Human reading comprehension involves constructing mental representations of the discourse content (Gernsbacher, 1997; Kintsch, 1998; Myers and O’Brien, 1998). As reader encounter new information, they integrate it with their prior knowledge (both prior discourse and general world knowledge). Because of the limited capacity of the reader’s working memory, comprehension proceeds in cycles in which only a small portion of the text is processed at a time (Kintsch and Van Dijk, 1978; Garrod and Sanford, 1981; Van Dijk et al., 1983). If this comprehension process is successful, the resulting representation forms a situation model, a network of semantic and causal relations that connects information from the discourse with reader’s prior knowledge.

The construction of a *coherent* situation model is the central goal of comprehension, achieved by integrating new information in both local and global contexts (Glenberg and Langston, 1992; O’Brien and Albrecht, 1992). Readers establish local coherence by integrating incoming information with the immediately preceding context (i.e., information in their short-term memory), and global coherence by linking incoming information to earlier, no-longer-active but relevant information (Albrecht and O’Brien, 1993; Glenberg and Langston, 1992; O’Brien and Albrecht, 1992). Researchers have long debated whether inferences support only local or also global coherence in comprehension (e.g., McKoon and Ratcliff (1992)). Albrecht and O’Brien (1993) found that the reader detects global incoherence in narratives immediately even when local coherence is maintained, indicating that despite capacity limitations in working memory, the reader maintains access to global situation models throughout the course of a narrative and continually integrates new information into this situation model during reading. This can be partially but not completely mitigated by qualifying informa-

tion (i.e., the inclusion of information that explains away the apparent incoherence). Further nuances have been explored in the literature, such as when past information that is no longer relevant conflicts (O’Brien et al., 1998), or when prospective information conflicts (Rizzella and O’Brien, 2024).

The present work uses the incoherent narratives studied by Rizzella and O’Brien (2024), which contributed to understanding how and when people integrate new information into mental representations of narratives as they read.

2.2 Coherence and comprehension in LLMs

Narrative comprehension in LLMs has been characterized as coherent, but brittle (Chang and Bergen, 2024), and there are debates as to the extent that robust situation modeling is possible under transformer-only architectures (Mahowald et al., 2024). One obvious challenge in situation modeling for LLMs is context window limitations: a model can only process a finite number of tokens, making it impossible to process long narratives such as those spanning multiple novels. While there are approaches that address context window limitations (Su et al., 2024), such approaches are computationally expensive. Efficient, human-like solutions to processing long narratives include hierarchical processing by incremental summarization (e.g., Song et al. (2024)). A second, more pernicious challenge is that LLMs have been found to make nonhuman-like mistakes in situation modeling, even when the context length poses no challenges. For example, generations can refer to non-existent discourse entities: “I doubt Michael baked a cake. It’s in the oven,” (Schuster and Linzen, 2022), and LLMs can produce highly incoherent narrative summaries (Goyal et al., 2022). Early evidence also suggests that LLMs have inconsistent abilities in making pragmatic inferences in texts and inferring relations between events (Chang and Bergen, 2024; de Langis et al., 2025).

Alongside studies analyzing LLM outputs to infer model knowledge and capabilities, a parallel branch of research answers questions about what LLMs truly understand by examining the information encoded in their internal representations. Probing classifiers are one such approach, revealing what knowledge is linearly accessible within networks (Alain and Bengio, 2018; Belinkov, 2022). One key question in this line of research is to what extent transformers trained on sequential prediction tasks develop coherent internal world models of the

processes generating their input sequences. Li et al. (2023) showed that a GPT variant trained to predict legal moves in the Othello game developed an emergent, and linear (Nanda et al., 2023), internal representation of the board state, despite never being explicitly given this information. Extending this to discourse, entity tracking studies find that maintaining entity states during narratives is fundamental to comprehension, yet some LLMs struggle with this (Kim and Schuster, 2023). Notably, only models pretrained on both text and code demonstrated robust entity tracking, suggesting that text-only pretraining might be insufficient (Kim and Schuster, 2023)

3 Paired Narrative Datasets

Our experiments use a dataset of paired coherent and incoherent stories based on materials from a reading psychology study by Rizzella and O’Brien (2024). The dataset contains 18 paired stories originally designed by psychologists for narrative comprehension research and undergone normative testing in human participants.

We chose this data set because it is the only dataset using the contradiction paradigm that has publicly available human psychometric data to our knowledge, thus enabling fine-grained comparison between human and LLM responses. Second, the incoherent scenarios presented in this dataset are well-studied; the majority were first conceived by Albrecht and O’Brien (1993) and have been replicated in multiple studies of human readers.

3.1 Story structure

All stories follow a fixed structure, as shown in Figure 2. Each coherent-incoherent pair consists of two almost completely identical stories, with the exception of the situation introduced after the introduction. We denote the consistent and inconsistent situations as s_c and s_i , respectively, and the critical target sentence as t . In the incoherent version, t conflicts with s_i , but notice that in the coherent version, t remains compatible with s_c . This design isolates the point of incoherence, which enables precise analysis of model and human responses to coherence violations.

