

Who Relies More on World Knowledge and Bias for Syntactic Ambiguity Resolution: Humans or LLMs?

So Young Lee[†], Russell Scheinberg[◇], Amber Shore[◇], Ameeta Agrawal[◇]

[†]Miami University, USA

[◇]Portland State University, USA

soyoung.lee@miamioh.edu

{rschein2, ashore, ameeta}@pdx.edu

Abstract

This study explores how recent large language models (LLMs) navigate relative clause attachment ambiguity and use world knowledge biases for disambiguation in six typologically diverse languages: English, Chinese, Japanese, Korean, Russian, and Spanish. We describe the process of creating a novel dataset – MultiWho¹ – for fine-grained evaluation of relative clause attachment preferences in ambiguous and unambiguous contexts. Our experiments with three LLMs indicate that, contrary to humans, LLMs consistently exhibit a preference for local attachment, displaying limited responsiveness to syntactic variations or language-specific attachment patterns. Although LLMs performed well in unambiguous cases, they rigidly prioritized world knowledge biases, lacking the flexibility of human language processing. These findings highlight the need for more diverse, pragmatically nuanced multilingual training to improve LLMs’ handling of complex structures and human-like comprehension.

1 Introduction

Natural language is inherently ambiguous, with single expressions often having multiple interpretations. This ambiguity poses significant challenges to both human cognition and computational models, especially in tasks requiring precise language understanding like machine translation, question answering, and dialogue systems. Miscommunication can arise when different listeners or readers interpret the same expression differently, making ambiguity resolution a critical area of research (Mehrabi et al., 2023; He et al., 2024; Niwa and Iso, 2024; Hatami et al., 2022; Tran et al., 2022; Futeral et al., 2023; Nath et al., 2024).

Among various types of ambiguity, this study focuses on syntactic ambiguity, specifically relative

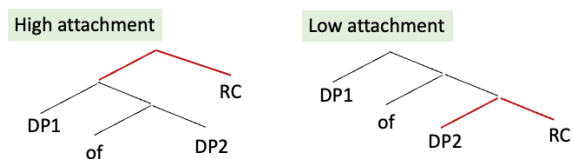


Figure 1: Syntactic Structures of DP1 Modification (left) and DP2 Modification (right) in English

clause (RC) attachment ambiguity. Syntactic ambiguity occurs when a sentence’s structure allows for multiple grammatical interpretations. RC attachment ambiguity arises when a relative clause can attach to more than one determiner phrase (DP), leading to different possible meanings. For instance, in (1), the RC *who had a beard* could refer to either the local DP (DP2) *the man* or the non-local DP (DP1) *the son*.

- (1) The doctor met the son_{DP1} of the man_{DP2} [who had a beard]_{RC}.
- DP2 modification** (linearly local, structurally low attachment (LA)): The person who had a beard is the man.
 - DP1 modification** (linearly non-local, structurally high attachment (HA)): The person who had a beard is the son.

Figure 1 illustrates this syntactic ambiguity with two distinct syntax trees. Previous studies have demonstrated that attachment preferences are language-specific, influenced by multiple factors in the resolution of syntactic ambiguities (Clifton Jr et al., 2003; Fodor, 1998; Grillo and Costa, 2014; Hemforth et al., 1996; Traxler et al., 1998). While world knowledge and biases often guide interpretation, the extent to which these factors override inherent syntactic parsing preferences remains underexplored — especially in large language models (LLMs) (OpenAI et al., 2024; Touvron et al., 2023). Moreover, much of the research on LLMs and am-

¹Dataset available at <https://github.com/PortNLP/MultiWHO>.

