

How LLMs Comprehend Temporal Meaning in Narratives: A Case Study in Cognitive Evaluation of LLMs

Karin de Langis¹, Jong Inn Park¹, Andreas Schramm²,
Bin Hu¹, Khanh Chi Le¹, Dongyeop Kang¹

¹University of Minnesota, ²Hamline University

Correspondence: dento019@umn.edu

Abstract

Large language models (LLMs) exhibit increasingly sophisticated linguistic capabilities, yet the extent to which these behaviors reflect human-like cognition versus advanced pattern recognition remains an open question. In this study, we investigate how LLMs process the temporal meaning of linguistic aspect in narratives that were previously used in human studies. Using an Expert-in-the-Loop probing pipeline, we conduct a series of targeted experiments to assess whether LLMs construct semantic representations and pragmatic inferences in a human-like manner. Our findings show that LLMs over-rely on prototypicality, produce inconsistent aspectual judgments, and struggle with causal reasoning derived from aspect, raising concerns about their ability to fully comprehend narratives. These results suggest that LLMs process aspect fundamentally differently from humans and lack robust narrative understanding. Beyond these empirical findings, we develop a standardized experimental framework for the reliable assessment of LLMs’ cognitive and linguistic capabilities.

1 Introduction

Advanced state-of-the-art large language models (LLMs) exhibit strikingly human-like behavior. However, the extent to which these models demonstrate genuine understanding, rather than sophisticated statistical pattern matching, remains an open question (Maleki et al., 2024; Li, 2023; Bender et al., 2021; Das et al., 2024). A growing body of research attempts to address this issue by applying experimental methodologies from cognitive science to evaluate cognitive and linguistic capabilities of LLMs (Ivanova, 2025). We contribute to this effort by investigating how LLMs comprehend the semantic and pragmatic temporal meanings expressed by *aspect* in narratives (e.g., Fig. 1).

Aspect is a complex two-component linguistic phenomenon that allows the speaker to present the

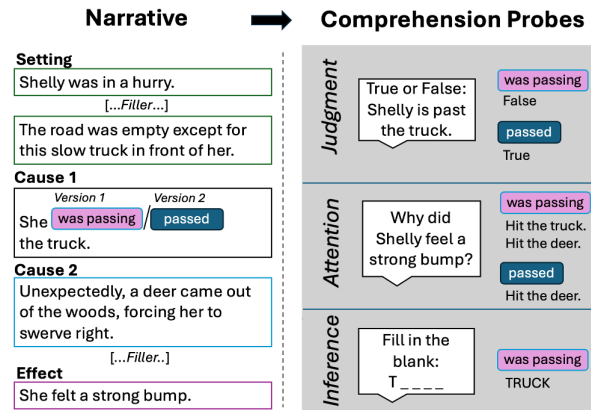


Figure 1: We examine how LLMs understand differences in aspect by presenting LLMs with narratives that have a key word either in the imperfective or perfective (e.g., “was passing” vs. “passed”) followed by comprehension probes adapted from previous human studies.

inherent temporal structure of a given event (known as lexical aspect¹) with different focuses (referred to as grammatical aspect²). Crucially, the imperfective aspect (e.g., “was passing”) presents an event as ongoing, without reaching a final state. In narrative contexts, this has been shown to *facilitate causal inference*, as ongoing events are more likely to be retained in working memory and integrated into a reader’s mental model of the story (Magliano and Schleich, 2000; Mozuraitis et al., 2013; Schramm and Mensink, 2016).

Because of its nuanced role in human language processing, aspect provides a valuable test case for probing LLM capabilities in semantic and pragmatic comprehension. Do LLMs, like humans, recognize aspect as a cue for retaining an event as open-ended or closing it off as completed? Do

¹For instance, the lexical aspect of “walk in the park” and “walk to the park” differ, since “walk to the park” has a natural endpoint and “walk in the park” does not.

²The two main grammatical aspects in English are the perfective (e.g., “walked”) and the imperfective (e.g., “was walking”).

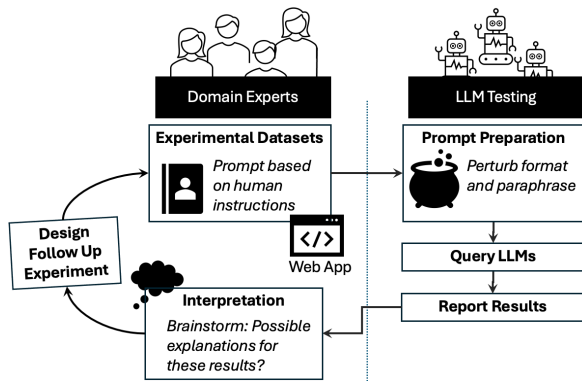


Figure 2: A conceptual overview of our Expert-in-the-Loop probing pipeline for assessing cognitive abilities of LLMs, designed in close collaboration with domain experts from cognitive science.

they grasp the causal implications conveyed by aspectual distinctions, and can they recognize the non-prototypical use of an open-ended event in a narrative? To answer these questions, we conduct a series of experiments using narrative stimuli and experimental paradigms previously developed by cognitive linguists to study similar questions in humans. To support these experiments, we build an Expert-in-the-Loop probing pipeline for LLM assessments based on experimental materials from existing studies (Fig. 2).

Our results indicate that while LLMs have some understanding of aspectual semantics, their comprehension remains weaker than expected and their cognitive performance across various measures is inconsistent (i.e., no model maintains human-like responses across all measures). Moreover, we find that many state-of-the-art LLMs struggle with the pragmatic causal narrative and backgrounding implications of imperfective aspect, suggesting an over-reliance on prototypical linguistic structures rather than a flexible, context-driven interpretation. These findings contribute to the growing understanding of LLM cognition (Apidianaki et al., 2024; Ohmer et al., 2024) and pragmatic reasoning (Beuls and Van Eecke, 2024; Sravanthi et al., 2024), underscoring the limitations of current LLMs and highlighting the need for further research into the cognitive foundations of LLM behavior.

2 Related Work

Recent advances in LLM have led to an emerging area of research that assesses abilities and characteristics of LLMs by leveraging datasets and methods from human studies (Ivanova, 2025; Hagendorff

et al., 2023). These works study a variety of domains: cognition (Roberts et al., 2024b; Echterhoff et al., 2024), causal reasoning (Dasgupta et al., 2022; Binz and Schulz, 2023), decision-making (Hagendorff et al., 2023; Coda-Forno et al., 2024), philosophy of mind (Ullman, 2023; Echterhoff et al., 2024), psycholinguistics (Bazhukov et al., 2024; Lee et al., 2024), and emotion and personality (Coda-Forno et al.; tse Huang et al., 2024).

Measurement Types and Challenges Studies assessing LLMs’ cognitive abilities typically rely on two types of indicators: **explicit outputs**, such as self-report responses of the model (Dasgupta et al., 2022; Hagendorff et al., 2023) and counterfactual prompting (Roberts et al., 2024a), or **implicit signals**, such as token probabilities (Lee et al., 2024; Ullman, 2023) and attention scores (Bazhukov et al., 2024). For example, Roberts et al. (2024a) investigated the Fan Effect (Anderson, 1974), in which people take longer to retrieve information from memory when there is a large number of facts linked to that concept, in LLMs by measuring the token probability as a proxy for memory retrieval difficulty. Most studies employ a single measure, but it is ideal to collect several measures to establish convergent evidence of a given cognitive effect in LLMs.