3.2 Incoherence taxonomy and annotation

Incoherence occurs when it becomes impossible to hold one consistent interpretation of the text, but there are many possible sources of this interpretive

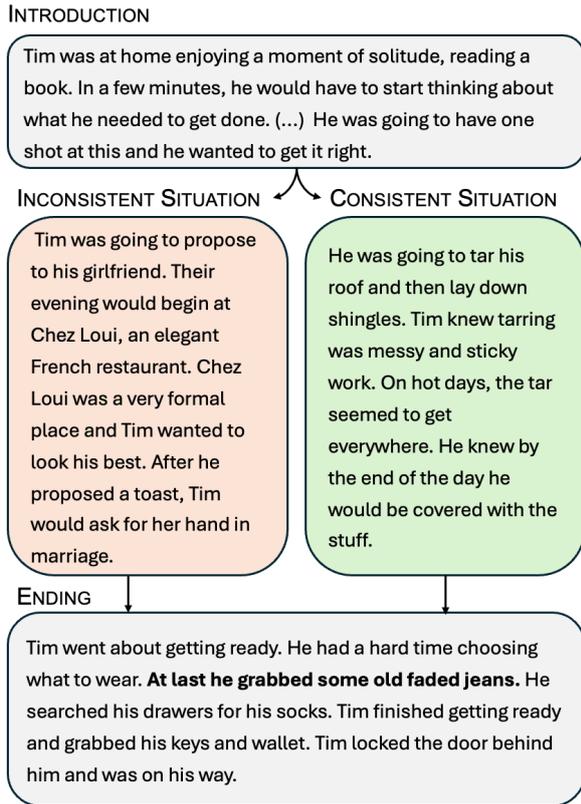


Figure 2: An exemplary incoherent-coherent story pair. Both versions share an identical introduction and conclusion, but differ in the middle situation: the incoherent version includes the inconsistent situation. Note the boldface **target** sentence, which conflicts with the inconsistent situation only. Some introductory text is omitted for brevity, denoted by (...).

breakdown (Reinhart, 1980). In particular, we note two distinct patterns in the dataset:

- *Trait-behavior conflict*: occurs when a strongly established character trait is later violated by that character’s behavior. For example, in Figure 1, Mary’s established vegetarian character trait conflicts with her later choice to order cheeseburger.
- *Event-setting conflict*: occurs when an event contradicts a strict expectation established by the narrative’s setting or social context. For example, in Figure 2, social expectations about the established setting of a formal restaurant conflicts with Tim dressing in old faded jeans.

Figure 3 also illustrates these two types of incoherence. We focus on this distinction because we hypothesize that LLMs may approach these two types of conflicts differently. Event-setting conflicts center on violation of the narrative situ-

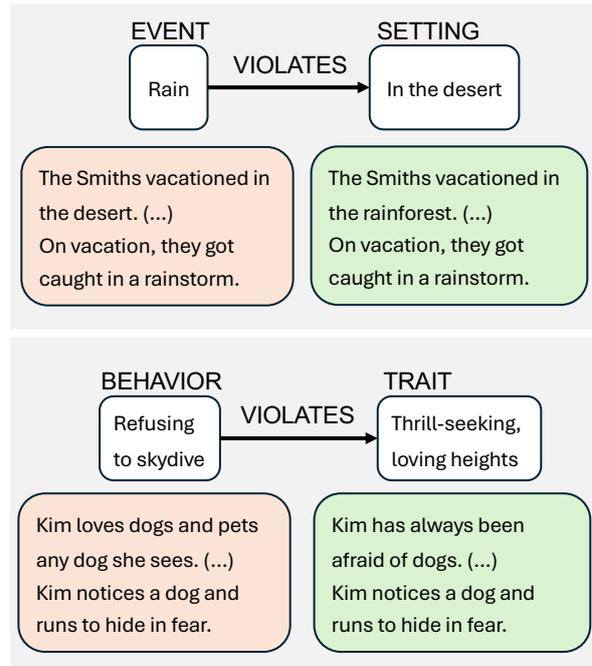


Figure 3: We identify and annotate two types of incoherence in the dataset: event-setting and trait-behavior. They differ in the target of the violation; specifically, trait-behavior is incoherent because it violates a strongly established character trait, while event-setting is incoherent because it violates some facet of the established narrative scene (e.g., the geographical location or the social norms).

ational context. Because all stories are set in realistic environments, such conflicts may engage the model’s pretrained world knowledge. In contrast, trait-behavior conflicts are centered around a character whose personality traits are defined only within the story itself, which may limit the influence of background knowledge and instead rely on internal narrative consistency.

To obtain ground truth labels for the incoherence type, we have three collaborators independently label each story’s conflict type based on the above definitions. The inter-annotator agreement is calculated with Krippendorff’s $\alpha = 0.835$, indicating strong agreement. Among the incoherent stories, 40% were even-setting conflicts ($n = 10$), and 60% were trait-behavior ($n = 15$).

3.3 Dataset augmentation

The original materials from Rizzella and O’Brien (2024) include 18 story pairs. To improve the reliability and precision of our conclusions, we augment it with seven additional pairs, resulting in a larger sample size of 25 story pairs (50 stories total).

The new stories are generated using a human-in-

the-loop writing process with Claude (Anthropic, 2025), followed by iterative refinement incorporating feedback from psychologists. All experimental results shown, unless otherwise stated, use the augmented dataset. The reported findings hold for both the original and augmented datasets, and experimental results for the original stories alone can be found in Appendix.

4 Experimental Setup and Methods

All experiments follow a consistent procedure for model inference and, in prompting experiments, systematic variation of prompts.