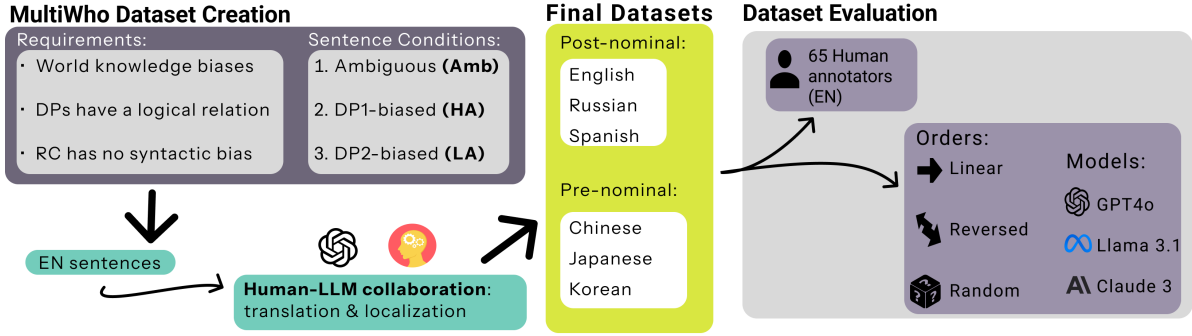


Figure 2: **MultiWho Dataset**: The dataset creation started with a list of **requirements** and three different **conditions**. Using a collaborative **human-LLM** process, we started with developing English sentences and continued through translation and localization, resulting in a multilingual dataset across **six languages**. While not all sentences are pragmatically equivalent in all languages, they are structurally equivalent with regard to our requirements. These datasets were **evaluated** in two ways: the English dataset was evaluated by 65 human annotators for ambiguity/DP-bias, and all 6 datasets were evaluated for ambiguity/DP-bias in three different answer order settings by LLMs.

biguity resolution has focused on English and other Indo-European languages, leaving languages with different syntactic structures less examined (Yuan et al., 2023; Cai et al., 2024).

This study explores how humans and LLMs resolve RC attachment ambiguities across languages and whether world knowledge overrides syntactic preferences.² We investigate if LLMs show consistent attachment patterns across languages and how they align with human processing. We selected six languages—English (EN), Chinese (ZH), Japanese (JP), Korean (KO), Russian (RU), and Spanish (ES)—for their syntactic diversity, particularly regarding the position of RCs. For instance, EN, ES, and RU use postnominal RCs that follow the DP they modify, as shown in (2):

- (2) DP RC: (English)
the man [who ran the marathon]

In contrast, JP, KO, and ZH use prenominal RCs that precede the DP they modify, as illustrated in the KO example (3):

- (3) RC DP: (Korean)
[마라톤을 땀] 남자
[marathon ran] man

Figure 2 presents a schematic figure of the components of this study. Our research addresses the following questions.

²As defined in the previous studies (Kecskes, 2023; Sayeed and Zarcone), world knowledge and bias, as used in our study, refer to societal norms, plausibility, and stereotypes. These factors shape predictions but are distinct from semantic cues, which are usually narrowly defined in linguistic contexts (e.g., animacy or thematic roles).

1. What are the syntactic preferences of LLMs to resolve the ambiguities in RC attachment in different languages? Will there be a certain pattern depending on the syntactic differences among languages?
2. How does the sensitivity of LLMs to world knowledge and biases compare to that of humans when resolving syntactic ambiguities?
3. Does the presentation order (linear, reversed, random) of the possible response choices (DP1 or DP2) influence LLM attachment preferences?

To briefly summarize our results, in ambiguous cases in English, we found that humans exhibit a strong Low Attachment (LA) preference, with an HA (High Attachment) response rate of approximately 0.2, while LLMs show an even stronger LA preference of around 0.1. Extending this observation to other languages, LLMs default to LA preference as well, aligning with human preferences in EN and ZH but contrasting with the HA preferences observed in humans for KO, JP, RU, and ES. In unambiguous EN cases, humans demonstrate a stronger and more consistent adherence to DP2-bias (0.88) compared to DP1-bias (0.64). This reflects a natural tendency toward LA structures and showcases their flexibility in adjusting interpretations despite inconsistencies in world knowledge. In contrast, LLMs display near-perfect accuracy for both DP2- and DP1-biases, suggesting a greater sensitivity to world knowledge biases and a rigidity in interpretations that may reinforce existing social stereotypes. Similarly, in unambiguous cases across other languages, LLMs show high accuracy

Condition	Example Sentence
Ambiguous	The doctor met the son of the man who had a beard.
DP1 Biased	The doctor met the son of the woman who had a beard.
DP2 Biased	The doctor met the daughter of the man who had a beard.

Table 1: Example set of English stimuli for RC attachment ambiguity from MultiWho dataset. In the ambiguous condition, the absence of clear world knowledge and biases allows both DP1 *the son* and DP2 *the man* to serve as equally plausible referents for the RC. In the DP1-biased condition, the less plausible scenario of a woman having a beard leads to a preference for attaching the relative clause to DP1 *the son*. In contrast, the DP2-biased condition exploits the greater plausibility of a man having a beard, which favors attachment to DP2 *the man*.

for both biases, with higher accuracy for DP2, mirroring the patterns observed in humans in EN.