Translating human-designed experiments to LLMs poses several challenges. A key issue is that LLM responses are highly sensitive to prompt contents. Some studies address this by introducing prompt perturbation methods (Coda-Forno et al.; Bazhukov et al., 2024) to account for variability in model outputs, and Roberts et al. (2024b) propose injecting variability by perturbing model weights through the addition of dropout layers. Rather than adapting datasets from human studies, several works instead assess LLM cognition using existing NLP task datasets (Ying et al., 2024; Lu et al., 2024; Shah et al., 2024) or curating new datasets tailored to specific research questions (Jones and Steinhardt, 2022; Lal et al., 2024; Lv et al., 2024; Joshi et al., 2024; Jung et al., 2024; Lee and Lim, 2024). This approach allows researchers to tailor materials to LLMs, but sacrifices the high quality offered by datasets designed by social scientists. This approach also introduces challenges in comparing LLMs to humans due to a lack of pre-existing human data. Some studies collect human performance data for comparison (Lee and Lim, 2024), but many omit this step (Ying et al., 2024;

Joshi et al., 2024). High quality human comparisons are important to probe details of LLM cognitive processes, since LLMs may achieve similar results to humans through fundamentally different cognitive mechanisms (Ivanova, 2025).

3 Preliminary on Aspect and Humans

Aspect is a particularly subtle meaning domain that is especially challenging for learners of English as a second language (Andersen, 1996; Bardovi-Harlig, 2000). The events in this study belong to the lexical aspect category “Accomplishment” and therefore come with an initial state, transition point, and final state. Accomplishment events prototypically pair up with *perfective* grammatical aspect because the perfective also has an endpoint (Salaberry, 2024). Thus, it takes an advanced command of the language to combine the Accomplishment events in this study with an imperfective aspect that focuses on the initial state rather than the final state: such combinations are ultimately tied to **pragmatic text-level decisions** by speakers (Hopper, 1982). For instance, the non-prototypical coupling of imperfective aspect with Accomplishment events can signal to readers that something unusual may happen (Hopper, 1982). Thus, for native speakers or non-natives with advanced proficiency, aspect significantly influences narrative comprehension, shaping how readers process ongoing narratives and form situation models (Zwaan et al., 1995), as well as how they incorporate events into their global discourse model (Mozuraitis et al., 2013).

Aspect also plays a critical role in the comprehension of **causal relations** within narratives. Specifically, events marked by the *imperfective aspect* are more likely to be interpreted as *potential causes* for subsequent events. Consider the example in Table 3: When imperfective aspect is used (“Rob *was washing* the dishes”), the action remains ongoing, making it a plausible cause of the subsequent loud noise. When perfective aspect is used (“Rob *washed* the dishes”), the action is interpreted as completed, reducing its potential causal relevance to the unfolding narrative.

3.1 Linguistic Design of Narrative Stimuli

The narratives used in this study are adapted from linguistic research on aspect comprehension in English (Schramm, 1998). They contain an unusual combination of imperfective aspect and Accomplishment, providing a good opportunity to test

Intro	Rob and Alisha had a nice system going.
Filler	Each day they split up their duties and rotated them. Today, Alisha took care of the living room, and Rob was down for kitchen.
C_1	Rob <i>*was washing*</i> the dishes.
C_2	Alisha watered the plants and started arranging flowers in one of their special vases.
Filler	She felt bad because it had been a terrible day for her. She overslept, missed her bus, and was reprimanded by her boss. In her rush she forgot her purse. And on top of everything her lunch date did not show.
Effect	Suddenly there was a loud noise.

Table 1: Structure of narratives designed by linguists to probe the impact of aspect on causal inferencing. The end effect has two possible causes, C_1 and C_2 . When C_1 is imperfective (rather than perfective), it has a higher potential connection to the final effect.

whether LLMs are able to handle temporality in a fashion that parallels the behavior of human native speakers. Our study employs the inherent temporal structure of Accomplishment events to examine whether LLMs exhibit aspect-driven inference and processing patterns when reading narratives. To this end, narratives follow this format (e.g., Table 3):

1. Early in the narrative, a potential cause event, C_1 is presented. The event is comprised of a prototypical closed-ended two-state Accomplishment situation (e.g., “dishes unwashed” → “dishes washed”) (Klein, 1994).
2. A second potential cause, C_2 , follows C_1 .
3. A surprise-effect event, which can be causally linked to either C_1 or C_2 , ends the narrative.

This design allows us to influence the causal connection of C_1 to the final surprise event solely by changing the *aspect markings* of the verb in C_1 . When C_1 is imperfective (*was washing*), it remains conceptually active as the reader processes the subsequent text, making both C_1 and C_2 plausible causes of the effect. When C_1 is perfective (*washed*), it is interpreted as completed and recedes from the reader’s focus, making C_2 the more likely inferred cause.

3.2 Genuine Cognitive Comprehension

To compare the performance of LLMs with that of humans in the arena of cognitive processes involved in understanding (Apidianaki et al., 2024; Ohmer et al., 2024), we investigate how LLMs first cognitively process temporal meanings expressed by linguistic aspect in narratives in working memory and eventually understand them in episodic memory. In humans, aspect information is encoded in working memory during reading. It influences the integration of subsequent events into the discourse model, and is modulated by world knowledge regardless of age (Mozuraitis et al., 2013). When the critical sentence is in the imperfective (“Rob was washing the dishes”), it is stored “in focus” in episodic memory after five sentences and can later be reactivated in working memory, because there is no clear endpoint (Schramm and Mensink, 2016). This in turn implies continued relevance as a potential cause for later events (“Suddenly there was a loud noise”). By contrast, when the critical sentence is in the perfective (“Rob washed the dishes”), the event itself is less accessible in episodic memory and less likely to be reactivated in working memory, since it is presented as completed, making it an unlikely or even impossible cause of the later surprise event. Only the after-effect of the event is available (Hart and Albarracín, 2009). Furthermore, to signal in-focus status requires indirect, pragmatic understanding of aspect in a narrative context. Open-ended imperfective aspect in narratives in the foreground is highly unusual (Hopper, 1982), because events are in the past and are prototypically narrated in the perfective aspect. This non-prototypicality should lead to greater attention and in turn activation in participants’ working memory, which can then be used for causal inferencing.

After processing in working memory during reading is complete, we investigate the pragmatic information available in episodic memory in humans. We query the discourse model regarding the cause for the surprise effect at the end of a story (“What was the reason for the loud noise?”). When the cause is in the open-ended imperfective, it is predicted to be stored accordingly more frequently in episodic memory than in the perfective.

It is also important for an assessment of human versus LLM performance to compare humans’ cognitive processes involved in processing sentence-level semantic as opposed to text-level pragmatic

meaning with those of LLMs. To accomplish this, we shifted to a semantic measure of aspect awareness at the sentence level. Truth value judgments obtained from humans’ episodic memories ascertain whether a causal event was closed off (“dishes washed” in the perfective) or open-ended (“dishes unwashed” in the imperfective). In the former case, it would be less available for further cognitive processing, such as causal inferencing, than in the latter case.

4 Expert-in-the-loop Probing Pipeline

To systematically investigate LLM behavior in comparison to human cognition, we develop an Expert-in-the-Loop probing pipeline. This pipeline facilitates controlled behavioral experiments with LLMs by integrating experimental material input, automatic prompt perturbation, and response parsing for preliminary result evaluation. Our philosophy is that the pipeline should be applied iteratively, to gather converging evidence, with experts assessing the intermediate results between iterations.

Overview Our pipeline is designed to ensure robust and interpretable LLM behavioral evaluations. The process consists of three core stages:

- Prompting – Constructing LLM prompts *faithful* to human experimental designs.
- Prompt Paraphrasing – Introducing controlled variations to ensure *robustness*.
- Models and Inference – Running experiments across multiple LLMs for supporting *model agnostic* behaviors.

Below, we describe each stage in detail.

Prompting. For each behavioral experiment, we develop a task-specific prompt based on the instructions used in a corresponding human study. While instructions are modified to optimize clarity for LLMs (e.g. converting “circle your answer” into “respond with your answer,” and applying structured formatting), we maintain fidelity to the original human task to enable valid comparisons between LLM and human responses. Each prompt explains the task in clear, structured language, with an example to illustrate how the task should be completed. Our pipeline also enforces two distinct components in prompts: general instructions, which provides the task details and an example, and a data format, which specifies how instances should be formatted. The general prompt serves as an instruction to the model, while the data format orga-

nizes task-specific input representations. For example, a general prompt may state: “*Read the stories and follow the general instructions...*”, while a corresponding data format could be structured as “*Stories: <story>, Question 1: <question1>*”. Finally, prompts are also formatted to display only one stimulus at a time to avoid order effects. For further prompt details see Appendix A.3.