Model Inference. We use four open-source instruction-tuned LLMs across two model families and sizes: Llama3.1 with 8B and 70B parameters (Dubey et al., 2024), and Qwen2.5 with 7B and 72B parameters (Yang et al., 2024). We also test two reasoning variants from each model family: a Llama-8B model distilled from DeepSeek-R1 (DeepSeek-AI et al., 2025) and Qwen3-8B with reasoning enabled. We use the transformers library (Wolf et al., 2020) for model inference, applying 4-bit quantization (Dettmers and Zettlemoyer, 2023) to the larger models to meet computational constraints. Otherwise, default model configurations and generation hyperparameters are used in all cases. All experiments are completed on a Linux server equipped with one Nvidia A40 GPU (48GB).

Prompt Variation. Our experiments involving model generation systematically vary prompts and instructions to capture a more robust range of model responses. Specifically, we use FORMAT-SPREAD to vary punctuation and whitespace (Sclar et al., 2024), and we write three paraphrases of each instruction. The combination of these prompt variations yields 30 prompts per stimulus, giving us a better estimate of LLM performance. Specific prompts used can be found in §B.1.

5 Results

To answer our central question *Are LLMs able to reliably recognize incoherence in narratives?*, we perform extensive analyses: We first investigate how models encode narrative coherence internally (§5.1, §5.2) and then examine their explicit behavioral judgments (§5.3, §5.4). Finally, we compare model sensitivity across the two incoherence types and relate these findings to human data (§5.5).

5.1 LLMs assign lower likelihoods to incoherent events

We begin by testing whether LLMs recognize incoherent events as less probable within the narrative. For each story pair, we compute the perplexity of the target sentence t given the story context: $\text{PPL}(t|s_i)$ and $\text{PPL}(t|s_c)$ for the incoherent and coherent versions, respectively. Across all models, $\text{PPL}(t|s_i) > \text{PPL}(t|s_c)$, indicating that LLMs systematically estimate incoherent continuations as less likely or higher perplexity (Figure 5). This result provides initial evidence that LLMs encode a sensitivity to narrative coherence at the token-probability level.

5.2 Incoherence is recoverable from target hidden state

To test whether this coherence sensitivity is explicitly represented within model activations, we conduct a probing study (Belinkov, 2022; Alain and Bengio, 2018). We frame detection of incoherence as a binary classification task and fit a linear probe on representations extracted from the hidden state of each layer at the last token in t . Probes are evaluated with 10-fold cross-validation at the *paired* story level, ensuring that coherent and incoherent versions of a story always appear in the same fold.

Probes at the last token of the target sentence achieve high accuracy ($\approx 80\%$) across models, confirming that coherence information is linearly decodable from internal representations (Figure 4). The most informative layers tend to be mid-to-upper layers, consistent with prior findings on semantic integration (Hewitt and Manning, 2019). When probing the final story token, accuracy tends to dramatically decline, especially for smaller models. It is possible that coherence distinctions decay as the model processes later narrative content. Alternatively, coherence information may be encoded locally in the residual stream at the target token position and but remain accessible via attention mechanisms. The information would then not automatically be propagated forward to subsequent tokens, leading to decreased probing performance when probing the last sentence.

5.3 Explicit ratings of coherence and quality show weak behavioral sensitivity

Coherence Rating. We next assess whether the models’ explicit judgments reflect their internal

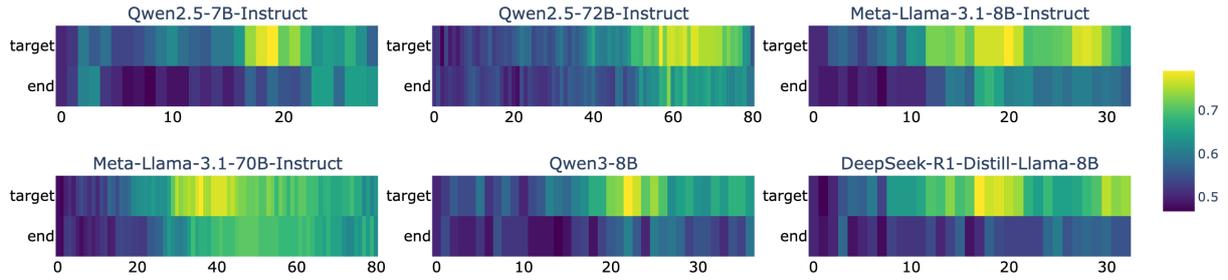


Figure 4: Mean accuracies across 10-fold cross-validation for probing hidden state representations to identify incoherent narratives. We probe both at the end of the target sentence that contains the incoherence in the incoherent story version (target), and at the conclusion of the story (end). The x axis denotes model layer. All models show strong separation at the target location, but by the story’s end, separation is notably weaker, with smaller models in particular near chance ($\approx 50 - 60\%$ accuracy). Llama3.1-70B has the best performance at the story’s end, and it also demonstrates the best understanding of coherence in responses to rating questions.

Model	Con. Rating (max 7)	Incon. Rating (max 7)	Diff. (\uparrow)	Con. Rating (max 7)	Incon. Rating (max 7)	Diff. (\uparrow)
<i>Instruction-tuned</i>						
→ Llama3.1 8B	6.26 (0.10)	6.03 (0.11)	0.00	4.49 (0.09)	4.39 (0.11)	0.10
→ Llama3.1 70B	6.47 (0.08)	4.61 (0.16)	1.66	4.97 (0.06)	4.58 (0.07)	0.39
→ Qwen2.5 7B	6.63 (0.06)	6.61 (0.06)	0.04	5.35 (0.12)	5.16 (0.11)	0.19
→ Qwen2.5 72B	6.21 (0.08)	5.53 (0.11)	0.73	4.71 (0.06)	4.50 (0.06)	0.21
<i>Reasoning-enabled</i>						
* Qwen3 8B	6.17 (0.11)	5.00 (0.14)	<u>1.18</u>	4.06 (0.16)	3.75 (0.17)	<u>0.31</u>
* R1-Distill-Llama 8B	6.17 (0.06)	5.97 (0.07)	<u>0.28</u>	5.15 (0.06)	5.09 (0.07)	<u>0.06</u>

(a) Mean LLM coherence ratings.