This study’s main contributions are threefold: (1) we create a new dataset for assessing the performance of LLMs in resolving ambiguities across multiple languages, (2) we describe the iterative process of linguist-LLM collaboration for generating such a dataset, and (3) we analyze how LLM performance compares to human processing patterns, identifying the limitations of LLMs and providing insights for future improvements.

2 MultiWho Dataset

We introduce MultiWho, a dataset designed to examine how humans and LLMs utilize world knowledge and biases to resolve syntactic ambiguities. MultiWho comprises a total of 1728 sentences, with 96 sets, spanning three categories (ambiguous, DP1-biased, and DP2-biased), in six languages, EN, ES, JP, KO, RU, and ZH. A sample English set is shown in Table 1.

Recently, LLMs’ fluency and instruction-following ability have prompted interest in human-LLM collaboration for dataset creation (Long et al., 2024; Wang et al., 2024). However, even advanced LLMs such as GPT-4o and Claude 3 Sonnet were unable to consistently generate sentences that met our criteria. This led us to develop a more flexible paradigm for specialized dataset creation, where human experts leverage LLMs for assistance and support rather than complete solutions. Our approach involved (i) brainstorming with LLMs to generate options meeting individual constraints, (ii) manually combining these elements to create sentences satisfying all criteria, (iii) using LLMs for validation and discussion throughout the process, (iv) employing LLMs’ fluency and cultural knowledge for initial drafts in multiple languages, and (v) validating and refining these drafts with help from native speakers.

This collaborative method allowed us to work ef-

ficiently while maintaining high standards of quality and consistency across diverse languages.

2.1 Dataset Creation

We started in EN with the construction [DP1 of DP2 RC], where the RC could syntactically attach to DP1 (HA) or DP2 (LA) (see Figure 1). A main subject and verb are added to complete the sentence. Importantly, we consistently use DP1 and DP2 to refer to the *structural positions in the syntax tree*, not the linear order of the words in the sentence, which varies by language: in postnominal languages (EN, ES, RU), the linear order *matches* the structural order: DP1 DP2 RC, while in postnominal languages (ZH, JP, KO), the linear order *is reversed*: RC DP2 DP1.

Sentences were designed to fall into one of three congruency categories:

- **ambiguous**, where no clear bias cue favors one attachment over the other;
- **DP1-biased**, where the bias content of the sentence favors attachment to DP1; and
- **DP2-biased**, where world knowledge and biases favor attachment to DP2.

Constraints for potential linguistic factors Bias is achieved by the RC applying differentially to the DPs (see Table 1). We further enforced the following constraints. **Morphosyntactic constraints:** To prevent biases from word length and animacy, DPs must be single words referring to humans. Additionally, the RC should exclude any morphosyntactic markers, such as grammatical number or gender, that would require agreement with either DP (e.g. “the **sons** of the doctor who **were** at home”). **Semantic relation between DPs:** The *DP1 of DP2* phrase must form a plausible relationship to exclude expressions like “the firefighter of the baby”. **Naturalness:** The subject and main verb of the sentence must be plausible in light of the rest of the sentence.

Categories of world knowledge and biases We employed various world knowledge and biases to guide RC attachment, such as **gender**: physiological differences (e.g., ‘giving birth’ for female bias) and gender roles (e.g., ‘participating in a beauty contest’ in RU); **age**: e.g., ‘the brother of the baby who was driving a car’, and **profession**: e.g., ‘the interpreter of the judge who was monolingual’. We recognize that some of these biases are based on societal stereotypes, but we utilize them as they reflect current linguistic associations and world knowledge (also see Section 7, Limitations).

2.2 Iterative Linguist-LLM Collaboration

While they can generate outputs quickly, the current generation of LLMs is unable to respect the large number of constraints simultaneously. For example, when reminded that “the firefighter of the baby” breaks the semantic relation constraint, it might generate “the toy car of the baby”, breaking the single word and animacy constraints. LLMs also often accept artificial-sounding sentences, and are hampered by their demonstrated yes-response bias (Dentella et al., 2023), particularly in languages other than English. The iterative process involved a linguist using the LLMs where possible while recognizing and working around its limitations. The linguist offset the weaknesses by verifying constraints, improving readability and naturalness (with help of native speakers), and correcting otherwise implausible sentences.