Prompt Paraphrasing. Prompt characteristics can significantly influence model responses. We therefore introduce controlled perturbations to both the general instruction component and the data format component of the prompt. Sclar et al. (2024) show that formatting influences model responses; we adapt their structured formatting protocol FORMATSREAD to introduce 9 additional data format variations that alter whitespace, casing, ordering, and punctuation. To alter the general instructions, we follow Wahle et al. (2024) who show that paraphrasing can also affect model responses. We apply paraphrasing categories of (1) discourse-based syntax and structural changes and (2) semantic changes to rephrase the general instructions while retaining meaning to create three versions of the general instructions. In all, we have a total of 30 unique prompt variations (3 general prompt × 10 data format) per prompt. By evaluating responses across these 30 variations, we avoid findings that are artifacts of a single prompt phrasing.

Models and Inference. We conduct inference on seven LLMs. Six are open weight models: Gemma2 with 9B and 27B parameters (Riviere et al., 2024), Llama3.1 with 8B and 70B parameters (Dubey et al., 2024), and Qwen2 with 7B and 72B parameters (Yang et al., 2024). One proprietary model is GPT4-o (Achiam et al., 2023).

We use the Hugging Face library (Wolf et al., 2020) for generation with open-source models, applying 4-bit quantization for efficiency when working with models exceeding 9B parameters. We use the OpenAI API to query GPT-4o. Default generation hyperparameters are used in all cases. Responses are parsed based on the format specified in the prompt, such that answers can be automatically extracted for analysis.

5 Experiments and Results

We examine three probes that evaluate various aspects of the narratives described in Section 3: truth value judgments (§5.1), word completion tasks

(§5.2), and open-ended causal questions (§5.3). The experiments query the LLMs based on the two versions (perfective/imperfective event in Cause 1) of the 16 narrative described in §3. All experiments are facilitated by our Expert-in-the-Loop pipeline, meaning prompts are perturbed such that we have 30 versions of each prompt. Finally, we perform statistical analyses on the LLM responses to evaluate our hypotheses. Analyses used multilevel model analysis through the R package lme4 (Bates et al., 2015). Details of these analyses can be found in Appendix A.1.

5.1 Experiment 1a: Truth Value Judgments

To assess whether LLMs understand the semantics of perfective and imperfective aspect in narratives, we probe their ability to infer completion in the Accomplishment events described earlier. Grammatical aspect determines whether an Accomplishment’s final state is explicitly expressed (perfective) or whether the event is left open-ended (imperfective). For instance, the imperfective statement “John was walking home” does not entail the resulting final state “John at home.”

This experiment tests whether LLMs correctly infer an event’s final state by evaluating truth-value judgments on the following two statements: (i) “John at home” (Positive inference), or (ii) “John not at home” (Negative inference). Semantically, the positive final state is true in the perfective, and the negative final state is true in the imperfective.

Semantics: Truth Value Judgments

Story: ...Lena ran down the stairs...

Truth Value Judgment: Lena downstairs

Setup. LLMs are presented with narratives and two final state formulations (positive or negative) associated with the event in Cause 1. LLMs are instructed to “*respond with whether you think the phrase is ‘True’ or ‘False’ or ‘Both’ or ‘Can’t Decide’ with respect to the story,*” where the phrase is the final state. These are adapted from the instructions in Schramm et al. (2022).

Findings. We hypothesized that LLMs, given their strong syntactic and grammatical knowledge, would correctly identify the semantic implicature of grammatical aspect for Accomplishments’ final states. Results (see Table 2) confirm that LLMs

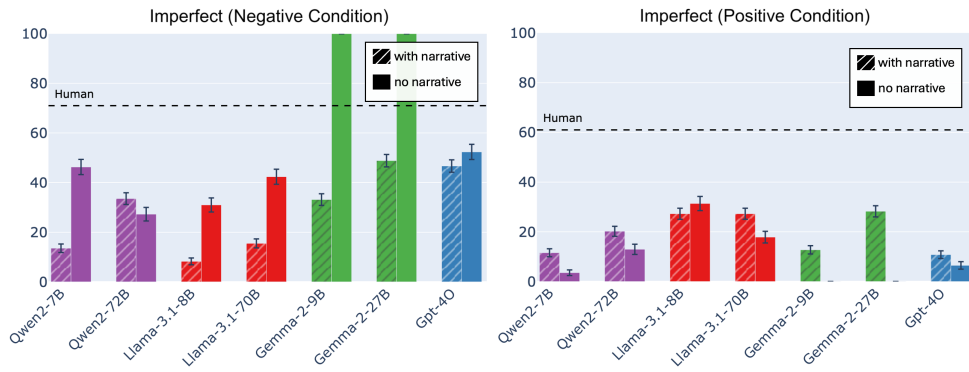


Figure 3: Accuracy in semantic truth-value judgments for events marked with imperfective aspect for LLMs, when the events are embedded within a narrative (shaded bars) versus not. For imperfective events, LLMs have much lower accuracy rates than humans when judging whether the event’s resulting final state is valid. Further, LLMs seem to be heavily affected by the presence or absence of a narrative, especially when judging the negative polarity of final states. Notably, the presence or absence of a narrative changes responses in inconsistent directions. Error bars represent the standard error.

Aspect	Polarity	Target	LLM	Human
Perfect	Positive	True	88%	88%
Imperfect	Negative	True	18%	71%
Perfect	Negative	False	89%	93%
Imperfect	Positive	False	35%	61%

Table 2: LLM semantic understanding of aspect. Narratives include an event in either the imperfective ("Lena was running downstairs") or perfective ("Lena ran downstairs") aspect. LLMs then evaluate the truth of the resulting state with positive or negative polarity ("Lena (not) downstairs"). The target column shows the semantically correct truth value. **Results show that LLMs perform significantly worse than humans in non-prototypical imperfective conditions.** Human data can be directly accessed in [Schramm and Mensink \(2024\)](#).

perform well for perfective events but struggle significantly with non-prototypical aspect (i.e., imperfective). Although human truth-value judgments ([Schramm and Mensink, 2024](#)) also deviate some of the time from the semantically expected responses in non-prototypical imperfective aspect conditions within this narrative context, LLMs exhibit significantly lower accuracy (18% compared to 71% for humans). Notably, the truth-value judgment task is not forced;³ LLMs chose neither true nor false in 9% of responses.

Experiment 1b: Influence of Narrative Context

The narrative context in Experiment 1a may pose a challenge for LLMs since the Accomplishment-imperfective combination is rarely seen in narratives, which could interfere with the LLMs’ ability to display its grammatical knowledge in its

³“Both” or “Can’t Decide” are also options.

response. To address this possibility, we repeat the experiment with the narrative context omitted. However, with no narrative context performance in the imperfective condition remains low (see Figure 3), reinforcing that the difficulty is inherent to aspect processing and not an artifact of the unusual narrative structure. Interestingly, LLMs show both positive and negative shifts in accuracy across models when narratives are removed, suggesting variability in how different architectures comprehend aspect with and without context.

Findings. In summary, our experiments reveal that LLMs have significant sensitivity to narrative context when evaluating events in the imperfective aspect. Regardless of the presence of narrative context, **LLM responses are significantly less accurate for imperfective events** (compared to perfective). The perfective events in narratives are much more common, suggesting that LLMs may rely on high-frequency linguistic patterns more than a deep mastery of aspect semantics in their responses.

This experiment evaluates self-reported interpretations from LLMs, which may not fully capture their understanding of aspect semantics and grammatical knowledge. However, it is notable that their responses in these experiments diverged sharply from humans’ despite the variety of prompts presented in the experiments.