(b) Mean LLM quality ratings (aligned rows).

Table 1: Comparison of model ratings for (a) narrative coherence and (b) narrative quality across coherent and incoherent story versions. Ratings use a 7-point Likert scale (7 = fully coherent / highest quality, 4 = neutral). Each value reflects the mean across 30 prompt trials per story. Reasoning models are marked with *.

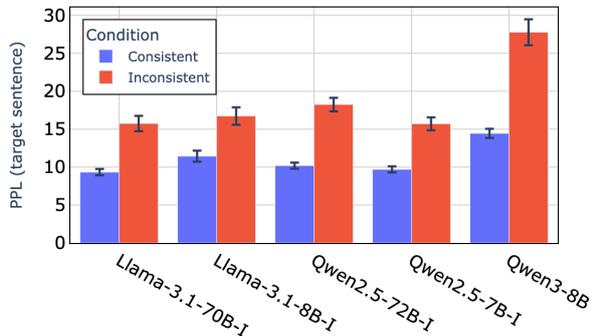


Figure 5: LLMs assign higher perplexity (i.e., lower likelihood) to incoherent events relative to coherent ones.

sensitivity. Each model rates story coherence on a 7-point Likert scale (7 = fully coherent), and we sample ratings across 30 perturbed prompts per story. Contrary to their internal representations, most models show minimal behavioral distinction between the two story types: the best-performing model (Llama3.1-70B) achieves a mean 1.66 point separation between coherent and incoherent stories, and four out of six models have less than one point

mean difference between coherent and incoherent stories (Table 1 (a)). Thus, while models encode coherence internally, this sensitivity is at best weakly expressed in their overt ratings.

Since it is possible that in a forced-choice scenario LLMs will show a stronger behavioral distinction between incoherent and coherent narratives, we additionally prompt models with a true or false evaluation of story coherence, but we do not find that this improves model performance (Figure 6).

Quality Rating. When prompted to rate narrative quality, models again exhibit limited differentiation: no model shows more than a one-point mean difference between coherent and incoherent versions (Table 1b). We do find that quality and coherence ratings have a weak to moderate positive correlation across all models (Pearson’s $r = 0.33, p < 0.05$), indicating that models do associate the two concepts, as expected.

		Probe Target	Probe End	Perplexity Target	Coherence Response	Quality Rating
Llama3.1	8B-Instruct	0.850	0.667	0.880	0.594	0.595
	70B-Instruct	0.942	0.817	0.880	0.765	0.672
Qwen2.5	7B-Instruct	0.800	0.683	0.689	0.720	0.561
	72B-Instruct	0.958	0.817	0.830	0.707	0.621
Reasoning	R1-Distill-Llama-8B	0.850	0.600	0.845	0.517	0.470
	Qwen3-8B	0.842	0.625	0.772	0.737	0.414

Table 2: Accuracy of various approaches in separating incoherent and coherent stories with LLMs. Measures based on internal model representations—perplexity and probe—work best, but these rely on *a priori* knowledge of the incoherent target sentence. Coherence ratings are only above chance for large models.

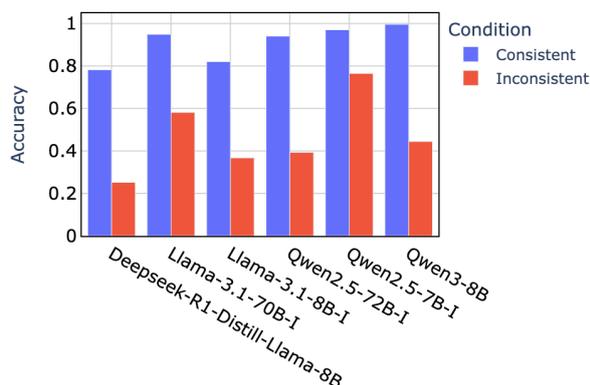


Figure 6: Model accuracies by condition when responding True or False to whether a story is coherent.

5.4 Reasoning models and metacognitive awareness

LLMs may be lacking in the meta-cognition required to transform an internal detection of incoherence into a coherence or quality rating. Reasoning-enabled variants are designed to generate explicit “thought strings” before answering. Thus, we test whether this reflective reasoning leads to stronger behavioral distinctions between the consistent and inconsistent stories.

While reasoning models occasionally verbalize the link between quality and coherence, their final ratings do not substantially improve (c.f. Qwen3-8B and Qwen2.5-7B, or Llama3.1-8B and R1-Distill-Llama-8B in Table 1).

Qwen3-8B thought (emphasis added)

So, the question is about the quality. What makes a story high quality? Usually, it’s about **coherence**, character development, plot, and whether it conveys a clear message or has a purpose. Let me break it down.