Although LLMs are unable to uphold all constraints at once, we had some success in simplifying the task by reducing the number of constraints. For example, we asked GPT-4o to verify the intended bias, providing it with only the DPs and the RCs and allowing it to focus only on the relations of DP1 and DP2 to RC. For example, given the ambiguous-condition sentence with DP1 ‘schoolfriend’, DP2 ‘preschooler’ and RC ‘was learning to use the potty’, GPT-4o flagged that “The term ‘schoolfriend’ suggests an older child who is likely past the potty training stage”, so we changed ‘schoolfriend’ to ‘brother’. However, the check’s accuracy was inconsistent. While it accepted the bias of ‘the neighbor of the boy who was a midwife’ in EN, it rejected it in ZH, claiming that “engaging in midwifery work can apply to any individual regardless of gender”³.

³While male midwives do exist, a boy in this role is clearly implausible. Here, GPT-4o’s debiasing mechanism may have overridden the practical implausibility of the situation.

2.3 Multilingual Sets

Starting from EN sentences, we used interactive processes with GPT-4o and Claude 3 to create initial versions in ES, JP, KO, RU, and ZH. However, this process quickly revealed that biases in EN do not always translate directly or maintain their relevance in other languages and cultures. As a result, we shifted from translation to adaptation, creating language-specific sentences that preserved the intended semantic relationships and ambiguities while respecting the linguistic and cultural norms of each language (see Appendix A).

2.4 Dataset Validation

In order to review and adjust the sentences to ensure accuracy, cultural appropriateness, and preservation of the intended ambiguities, native-speaking professional teachers, translators, and researchers volunteered or were hired, including one of the authors, a professional translator but non-native speaker of several languages. This process often resulted in sentences that diverged significantly from the original EN versions, tailored to each language’s specific linguistic and cultural context.

3 Experiment 1: Humans

To directly compare LLM and human performance in RC ambiguity resolution, we conducted a forced-choice experiment with human subjects using the EN dataset.

3.1 Participants and Procedure

Sixty-five native EN speakers (mean age: 31 years; age range: 18–80 years) were recruited through the online platform Prolific. Participation was restricted to individuals whose first language was EN and who were residing in the United States at the time of the experiment. No participants reported a clinical history of hearing or auditory processing issues, reading difficulties, or prior brain surgery. Participants were compensated at a rate of 15 USD per hour, upon successful completion of the task.

We conducted a forced-choice experiment using web-based survey platform PCIBex Farm (Zehr and Schwarz, 2018). As in example (4), after each sentence, participants were presented with a comprehension question and response options probing their interpretation of the sentence on a separate screen. The order of response options was counter-balanced.

- (4) I saw the daughter of the woman who bought a dress.
 ‘Who bought a dress?’
 1. the daughter 2. the woman

In this experiment, we tested 96 sets of EN items. These target items were distributed across three groups using a Latin square design, and fillers were included to maintain balance in the experimental conditions. The entire experiment took approximately 35 minutes to complete on average. The study protocol was approved by the Institutional Review Board (IRB).

3.2 Analysis

Following established psycholinguistic norms and previous studies on relative clause attachment ambiguities, an HA (non-local attachment) response rate above 0.5 is interpreted as a preference for HA, while a rate below 0.5 indicates a preference for LA (local attachment). Our analysis treated ambiguous and unambiguous cases separately. In the case of ambiguous sentences, we explored which attachment preference was favored, assessing whether there was a consistent inclination toward one attachment choice over the other. In the forced-choice task, participants were presented with two response options: low attachment (LA) and high attachment (HA). Consequently, the response rates for LA and HA sum to 1. In our analysis, responses were coded as binary values, with 0 representing LA and 1 representing HA. The HA response rate was calculated directly, and the LA response rate was obtained as 1HA rate. This allowed us to validate the observed LA preference in EN, which aligns with established findings from earlier psycholinguistic research (Cuetos and Mitchell, 1988; Mitchell et al., 2012). For the unambiguous cases, we examined whether the participants’ responses corresponded with the intended biases (DP1 or DP2), ensuring the cues were correctly followed. It is important to note that human responses confirmed that biases generally led to the intended interpretation—either DP1 or DP2 attachment—in all but one of the sets.

3.3 Results and Discussion

In ambiguous conditions, the HA response rate — calculated as the number of HA choices divided by the total number of responses — was 0.268. This rate, significantly below the 0.5 threshold, reaffirms the LA preference in EN, consistent with findings reported in previous studies (Cuetos and

Mitchell, 1988; Gilboy et al., 1995; Frazier and Clifton, 1997). When the participants’ answers matched the expected responses for DP1- and DP2-biased conditions, the bias-aligned answer rates revealed a clear difference. In the DP1-biased condition, the answer rate was 0.641, indicating participants generally aligned with the DP1 bias, but with noticeable variability ($sd = 0.47$). In contrast, the DP2-biased condition showed a much higher answer rate of 0.885, reflecting a stronger and more consistent adherence ($sd = 0.31$) to the intended DP2 bias.