5.2 Experiment 2: Word Completion Task

Do LLMs, like humans, treat aspect in narratives as a temporary signal to eventually retain or encode an event? Aspect in narratives influences tempo-

rary human working memory and permanent event encoding: Prior studies show that when an event is described with imperfective aspect, humans infer its ongoing relevance and are more likely to retain it in working memory (Mozuraitis et al., 2013). To investigate whether LLMs exhibit a similar response, we conduct a word-completion experiment, a widely used measure of concept activation in working memory (Tiggemann et al., 2004) that has also been used in linguistic aspect studies (Schramm and Mensink, 2016).

Word Completion Task

Story: ... she *was eating* the apple ...

Complete the word: A P _ _ _

If an event remains active in working memory, its associated key noun is expected to be retrieved more frequently. The frequency of specific word completions serves as an indicator of **concept accessibility**, reflecting how strongly a word is activated by underlying cognitive processes. Unlike other measures of word activation, such as timed lexical decision tasks, word completion is a feasible assessment tool for LLMs. This task is grounded in the priming effect, a phenomenon in human cognition where exposure to a stimulus influences subsequent recall. Prior research suggests that LLMs exhibit similar priming-based activation patterns, making word completion a viable proxy for assessing concept accessibility in LLMs (Roberts et al., 2024a; Jumelet et al., 2024).⁴

Setup. In addition to a narrative, LLMs are presented with two partial words (e.g., A P _ _ _) and asked to fill in the blanks to create a complete word (e.g., A P P L E). One partial word is a distractor, and the other can be completed with a word from the event in the critical Cause 1 section of the narrative. Crucially, the *location* of the partial word question is also varied to test the extent to which the cognitive processes in LLMs are similar to those in humans’ working memory. In the “Near Cause 1” condition, the word completion probe is located after the sentence with the target word, allowing us to estimate the extent to which the word is activated because the imperfective is nonprototypical at the Cause 1, thus putting the event in

⁴However, because LLM attention mechanisms differ fundamentally from working memory in humans, LLM responses to this task should be interpreted with caution.

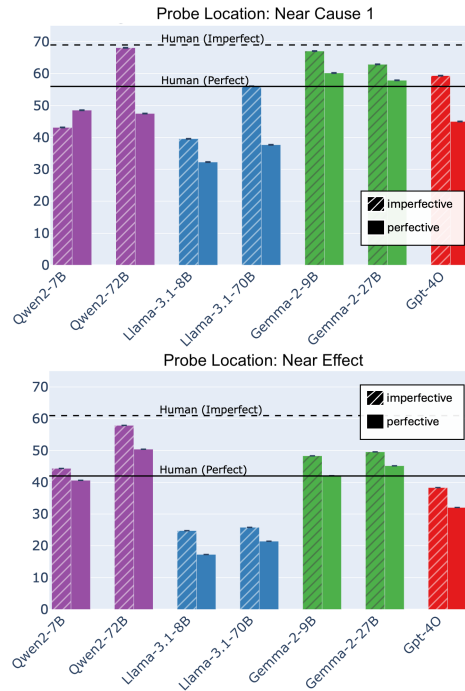


Figure 4: Frequencies at which word completion rates match the target word from Cause 1 across models. Shaded bars are for imperfective aspect in Cause 1. LLM completions have significantly higher match frequencies when the probe directly follows Cause 1 (top) and are reduced after the effect (bottom).

focus. In the “Near Effect” condition, the word completion probe is located after the effect, and we investigate the extent to which the target word is reactivated by the surprise effect at the end of the narrative, as would be expected when Cause 1 has the unbounded imperfective aspect.

Findings. Our results show that LLM responses roughly align with human patterns when the word completion occurs near Cause 1. However, the frequency of responses that point to the target word decrease (on average 33%) when the completion is placed near the Effect, and differences between imperfective and perfective shrink (Figure 4). This suggests that while LLMs may in the short term attend more to pragmatic nonprototypicality information regarding imperfective narrative, they lack **distal causal narrative integration capabilities** akin to human memory mechanisms. Additionally, because imperfective aspect in the earlier foregrounded events is highly non-prototypical, LLMs may be attending to these events due to their statistical rarity, rather than demonstrating an implicit understanding of their pragmatic implications.

5.3 Experiment 3: Open-ended Causal Questioning

As discussed in Section 3, imperfective aspect in narratives carries causal implications, influencing how humans interpret event relationships in episodic memory. In this experiment, we examine whether LLMs’ causal inferences are similarly affected by aspect. Recall the narrative structure described in § 3, in which a surprising effect is introduced, and there are two potential causes: When Cause 1 is presented in the imperfective aspect, it is more likely to be perceived as ongoing, consequently increasing its likelihood of being inferred as the cause. When Cause 1 is presented in the perfective aspect, it is more likely to be perceived as completed, reducing its consideration during causal inference.

Open-ended Causal Questioning

Story: ... Suddenly there was a loud noise.

Question: Why was there a loud noise?

Human studies confirm that imperfective aspect significantly increases the likelihood of Cause 1 being inferred as the event’s cause (Schramm and Mensink, 2016). We assess whether LLMs exhibit similar causal reasoning patterns.

Setup. An open-ended causal question follows the narrative, asking the LLM to infer the most plausible cause of the final effect. We measure how frequently Cause 1 and Cause 2 are identified as causing the effect. Since manual annotation of the open-ended LLM responses is impractical at scale, we employ OpenAI’s API to assist with response classification (see Appendix A.3 for details). To ensure reliability, one author manually coded 128 responses for comparison. We measured inter-annotator agreement between the manual and automated annotations using Cohen’s Kappa and found $\kappa = .93$, indicating strong reliability.

Findings. The results of LLM causal inferences are presented in Figure 5 and compared to previous human results (Schramm and Mensink, 2016). We observe that LLMs, like humans, are also more likely to infer that Cause 1 caused the effect in the imperfective condition (Fig 5). However, most LLMs make this inference substantially less frequently than humans do.

The LLMs that score closest to humans on the

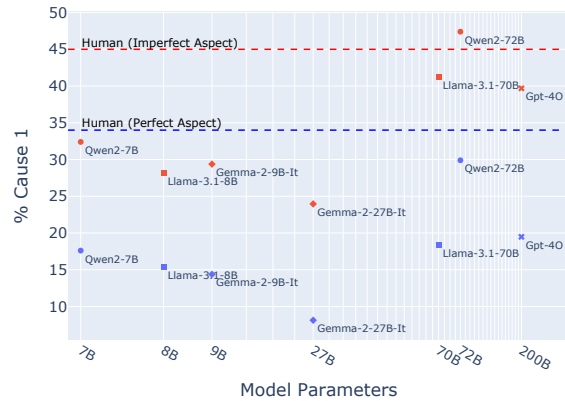


Figure 5: As LLM parameter size increases, there is a trend towards more human-like causal inferences with respect to the Cause 1 event when Cause 1 is in the imperfective. When Cause 1 is in the perfective, LLMs are consistently below human causal inference rates.

open-ended causal inference question are notably not necessarily the highest scores in our other aspect-related experiments (e.g., Llama-3.1-70B does well in this task, but was not close to humans in Experiments 1a, 1b, and 2). Qwen2-72B also has a strong performance, and notably also did very well in the word completion task in the critical “Near Effect” condition, suggesting that the word completion task may be a valuable probe for LLMs. Finally, all LLMs tested were much less likely than humans to offer a *perfective* Cause 1 as their answer. This may again point to an over-reliance on prototypicality, or perhaps to a lack of distal causal capabilities – LLMs tend to lose attention on Cause 1 by the end of the narrative, as evidenced by the “Near Effect” condition results in Experiment 2.