Moreover, these models sometimes produce

false positives when attempting to detect incoherence, such as misinterpreting consistent details as contradictions or hallucinating nonexistent incoherent events. This suggests that although reasoning prompts activate coherence-related concepts, they do not reliably translate internal sensitivity into accurate explicit judgments., e.g.:

Qwen3-8B thought (false positive)

I need to check if there’s any inconsistency. The story mentions Bill does additional workouts before and after his walks. That’s a bit odd because if he’s walking, maybe he’s doing it as part of his exercise.

5.5 Comparison to Human Readers: Sensitivity to Incoherence Types

Across analyses, models show a tendency toward *greater sensitivity to event–setting incoherence* compared to trait-behavior incoherence. Perplexity differences are larger for event–setting conflicts (Figure 8), and probe accuracies are also consistently higher for this type (Figure 7). This pattern aligns with our hypothesis that event–setting incoherence leads to higher engagement of pre-trained world knowledge, whereas trait–behavior incoherence depends more on within-story reasoning, which models handle less consistently.

In contrast, we do not find evidence that human readers show a similar sensitivity. We examine human reading times² of the target sentence of each story using the publicly available data collected by Rizzella and O’Brien (2024). On average, reading times were similar for setting-event stories ($M = 1967.91$, $SD = 478.80$) and trait-behavior stories ($M = 1913.62$, $SD = 487.55$). Reading times

²There is a well-established relationship between language model surprisal and human reading times (Wilcox et al., 2020; Oh et al., 2022).

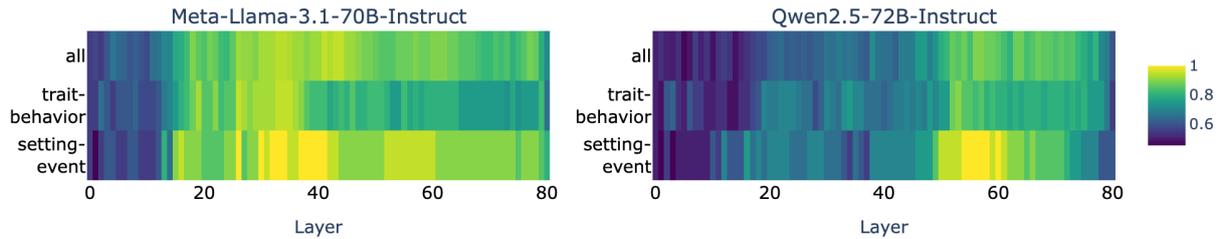


Figure 7: Probe accuracy in detecting incoherence using a linear regression train on hidden states, averaged across 10-fold cross validation. We conduct a probing study over three datasets: (1) only trait-behavior incoherence stories, (2) only setting-event incoherence stories. Probes generally achieve better separation for the setting-event stories.

are measured in milliseconds. A paired-samples t-test confirmed this difference was not statistically significant, $t(29) = 0.65, p = .523, M_{diff} = 54.29, 95\%CI = [-117.39, 225.97]$. Since the human study was not designed to test story type effects, individual participants did not necessarily read all possible combinations of story type and consistency conditions (e.g., inconsistent \times setting-event). Therefore, we do not test for possible interaction effects between story category and consistency due to limited statistical power.

5.6 Disentangling incoherent and unexpected events

It is possible that the discrepancy between model representations at incoherent events and final model responses stems from the model recognizing an unexpected event in the story, but not actually identifying it as *incoherent*. While incoherent events are unexpected, not all unexpected events are necessarily incoherent; in fact, unexpected events are oftentimes the key to a good story (Kintsch, 1980). Thus, LLMs confounding unexpectedness and incoherence would be highly undesirable.

To further investigate this possibility, we utilize the *Qualified* version of the inconsistent stories, introduced in O’Brien et al. (1998). The qualified story is identical to the incoherent story, but it includes a qualifying elaboration that makes the story coherent. For instance, consider the qualifying elaboration in the story of Mary, the cheeseburger eating vegetarian:

At the party where she met Mary, Joan overheard people talking about someone who had been a strict vegetarian for ten years. Joan made a mistake and assumed they were talking about Mary. Joan had thought that Mary’s favorite restaurant had fantastic health food. She thought Mary was a health nut and wouldn’t eat

anything which was fried or cooked in grease. She also thought Mary’s favorite food was cauliflower.

The elaboration makes it so that the target event (“Mary ordered a cheeseburger and fries”) is not incoherent. If models truly have separate representations of incoherent and coherent stories, then these qualified coherent stories should also be separable from the incoherent stories in which the target event (“Mary ordered a cheeseburger and fries”) is very unexpected. However, we find that probes have a much lower accuracy when separating the qualified and incoherent stories, relative to the coherent and incoherent stories (Table 3).

The subtle difference between the qualified coherent story and the incoherent story is an interesting test case for internal representations of coherence in LLMs. These results show that while large models (Llama3.1 70B and Qwen2.5 72B) are above chance at distinguishing the qualified and incoherent stories, the separation is substantially less pronounced than it is in the coherent and incoherent stories, suggesting that LLMs do not have a particularly robust internal representation of narrative coherence.

6 Discussion

Claims that LLMs truly understand narratives ultimately hinge on whether they can monitor and maintain coherence across a story. Our experiments aim to shed light on whether LLMs actually recognize contradictions within a narrative or simply respond to shallow statistical patterns.

Our findings indicate that, even for the relatively short and simple narratives studied here, the LLMs tested – especially the smaller ones at 7-8B parameters – have difficulty reliably separating incoherent and coherent narratives. Additionally, all LLMs tested fail to consistently identify coherent stories as higher quality, raising questions about

Model	Incoherent/ Qualified	Incoherent/ Coherent	Coherent/ Qualified
Llama3.1 8B	.58	.84	.48
Llama3.1 70B	.64	.92	.54
Distill Llama	.41	.81	.44
Qwen2.5 7B	.44	.83	.43
Qwen2.5 72B	.62	.89	.53
Qwen3 8B	.48	.80	.53

Table 3: Linear probes for separating incoherent / qualified, incoherent / coherent, and coherent / qualified stories. In particular, if models have advanced understanding of coherence, they should achieve high separation between not only incoherent and coherent stories, but also incoherent and qualified stories. However, accuracy for incoherent and qualified stories is lower, and often relatively close to chance (.5). For each layer, we performed 10-fold cross validation and measured test accuracy. The reported scores reflect the layer that achieved the best classification performance.

LLM reliability on tasks involving understanding or evaluating storytelling. Despite these poor responses on various rating questions, probing does indicate that models’ internal representations distinguish between coherent and incoherent stories, bringing *why* they fail to give reasonable coherence and quality ratings into question.

One possible contributing factor is suggested by the drop in performance between probes on the token at the end of the target incoherence location and probes on the token ending the story: while LLMs may temporarily represent incoherence, they may lose access to that representation at later locations in the story. Interestingly, Llama-70B has the highest performance on the story-end probe and also has the best rating responses, which supports this interpretation.

Another possible explanation for the discrepancy between model internal states and responses is that the model representations are in fact capturing an attribute highly associated with incoherence, such as unexpectedness. To investigate this possibility, we test versions of the inconsistent stories with a *qualifying elaboration* that makes the stories coherent. Internal model representations have substantially worse separation between the qualifying coherent and the incoherent stories, suggesting that narrative coherence may not constitute a robustly encoded feature within these models.

Additionally, we note an apparent tendency for models to respond more to incoherence based on

event-setting conflicts, rather than trait-behavior conflicts. This is puzzling as humans do not appear to share this tendency.³ It is possible that background world knowledge modulates the models’ response heavily in some event-setting instances. Future work investigating, e.g., how LLMs respond to event-setting conflicts in fantasy worlds with narrative-specific rules would be another important direction for isolating this phenomenon.

7 Conclusion

We investigate how well six LLMs identify incoherence in narratives using a combination of metrics based on model internal representations and responses to prompts. To facilitate our experiments, we adapt and augment a paired narrative dataset originally used in human research. Our results indicate that while internal representations achieve some separation – albeit imperfect – of incoherent and coherent narratives at the location of incoherence, this separation largely disappears by the end of the story. We also find that this internal separation is less reliable at distinguishing between coherent and incoherent stories when the coherent story requires a qualifying elaboration to maintain consistency. We also find that LLM responses identify the incoherence at generally lower rates than we find via perplexity or probing at the target. Together, our results suggest that LLMs lack a robust coherence-monitoring process during narrative comprehension.

8 Limitations

Our work studies narratives made for psychological studies rather than narratives drawn from naturally occurring stories. This also necessarily means that the dataset is limited in size, as it is hand-crafted. The advantage of this approach is that the stories are well designed to isolate the attribute of interest (coherence), but the results must be interpreted with the above caveats in mind.

Our analysis also does not investigate any proprietary models, as we rely on access to model hidden states. In a similar vein, while we examine six models, we cannot be certain our results generalize more broadly.

³The result hints at low understanding of character incoherence; it may align with an earlier finding that language model summaries make the most coherence errors with respect to character references (Goyal et al., 2022).

Ethics Statement

Our analyses of human data use only publicly available data that do not contain personally identifiable information.

Additionally, the application of LLMs in creative domains such as literature raises ethical concerns regarding the potential displacement of human labor and creative expression. There is a risk of societal impact if LLMs begin to assume roles traditionally held by human writers and storytellers.

Acknowledgements

The authors gratefully acknowledge the support of the University of Minnesota Graduate School through the Doctoral Dissertation Fellowship. We also thank the members of the University of Minnesota Text Group and NLP Lab for their helpful feedback.