6 Discussion

Our findings indicate that LLMs process aspect in narratives fundamentally differently from humans. While LLM truth value judgments in Experiment 1 mostly align with humans in prototypical cases, they deviate sharply in non-prototypical contexts. Interestingly, despite the poor truth-value judgments in non-prototypical cases, LLMs can readily articulate the definition of aspect in academic discussions. We speculate that aspect is one of several areas in which LLMs excel in declarative knowledge (e.g., reciting a definition) while failing in implicit application (e.g., making truth-value judgments following the definition). This contrasts sharply with humans, who may lack declarative knowledge of aspect but consistently demonstrate

implicit understanding in real-world judgments. Future research should explore how the tension between declarative and implicit knowledge extends to other linguistic and cognitive domains.

In Experiment 2, we observe that while LLMs seem to detect pragmatic non-prototypicalities within immediate contexts (“near Cause 1”), **LLMs tend to not maintain causal pragmatic focus on imperfective events** over the distance of the narrative (“near Effect”). If LLMs constructed temporary internal situation models as humans do, we would expect a more maintained focus. This raises broader questions about how LLMs track narrative elements, such as point of view, and temporal progression.

Experiment 3 also explores whether LLMs grasp the causal implications of aspectual distinctions by probing the contents of the narrative situation model. Our results suggest that larger models exhibit more human-like, albeit less robust, causal reasoning, even though in previous experiments larger models showed no clear improvements in aspect-based temporary pragmatic focus or semantic aspect comprehension. However, we see that while a substantial portion of humans (one-third) ascribe the event in Cause 1 to the effect even when Cause 1 is in the perfective, LLMs are significantly less likely to do so. This suggests that humans approach narrative interpretation with **more flexibility**, striving for coherence and constructing internal situation models in ways that LLMs do not.

Pragmatic Understanding. Overall linguistically, the performance of LLMs appears to be due to a lack of pragmatic context-level understanding (Beuls and Van Eecke, 2024). The semantic exploration of truth value conditions connected with prototypical and non-prototypical pairings of Accomplishment events possessing final states with bounded perfective aspect and unbounded imperfective aspect suggests that LLMs represent aspect distributionally, rather than based on the concepts expressed by aspect, similar to learners of English (Salaberry, 2024). The accuracy rates for LLMs are high in the bounded perfective when prototypically paired with bounded Accomplishment events with a final state, and they are low when the events are non-prototypically combined with unbounded imperfective aspect.

It is thus not surprising that LLMs also appeared to differ from humans when we compare their respective cognitive processes involved in the

understanding of (non)prototypicality and causality, echoing recent findings (e.g., Apidianaki et al. (2024) and Ohmer et al. (2024)). Although probes in close proximity to Cause 1 have accuracy rates in the right direction, the rates are lower for probes in close proximity to the final effect, and comprehension of aspectual meanings and the distal inference of causality largely failed. In the eventual cognitive situation model, shortcomings were also apparent. LLMs’ pragmatic understanding is generally lower than humans’, and the lack of a connection between cause and effect in the perfective aspect is exaggerated. This may be due to the distributional (rather than meaning-based) prototypical aspect understanding found on the semantic level. Speculatively, the near-human understanding of the imperfective in the situation model may be a “compensatory over-emphasis” to the non-prototypicality of this aspect rather than human-like behavior, since the semantic accuracy was fairly low.

Effects of Model Type and Size. We compare three state-of-the-art model families (Gemma, Llama, Qwen) and one proprietary model (GPT-4o), evaluating both small and large model sizes to examine differences in family architecture and parameter scaling. While model families show statistically significant differences across tasks, no model or model family demonstrates consistent mastery or failure in aspect-related processing. This suggests that LLMs do not process aspect using human-like mechanisms. Moreover, there is no clear relationship between model performance on structured aspectual tasks and their performance on open-ended causal inference tasks, indicating that **aspect comprehension and causal reasoning may be processed separately in LLMs**. Interestingly, we only observe a clear distinction between larger and smaller models in the open-ended causal inference experiment: larger models demonstrate a stronger grasp of imperfective causal implications. It remains uncertain whether further scaling will bridge the gap between LLMs and human-like aspectual processing.

Conclusion Through a series of targeted experiments, we assess LLMs’ cognitive capabilities in processing linguistic aspect. Our findings indicate that LLMs over-rely on prototypicality, exhibit inconsistent aspectual judgments, and often fall short of performing human-like aspect-mediated causal inference, raising concerns about their ability to fully comprehend narratives.

7 Limitations

This work relies on “self”-report from large language models in response to prompts designed to probe specific functionalities. The epistemology of LLM responses, i.e., how well LLM self-report reflects internal states, is unclear. We make efforts to mitigate this risk by conducting multiple experiments, paraphrasing our prompts, and measuring multiple signals from the responses.

This work also evaluates LLMs on narrative comprehension in controlled experimental states and without interaction. It is always possible that with more dialogue, or with different prompt engineering, results would change. However, if LLM comprehension of a topic is robust, it should not demand elaborate prompt engineering. With this in mind, we believe our results are informative, although they should be contextualized within the specific parameters of our experiments.

8 Ethics

We note that the theoretical framework for LLM probing discussed in this work, if misused or misapplied, may lead to incorrect conclusions about the nature of LLMs and LLM cognition. This can be misleading and in rare cases causes personal upset and concern (Allyn, 2022). We therefore maintain that the framework used in this work is best applied responsibly by researchers working in interdisciplinary teams.

Acknowledgments

We are especially grateful to Dr. Michael Mensink and Ahn Thu Tong for making the understanding of human data reported in this work possible. This research was supported in part by the Doctoral Dissertation Fellowship from the University of Minnesota for the first author.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.

Bobby Allyn. 2022. [The Google engineer who sees company’s AI as ‘sentient’ thinks a chatbot has a soul](https://www.npr.org/2022/06/16/1105552435/google-ai-sentient). <https://www.npr.org/2022/06/16/1105552435/google-ai-sentient>. Accessed: 2025-02-01.

Roger W Andersen. 1996. Primacy of aspect in first and second language acquisition: The pidgin/creole connection. *Handbook of second language acquisition*, pages 527–570.

John Robert Anderson. 1974. Retrieval of propositional information from long-term memory. *Cognitive psychology*, 6(4):451–474.

Marianna Apidianaki, Abdellah Fourtassi, and Sebastian Padó. 2024. Language learning, representation, and processing in humans and machines: Introduction to the special issue.

Kathleen Bardovi-Harlig. 2000. Tense and aspect in second language acquisition: Form, meaning, and use. *Language Learning: A Journal of Research in Language Studies*, 50:1.

Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. [Fitting linear mixed-effects models using lme4](#). *Journal of Statistical Software*, 67(1):1–48.

Maxim Bazhukov, Ekaterina Voloshina, Sergey Pletenev, Arseny Anisimov, Oleg Serikov, and Svetlana Toldova. 2024. [Of models and men: Probing neural networks for agreement attraction with psycholinguistic data](#). In *Proceedings of the 28th Conference on Computational Natural Language Learning*, pages 280–290, Miami, FL, USA. Association for Computational Linguistics.

Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.

Katrien Beuls and Paul Van Eecke. 2024. Humans learn language from situated communicative interactions. what about machines? *Computational Linguistics*, 50(4):1277–1311.

Marcel Binz and Eric Schulz. 2023. Using cognitive psychology to understand GPT-3. *Proceedings of the National Academy of Sciences*, 120(6):e2218523120.

Julian Coda-Forno, Marcel Binz, Jane X Wang, and Eric Schulz. 2024. [CogBench: a large language model walks into a psychology lab](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 9076–9108. PMLR.

Julian Coda-Forno, Kristin Witte, Akshay K Jagadish, Marcel Binz, Zeynep Akata, and Eric Schulz. Inducing anxiety in large language models increases exploration and bias.