References

- Guillaume Alain and Yoshua Bengio. 2018. [Understanding intermediate layers using linear classifier probes](#). *Preprint*, arXiv:1610.01644.
- Jason E Albrecht and Edward J O'Brien. 1993. Updating a mental model: Maintaining both local and global coherence. *Journal of experimental psychology: Learning, memory, and cognition*, 19(5):1061.
- Anthropic. 2025. [Claude 4.5 Sonnet](https://claude.ai/). <https://claude.ai/>. Retrieved [date].
- Yonatan Belinkov. 2022. [Probing classifiers: Promises, shortcomings, and advances](#). *Computational Linguistics*, 48(1):207–219.
- Tyler A Chang and Benjamin K Bergen. 2024. Language model behavior: A comprehensive survey. *Computational Linguistics*, 50(1):293–350.
- Karin de Langis, Jong Inn Park, Andreas Schramm, Bin Hu, Khanh Chi Le, Michael Mensink, Ahn Thu Tong, and Dongyeop Kang. 2025. How llms comprehend temporal meaning in narratives: A case study in cognitive evaluation of llms. *arXiv preprint arXiv:2507.14307*.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 181 others. 2025. [DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning](#). *Preprint*, arXiv:2501.12948.
- Tim Dettmers and Luke Zettlemoyer. 2023. The case for 4-bit precision: k-bit inference scaling laws. In *International Conference on Machine Learning*, pages 7750–7774. PMLR.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Simon Garrod and Anthony J Sanford. 1981. Bridging inferences and the extended domain of reference. *Attention and performance IX*, pages 331–346.
- Morton Ann Gernsbacher. 1997. Attenuating interference during comprehension: The role of suppression. In *Psychology of Learning and Motivation*, volume 37, pages 85–104. Elsevier.
- Arthur M Glenberg and William E Langston. 1992. Comprehension of illustrated text: Pictures help to build mental models. *Journal of memory and language*, 31(2):129–151.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. Snac: Coherence error detection for narrative summarization. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 444–463.
- John Hewitt and Christopher D Manning. 2019. A structural probe for finding syntax in word representations. 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 4129–4138.
- Najoung Kim and Sebastian Schuster. 2023. [Entity tracking in language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3835–3855, Toronto, Canada. Association for Computational Linguistics.
- Walter Kintsch. 1980. Learning from text, levels of comprehension, or: Why anyone would read a story anyway. *Poetics*, 9(1-3):87–98.
- Walter Kintsch. 1998. *Comprehension: A paradigm for cognition*. Cambridge university press.
- Walter Kintsch and Teun A Van Dijk. 1978. Toward a model of text comprehension and production. *Psychological review*, 85(5):363.
- Kenneth Li, Aspen K Hopkins, David Bau, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023. [Emergent world representations: Exploring a sequence model trained on a synthetic task](#). In *The Eleventh International Conference on Learning Representations*.
- Joseph P Magliano, Rolf A Zwaan, and Arthur C Graesser. 1999. The role of situational continuity in narrative understanding. *The construction of mental representations during reading*, pages 219–245.

- Kyle Mahowald, Anna A Ivanova, Idan A Blank, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. 2024. Dissociating language and thought in large language models. *Trends in cognitive sciences*, 28(6):517–540.
- Gail McKoon and Roger Ratcliff. 1992. Inference during reading. *Psychological review*, 99(3):440.
- Danielle S McNamara and Joe Magliano. 2009. Toward a comprehensive model of comprehension. *Psychology of learning and motivation*, 51:297–384.
- Jerome L Myers and Edward J O’Brien. 1998. Accessing the discourse representation during reading. *Discourse processes*, 26(2-3):131–157.
- Neel Nanda, Andrew Lee, and Martin Wattenberg. 2023. [Emergent linear representations in world models of self-supervised sequence models](#). In *Proceedings of the 6th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 16–30, Singapore. Association for Computational Linguistics.
- Edward J O’Brien and Jason E Albrecht. 1992. Comprehension strategies in the development of a mental model. *Journal of experimental psychology: learning, memory, and cognition*, 18(4):777.
- Edward J O’Brien, Michelle L Rizzella, Jason E Albrecht, and Jennifer G Halleran. 1998. Updating a situation model: a memory-based text processing view. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(5):1200.
- Byung-Doh Oh, Christian Clark, and William Schuler. 2022. Comparison of structural parsers and neural language models as surprisal estimators. *Frontiers in Artificial Intelligence*, 5:777963.
- Tanya Reinhart. 1980. Conditions for text coherence. *Poetics today*, 1(4):161–180.
- Michelle L Rizzella and Edward J O’Brien. 2024. When prospective information conflicts with current information. *Discourse Processes*, 61(9):421–438.
- Sebastian Schuster and Tal Linzen. 2022. When a sentence does not introduce a discourse entity, transformer-based models still sometimes refer to it. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 969–982.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2024. [Quantifying Language Models’ Sensitivity to Spurious Features in Prompt Design or: How I learned to start worrying about prompt formatting](#). In *The Twelfth International Conference on Learning Representations*.
- Woomin Song, Seunghyuk Oh, Sangwoo Mo, Jaehyung Kim, Sukmin Yun, Jung Woo Ha, and Jinwoo Shin. 2024. Hierarchical context merging: Better long context understanding for pre-trained llms. In *12th International Conference on Learning Representations, ICLR 2024*.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063.
- P Van den Broek, M Young, Y Tzeng, T Linderholm, H van Oostendorp, and SR Goldman. 1999. The construction of mental representations during reading. *The Construction of Mental Representations During Reading*, pages 71–98.
- Teun Adrianus Van Dijk, Walter Kintsch, and 1 others. 1983. Strategies of discourse comprehension.
- Ethan G Wilcox, Jon Gauthier, and Jennifer Hu. 2020. On the predictive power of neural language models for human real-time comprehension. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, 42 (0).
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, and 3 others. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Rolf A Zwaan and Gabriel A Radvansky. 1998. Situation models in language comprehension and memory. *Psychological bulletin*, 123(2):162.

A Appendix

B Licenses of Models

All licenses are fair use for this work.

1. LLaMA-3.1 models have a Llama 3.1 Community License.
2. Qwen models have an Apache 2.0 license.

B.1 Prompts

We use FORMATSREAD to vary whitespace and punctuation of stimuli data within the prompt elements in the following 10 ways:

- Story: (ISTORY|
(IQUESTION|

- Field: {}
Answer: {}
Story: (|STORY|)
(|QUESTION|)
- Field: {} <sep> Story: (|STORY|) <sep>
(|QUESTION|)
- Field - {}. Story - (|STORY|). (|QUESTION|)
- Field{̂}. Story{̂}(|STORY|). (|QUESTION|)
- FIELD- {}
STORY- (|STORY|)
(|QUESTION|)
- field:: {}
story:: (|STORY|) – (|QUESTION|)
- field - {}
story - (|STORY|) , (|QUESTION|)
- Field
{̂}
- {̂}
Story
{̂}(|STORY|)
{̂}(|QUESTION|)
- Field - {}
Story - (|STORY|)
(|QUESTION|)

B.2 Coherence Rating Prompts

Our three paraphrasings for coherence prompts (used in conjunction with above format spread to create 30 prompts).