Debarati Das, Karin De Langis, Anna Martin-Boyle, Jaehyung Kim, Minhwa Lee, Zae Myung Kim, Shirley Anugrah Hayati, Risako Owan, Bin Hu, Ritik Parkar, Ryan Koo, Jonginn Park, Aahan Tyagi, Libby Ferland, Sanjali Roy, Vincent Liu, and

- Dongyeop Kang. 2024. [Under the surface: Tracking the artifactuality of llm-generated data](#). *Preprint*, arXiv:2401.14698.
- Ishita Dasgupta, Andrew K Lampinen, Stephanie CY Chan, Antonia Creswell, Dharshan Kumaran, James L McClelland, and Felix Hill. 2022. Language models show human-like content effects on reasoning. *arXiv e-prints*, pages arXiv–2207.
- Karin de Langis, Jong Inn Park, Bin Hu, Khanh Chi Le, Andreas Schramm, Michael C Mensink, Andrew Elfenbein, and Dongyeop Kang. 2025. A framework for robust cognitive evaluation of llms. *arXiv preprint arXiv:2504.02789*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The LLaMA 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Jessica Maria Echterhoff, Yao Liu, Abeer Alessa, Julian McAuley, and Zexue He. 2024. [Cognitive bias in decision-making with LLMs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 12640–12653, Miami, Florida, USA. Association for Computational Linguistics.
- Thilo Hagendorff, Sarah Fabi, and Michal Kosinski. 2023. Human-like intuitive behavior and reasoning biases emerged in large language models but disappeared in ChatGPT. *Nature Computational Science*, 3(10):833–838.
- William Hart and Dolores Albarracín. 2009. The effects of chronic achievement motivation and achievement primes on the activation of achievement and fun goals. *Journal of personality and social psychology*, 97(6):1129.
- Paul J Hopper. 1982. Aspect between discourse and grammar: An introductory essay for the volume. *Tense-aspect: Between semantics and pragmatics*, pages 3–18.
- Anna A Ivanova. 2025. How to evaluate the cognitive abilities of llms. *Nature Human Behaviour*, pages 1–4.
- Erik Jones and Jacob Steinhardt. 2022. Capturing failures of large language models via human cognitive biases. *Advances in Neural Information Processing Systems*, 35:11785–11799.
- Nitish Joshi, Abulhair Saparov, Yixin Wang, and He He. 2024. [LLMs are prone to fallacies in causal inference](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 10553–10569, Miami, Florida, USA. Association for Computational Linguistics.
- Jaap Jumelet, Willem Zuidema, and Arabella Sinclair. 2024. Do language models exhibit human-like structural priming effects? *arXiv preprint arXiv:2406.04847*.
- Chani Jung, Dongkwan Kim, Jiho Jin, Jiseon Kim, Yeon Seonwoo, Yejin Choi, Alice Oh, and Hyunwoo Kim. 2024. [Perceptions to beliefs: Exploring precursory inferences for theory of mind in large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 19794–19809, Miami, Florida, USA. Association for Computational Linguistics.
- Wolfgang Klein. 1994. *Time in language*. Routledge.
- Yash Kumar Lal, Vanya Cohen, Nathanael Chambers, Niranjan Balasubramanian, and Ray Mooney. 2024. [CaT-bench: Benchmarking language model understanding of causal and temporal dependencies in plans](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 19336–19354, Miami, Florida, USA. Association for Computational Linguistics.
- Bruce Lee and Jaehyuk Lim. 2024. [Language models don’t learn the physical manifestation of language](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3554–3579, Bangkok, Thailand. Association for Computational Linguistics.
- Eun-Kyoung Rosa Lee, Sathvik Nair, and Naomi Feldman. 2024. [A psycholinguistic evaluation of language models’ sensitivity to argument roles](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3262–3274, Miami, Florida, USA. Association for Computational Linguistics.
- Russell V. Lenth. 2025. *emmeans: Estimated Marginal Means, aka Least-Squares Means*. R package version 1.10.7-100001, <https://rvlenth.github.io/emmeans/>.
- Zihao Li. 2023. The dark side of ChatGPT: Legal and ethical challenges from stochastic parrots and hallucination. *arXiv preprint arXiv:2304.14347*.
- Sheng Lu, Irina Bigoulaeva, Rachneet Sachdeva, Harish Tayyar Madabushi, and Iryna Gurevych. 2024. [Are emergent abilities in large language models just in-context learning?](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5098–5139, Bangkok, Thailand. Association for Computational Linguistics.
- Steven G Luke. 2017. Evaluating significance in linear mixed-effects models in r. *Behavior research methods*, 49:1494–1502.
- Yaojia Lv, Haojie Pan, Zekun Wang, Jiafeng Liang, Yuanxing Liu, Ruiji Fu, Ming Liu, Zhongyuan Wang, and Bing Qin. 2024. [CogGPT: Unleashing the power of cognitive dynamics on large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 6074–6091, Miami, Florida, USA. Association for Computational Linguistics.
- Joseph P Magliano and Michelle C Schleich. 2000. Verb aspect and situation models. *Discourse processes*, 29(2):83–112.

- Negar Maleki, Balaji Padmanabhan, and Kaushik Dutta. 2024. AI hallucinations: a misnomer worth clarifying. In *2024 IEEE conference on artificial intelligence (CAI)*, pages 133–138. IEEE.
- Mindaugas Mozuraitis, Craig G Chambers, and Meredyth Daneman. 2013. Younger and older adults’ use of verb aspect and world knowledge in the online interpretation of discourse. *Discourse processes*, 50(1):1–22.
- Xenia Ohmer, Elia Bruni, and Dieuwke Hupke. 2024. From form (s) to meaning: Probing the semantic depths of language models using multisense consistency. *Computational Linguistics*, 50(4):1507–1556.
- Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Jesse Roberts, Kyle Moore, Douglas Fisher, Oseremhen Ewaleifoh, and Thao Pham. 2024a. Large language model recall uncertainty is modulated by the fan effect. In *Proceedings of the 28th Conference on Computational Natural Language Learning*, pages 303–313.
- Jesse Roberts, Kyle Moore, Drew Wilenzick, and Douglas Fisher. 2024b. Using artificial populations to study psychological phenomena in neural models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18906–18914.
- M Rafael Salaberry. 2024. Thirty years of studying lexical aspect in L2 acquisition: Moving towards an expansive research agenda. *Ampersand*, 13:100185.
- Andreas Schramm, Verena Haser, Michael C Mensink, Jonas Reifenrath, and Parinaz Kassemi. 2022. Implicit textually enhanced processing of aspectual meanings in english learners with german as a first language. *Discourse Processes*, 59(7):520–552.
- Andreas Schramm and Michael C Mensink. 2016. Processing of aspectual meanings by non-native and native english speakers during narrative comprehension. *New approaches to English linguistics: Building bridges*, pages 251–280.
- Andreas Schramm and Michael C. Mensink. 2024. Intentional explicit meaning-based learning of aspectual meanings. <https://osf.io/7ksdq>.
- Andreas Hans Schramm. 1998. *Aspect and causal inferences: The role of temporal coherence in a process model of inference generation in text comprehension*. University of Minnesota.
- 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*.
- Raj Sanjay Shah, Khushi Bhardwaj, and Sashank Varma. 2024. *Development of cognitive intelligence in pre-trained language models*. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9632–9657, Miami, Florida, USA. Association for Computational Linguistics.
- Settaluri Sravanthi, Meet Doshi, Pavan Tankala, Rudra Murthy, Raj Dabre, and Pushpak Bhattacharyya. 2024. PUB: A Pragmatics Understanding Benchmark for Assessing LLMs’ Pragmatics Capabilities. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 12075–12097.
- Marika Tiggemann, Duane Hargreaves, Janet Polivy, and Traci McFarlane. 2004. A word-stem completion task to assess implicit processing of appearance-related information. *Journal of Psychosomatic Research*, 57(1):73–78.
- Jen tse Huang, Wenxuan Wang, Eric John Li, Man Ho LAM, Shujie Ren, Youliang Yuan, Wenxiang Jiao, Zhaopeng Tu, and Michael Lyu. 2024. *On the humanity of conversational AI: Evaluating the psychological portrayal of LLMs*. In *The Twelfth International Conference on Learning Representations*.
- Tomer Ullman. 2023. Large language models fail on trivial alterations to theory-of-mind tasks. *arXiv preprint arXiv:2302.08399*.
- Jan Philip Wahle, Terry Ruas, Yang Xu, and Bela Gipp. 2024. *Paraphrase types elicit prompt engineering capabilities*. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11004–11033, Miami, Florida, USA. Association for Computational Linguistics.
- 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, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 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, et al. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Jiahao Ying, Yixin Cao, Kai Xiong, Long Cui, Yidong He, and Yongbin Liu. 2024. *Intuitive or Dependent? Investigating LLMs’ Behavior Style to Conflicting Prompts*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4221–4246, Bangkok, Thailand. Association for Computational Linguistics.