- In this activity, you will be presented with a story to read. Afterwards, you’ll be asked a question about the story.
QUESTION: How coherent or consistent does the story seem on a scale from 1-7, with 1 being very incoherent and 7 being very coherent.
- In this exercise, you will go through a narrative, and upon finishing, you’ll answer a question about how good the story way.
QUESTION: Rate the quality of the story on a scale from 1-7, with 1 being poor coherence and 7 being high coherence. Respond only with a number between 1 and 7.

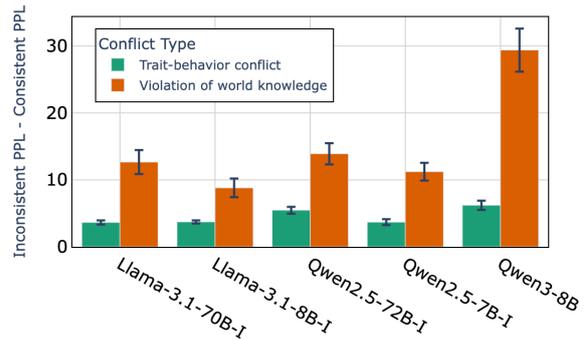


Figure 8: We find that there is a higher difference in perplexity for stories with event-setting incoherence.

- In this task you will read a story. At the end you will be asked a question about the story.

QUESTION: How coherent or consistent does the story seem on a scale from 1-7, with 1 being very incoherent and 7 being very coherent.

Quality prompts follow this format, replace “coherent” with quality and “incoherent” with low quality.

B.3 Coherence Forced Choice Prompts

Our three paraphrasings for coherence forced choice prompts (used in conjunction with above format spread to create 30 prompts).

- In this activity, you will be presented with a story to read. Afterwards, you’ll be asked a question about the story.

Story: (|STORY|)

Question: True or False: the story seems coherent and consistent. Only answer with True or False.

- In this exercise, you will go through a narrative, and upon finishing, you’ll answer a question about how good the story way.
- In this task you will read a story. At the end you will be asked a question about the story.

B.4 Perplexity differences

We find that perplexity differences are higher for stories with event-setting incoherence (Fig. 8).

C No Augmentations

Our main results with no augmentations to the dataset. Coherence ratings Table 4, quality ratings 5. Perplexity info is in Table 6

Model	Coherent	Incoherent
Distill-Llama-8B	6.142 ± 0.073	5.937 ± 0.082
Llama-3.1-70B-I	6.434 ± 0.094	4.681 ± 0.176
Llama-3.1-8B-I	6.240 ± 0.111	6.006 ± 0.131
Qwen2.5-72B-I	6.150 ± 0.091	5.624 ± 0.126
Qwen2.5-7B-I	6.649 ± 0.061	6.628 ± 0.064
Qwen3-8B	6.091 ± 0.127	4.947 ± 0.153

Table 4: Mean coherence ratings on original stories only.

Model	Coherent	Incoherent
Gemma-2-27Bi	4.037 ± 0.091	3.852 ± 0.109
Llama-3.1-70B-I	4.685 ± 0.065	4.426 ± 0.080
Llama-3.1-8B-I	4.875 ± 0.119	4.456 ± 0.107
Qwen2.5-72B-I	4.562 ± 0.062	4.352 ± 0.071
Qwen2.5-7B-I	5.150 ± 0.060	5.006 ± 0.071
Qwen3-8B	4.056 ± 0.164	3.747 ± 0.166

Table 5: Mean quality ratings on original stories only.

Model	Conflict Type	Coherent PPL	Incoherent PPL
Distill-Llama-8B	Trait-behavior conflict	12.113 ± 1.159	22.253 ± 2.062
Distill-Llama-8B	Violation of world knowledge	39.761 ± 7.062	84.328 ± 21.749
Llama-3.1-70B-I	Trait-behavior conflict	6.251 ± 0.456	11.650 ± 1.128
Llama-3.1-70B-I	Violation of world knowledge	15.958 ± 2.360	31.253 ± 6.545
Llama-3.1-8B-I	Trait-behavior conflict	6.472 ± 0.492	11.272 ± 0.749
Llama-3.1-8B-I	Violation of world knowledge	22.053 ± 4.347	35.819 ± 7.532
Qwen2.5-72B-I	Trait-behavior conflict	6.101 ± 0.357	13.286 ± 1.071
Qwen2.5-72B-I	Violation of world knowledge	14.509 ± 1.615	27.609 ± 4.989
Qwen2.5-7B-I	Trait-behavior conflict	6.224 ± 0.368	9.297 ± 0.790
Qwen2.5-7B-I	Violation of world knowledge	16.371 ± 1.831	25.058 ± 4.517
Qwen3-8B	Trait-behavior conflict	9.733 ± 0.882	16.768 ± 1.789
Qwen3-8B	Violation of world knowledge	23.747 ± 2.223	49.794 ± 8.290

Table 6: Non-augmented dataset perplexity means across coherent and incoherent stories for different types of incoherence.