Rolf A Zwaan, Mark C Langston, and Arthur C Graesser. 1995. The construction of situation models in narrative comprehension: An event-indexing model. *Psychological science*, 6(5):292–297.

A Appendix

A.1 Statistical Analyses

We conduct statistical analyses on experimental results to assess the significance of the experimental manipulations, and to control for random effects inherent in the individual narratives. In our analyses we set $\alpha = .01$. We use the lme4 package (Bates et al., 2015) in R to perform analyses and the emmeans package (Lenth, 2025) to calculate the estimated marginal means. F values are calculated using Satterthwaite’s method (Luke, 2017).

A.1.1 Experiment 1: Truth Value Judgments

This experiment has two independent variables, aspect and polarity of final state (positive or negative). To model LLM responses, we define aspect and polarity as fixed effects, and narrative type as a random effect. LLM responses are coded to binary variable indicating correctness. Our model is:

$$\text{accuracy} \sim \text{aspect} * \text{polarity} + (1 \mid \text{narrative})$$

We find a significant main effect of aspect, $F = 66.5, p < .01$, in which imperfective has lower LLM accuracy than perfective ($M = .41, SE = .02$ for imperfective, $M = .46, SE = .02$ for perfective). There is also a significant main effect of polarity, $F = 10363, p < .01$ with performance being much higher in the positive than in the negative condition ($M = .76, SE = .02$ for positive and $M = .12, SE = .02$ for negative) and a significant aspect and polarity interaction. The interaction between the two is also significant at $F = 661.5, p < .01$. All pairwise comparison between aspect and polarity condition combinations were significant (Bonferonni corrections for multiple comparisons were applied).

A.1.2 Experiment 2

This experiment has two independent variables, aspect (perfective or imperfective) and the probe location in the prompt (early or late). To model LLM responses, we define aspect and location as fixed effects, and narrative type as a random effect. LLM responses are coded to binary variable indicating presence or absence of the target word in the response. The model is:

$$\text{response} \sim \text{aspect} * \text{location} + (1 \mid \text{narrative})$$

We find a significant main effect of aspect, $F = 92.4, p < .01$, in which imperfective has lower accuracy than perfective ($M = .41, SE = .04$ for imperfective, $M = .49, SE = .04$ for perfective). There is also a significant main effect of probe location, $F = 282.2, p < .01$ with responses happening much more frequently in the early than in the late condition ($M = .52, SE = .04$ for early and $M = .39, SE = .04$ for late). The interaction between the two approaches significance $F = 5.9, p = .02$. All pairwise comparison between aspect and probe location condition combinations were significant (Bonferonni corrections for multiple comparisons were applied).

A.1.3 Experiment 3

This experiment has two independent variables, aspect and probe location. We define the same model as in Experiment 2.

We find a significant main effect of aspect, $F = 98.5, p < .01$, in which imperfective has lower accuracy than perfective ($M = .35, SE = .06$ for imperfective, $M = .17, SE = .06$ for perfective). Although we have no hypothesis for probe location, there is a significant main effect of probe location, $F = 25.3, p < .01$ with responses happening much more frequently in the late condition than in the early condition ($M = .22, SE = .30$ for early and $M = .39, SE = .04$ for late). The interaction between the two approaches significance at $F = 2.4, p = .02$.

A.2 Web Application

An overview of the highlights of our web application for collaborating with cognitive scientists is shown in Figure 7. Further details can be found in de Langis et al. (2025). Our collaboration pipeline is generalizable, allowing users to upload generic stimulus files that are the basis of the remainder of the pipeline. Users can select which parts of the stimulus files should be shown to the LLM during experiments and specify the independent variables and resulting groups that should be analyzed.

A.3 Prompts

The original prompts (i.e., before paraphrasing) are included for reference. The prompts for Experiment 1 are in Tables 5 and 6. The prompt for Experiments 2 and 3, word completion and causal question inference tasks, is in Table 3.

Evaluating open-ended causal question responses The prompt used in the automatic scor-




Figure 6: We assume as little structure as possible for experimental datasets to allow for generalizability to other domains. Datasets consist of multiple *groups* of stimuli, which have different independent variable values (users can indicate which *fields* are independent variables). Human studies often compare metrics across different stimuli groups to draw conclusions about the effects of independent variables.

ing of the LLM responses to the causal inference questions in Experiment 3 is shown in Table 4.

Experiment 1: Truth value judgment prompts

Experiment 1 used truth value judgment (TVJ) prompts used to probe semantic understanding of aspect in LLMs are in Tables 6 and 5.

A.4 Application Usage

Based on standard experimental procedures in human studies, we structure our expected experiment dataset format such that each row contains an individual stimulus and any number of fields describing that stimulus. Users can then identify which fields correspond to the *independent variables* of their study and define *groups* for their dataset based on those independent variables. (Dataset structure is illustrated in Figure 6.) Finally, users can specify *predictions* – what *dependent variable* should be measured, and how that measure is expected to vary between groups. Users also indicate the dependent variable(s) their study should measure in the LLM response: frequency of responses matching a target value, token probabilities of a target response, or average numerical value of a response (e.g., average rating for experiments that ask for numerical ratings). Instructions are modified for the language model.

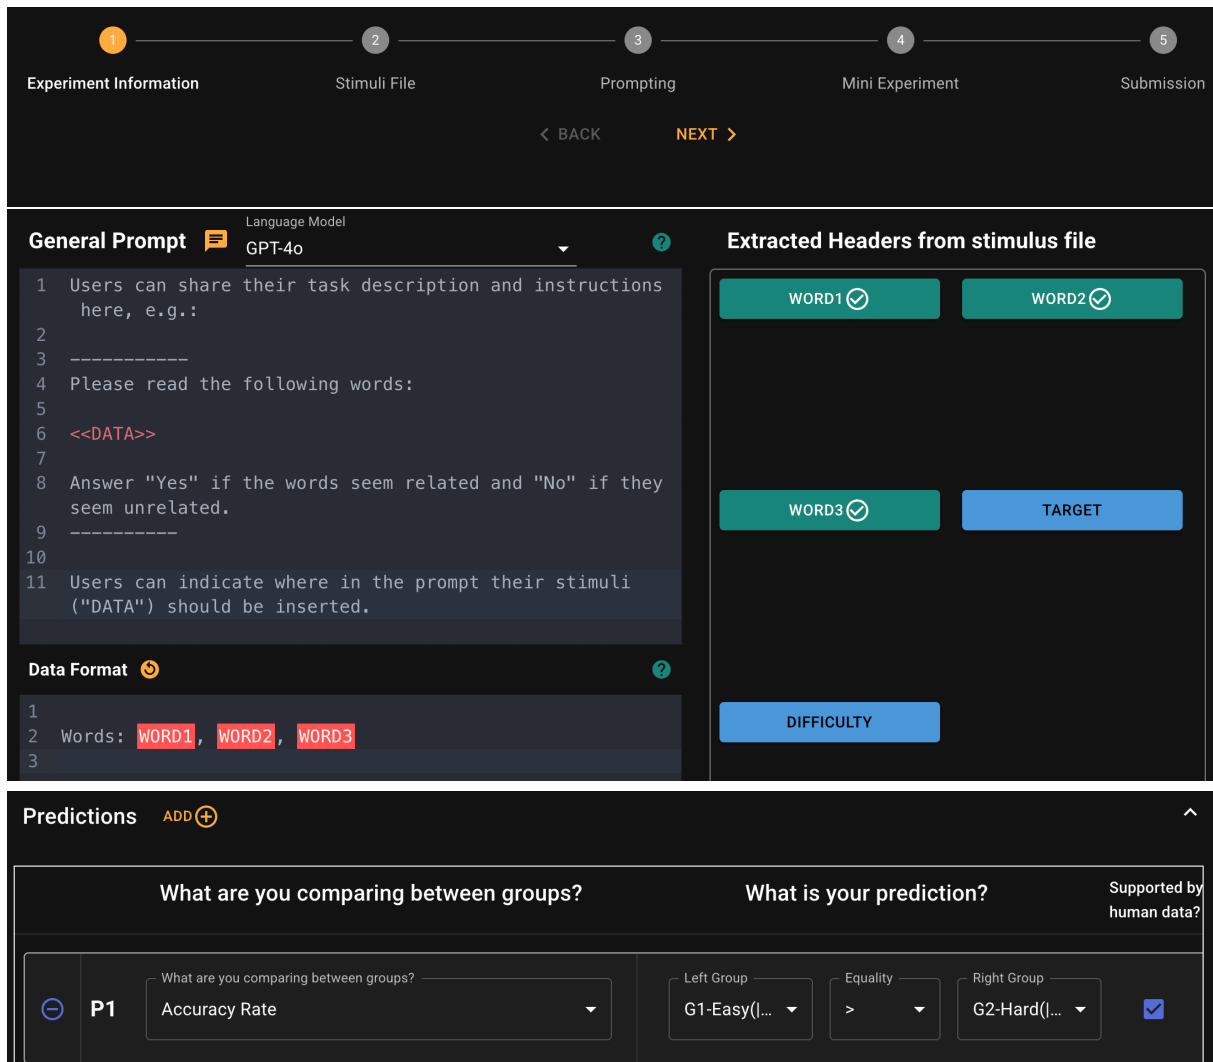


Figure 7: Important components of our web application for collaborating with cognitive scientists. **Top:** A navigation bar showing the pipeline that cognitive scientists complete to submit an experiment to the pipeline. **Middle:** Cognitive scientists can share task instructions and select which columns from their uploaded stimulus file should be included in the prompt (we then post-process the prompts and create 30 paraphrased versions). **Bottom:** Cognitive scientists identify independent variables from their uploaded stimulus file and define groups based on these independent variables.

General Instructions:

This study evaluates your ability to handle two simultaneous tasks:

- Task 1: Read and understand stories well enough to answer comprehension questions correctly.
- Task 2: Occasionally complete "word edges" as quickly as possible.

What is a "word edge"?

- A word edge consists of one or two starting letters of a word, followed by blanks. Your task is to fill in the blanks to form a valid English word.
- Complete the word with the first valid word that comes to mind. The word can be a compound word (e.g., SETUP or LUNCHBOX) or include endings (e.g., STICKS or BUSHES).
- Important:
 - Fill all blanks (no extra letters or fewer letters than the blanks provided).
 - Do not use personal names.

Example of a Word Edge:

- BR _ _ _ _ _
 - Correct answers:
 - BRACKET
 - BREAKIN
 - BRUSHES
 - Incorrect answers:
 - BREAKAGE (too many letters)
 - BRIDG_ (too few letters)
 - BRIDGIT (personal name)

Instructions for the task:

1. Read the stories provided.
2. Occasionally, you will encounter two partial word edges.
3. Some stories are split into parts, and the word edges may appear in the middle.
4. Other stories are shown in full, with word edges at the end.
5. Fill in the blanks for the first partial word that comes to mind.
6. Then, complete the second word edge.
7. Continue reading the story.
8. After each story, you will be asked a question to test your memory of its content.

Answer Format:

- Question 1:
{WORD1}
{WORD2}
- Question 1:
{ANSWER}

Note: Do not add any explanations or repeat stories/questions. Strictly follow the Answer Format provided.

Story Part 1: {{STORY PART 1}}

Question 1:
{{QUESTION 1}}

Story Part 2: {{STORY PART 2}}

Question 2: {{QUESTION 2}}

Table 3: Original prompt (prior to prompt perturbation) used for word completion and causal inference experiments.

You are evaluating the answer to a question about a story. You will see the story, an extracted sentence from the story, and the question and answer.

The extracted sentence will either be imperfect or perfect tense. We call this grammatical aspect, and it can be interpreted to have temporal meaning for the story (the aspectual meaning). The question will be about what *caused* an outcome at the end of the story.

Your job is to decide whether the answer refers to information from the extracted sentence (true/false). If the answer refers to the aspectual temporal meanings of the extracted sentence's verb or to its non-completion, respond with True. If the answer does not refer to information from the extracted sentence or indicated that that the verb has been completed, respond with False.

Think only within the given materials, and regardless of whether you agree or disagree with the answer, determine whether the answer refers to information from the extracted sentence. Provide a True or False response and explain your rationale clearly. Do not evaluate whether the answerer is right or wrong. Accept the answer as valid, and if the answer refers to the aspectual meanings of the extracted sentence's verb or to its non-completion, respond with True. If the answer does not refer to information from the extracted sentence or indicated that that the verb has been completed, respond with False.

```
[Story]
{{full_story}}

[Extracted Sentence]
{{extracted_sentence}}

[Question]
{{question}}

[Answer]
{{answer}}

[ExtractedConsideredCause (true/false)]:
rationale:
```

Table 4: Prompt that was used to automatically score responses to all LLM causal question responses via the OpenAI API.

This is a test designed to test your ability to manage two different tasks at once: Task 1: to read and comprehend about a story well enough to evaluate a phrase that may refer to the story, and Task 2: to assess how accurate your evaluation of the phrase is. The phrases are 2-5 words long and are not complete sentences. You should decide *whether the phrase is true* with respect to the story. You should respond with "True" or "False" or "Both" or "Can't Decide" according to the first understanding that comes to mind that fits the story. Specifically, you'll first read a story. When you get to the end of a story, the last sentence will repeat. Below this last sentence is a phrase that may refer back to the story. Immediately respond with whether you think the phrase is "True" or "False" or "Both" or "Can't Decide" *with respect to the story.* Make sure your response contains ONLY your truth value judgment (True/False/Both/Can't Decide). Here's an example:

STORY: Nancy and Chris were moving into a new apartment. They decided that Nancy would get the big bedroom and Chris would get the garage space. Nancy and Chris decided to move their furniture into their rooms first. When Chris left with the truck to get the furniture, Nancy PAINTED her wall.

LAST SENTENCE: Nancy painted her wall

PHRASE: whole wall with fresh paint

Is this phrase true with respect to the story?

Option 1 - True

Option 2 - False

Option 3 - Both

Option 4 - Can't Decide

Response: True

STORY: {{STORY TO INCLUDE IN PROMPT}}

LAST SENTENCE: {{LAST SENTENCE}}

PHRASE: {{PHRASE}}

{{OPTIONS}}

Response:

Table 5: Prompt that was used for Experiment 1a.

This is a task designed to test your ability to manage two different tasks at once: Task 1: to read and comprehend a sentence, and Task 2: to assess how accurate your evaluation of a phrase is. The phrases are 2-5 words long and are not complete sentences. You should decide *whether the phrase is true* with respect to the original sentence. You should respond with "True" or "False" or "Both" or "Can't Decide" according to the first understanding that comes to mind. Make sure your response contains ONLY your truth value judgment (True/False/Both/Can't Decide). Here's an example:

SENTENCE: Nancy painted her wall

PHRASE: whole wall with fresh paint

Is this phrase true with respect to the sentence?

Option 1 - True

Option 2 - False

Option 3 - Both

Option 4 - Can't Decide

Response: True

SENTENCE: {{SENTENCE}}

PHRASE: {{PHRASE}}

{{OPTIONS}}

Response:

Table 6: Prompt that was used for TVJ semantic experiment.