Our results indicate a clear LA preference in ambiguous conditions, while in unambiguous cases, humans effectively use world knowledge and biases to resolve structural ambiguities. However, discrepancies emerged depending on the targeted DP position (DP1 or DP2). Discrepancies based on DP position (DP1 or DP2) likely stem from a natural tendency to adopt LA structures in "DP1 of DP2 RC" constructions, even when world knowledge favors HA. This default to LA aligns with prior research suggesting humans favor syntactic simplicity in initial parsing (Fodor, 1978; Meng and Bader, 2000, a.o.).

Our findings further illustrate humans’ flexibility in adjusting interpretations when faced with world knowledge inconsistencies, even in rare or unconventional scenarios. For instance, in the sentence “*The doctor met the son of the woman who had a beard*,” humans still showed an LA preference, interpreting the relative clause as modifying “*the woman*” despite the unusual nature of a woman having a beard. This flexibility can be attributed to humans’ ability to adaptively incorporate context and pragmatics when processing ambiguous sentences, enabling them to accommodate infrequent but plausible real-world scenarios. The cognitive mechanism known as good-enough parsing (Ferreira and Patson, 2007) suggests that humans do not always strive for fully accurate interpretations but rather settle for interpretations that are sufficient for comprehension, even if they require adjusting expectations based on uncommon world knowledge. This adaptability may also be linked to the fact that humans draw upon a vast reservoir of experiences and cultural knowledge, allowing them to entertain even improbable interpretations when syntactic ambiguity arises, thus showcasing their unique capacity for flexible language processing.

4 Experiment 2: LLMs

4.1 Large Language Models and Procedure

We tested three widely used LLMs, using both proprietary models such as **Claude 3.5 Sonnet** (claude-3-5-sonnet-20240620) and **GPT-4o** (gpt-4o-2024-05-13), and open-source models such as the instruction-tuned **Llama 3.1** (Meta-Llama-3.1-70B-Instruct).

We conducted a forced-choice experiment, similar to Experiment 1, consisting of three components: sentence, question, and answer choices. The prompt used for asking the response is “Answer with a single number, 1 or 2, without commentary”. This was translated into each target language (see Appendix ?? for the full prompt texts), with the respective version used for each language. The experiment was repeated three times, with each iteration varying the order of the presented choices. We tested three different configurations:

- **Linear** order: The choices were presented in the same sequence as they appeared in the sentence, e.g., in postnominal languages such as EN, this consistently meant starting with DP1, while in prenominal languages, such as KO, it corresponded to DP2.
- **Reversed** order: The choices were presented in the opposite sequence to how they appeared in the original sentence.
- **Random** order: The choices were presented in a randomized sequence, with the randomization kept consistent across all test items and LLMs to ensure comparability.

Prior research indicates that the order in which options are presented can affect responses in both humans and computational models (Tversky and Kahneman, 1981; Smith and Kosslyn, 2007; Zheng et al., 2023a). The use of multiple configurations aimed to control potential order effects that might influence the LLMs’ responses.

4.2 Analysis

In our analysis, responses were examined separately for ambiguous and unambiguous cases to capture potential differences in model behavior. For ambiguous cases, we investigated the model’s preference between the two possible interpretations, assessing whether there was a consistent tendency. For unambiguous cases, we evaluated whether the model’s responses aligned with the intended bias

(DP1 or DP2). Only responses selecting the provided answer choices were analyzed.

The prompt in all queries comprised a sentence, a question, two possible answers, and a request for a response consisting of only the number 1 or 2. The language-specific requests are shown below.

Lang	Simple response request
EN	Answer with a single number, 1 or 2, without commentary.
KO	설명 없이 숫자 1 또는 2 로만 답하세요.
SP	Responde con un solo número, 1 o 2, sin comentarios.
RU	Ответь одним числом, 1 или 2, без комментариев.
JP	説明なしで、1 か 2 の数字で答えてください。
CH	请用 1 或 2 的数字回答。不要评论。

GPT-4o and Claude 3 responded consistently with 1 or 2, as requested while all outliers came from Llama-3.1, indicating slightly more variability, especially in ZH and KO. In rare cases, Llama-3.1 responded with texts (e.g., " 여동생" [younger sister] or "答案: 1. 娃娃" [answer: 1. baby]). Of 15,552 responses, seven (0.045%) were failures where Llama-3.1 chose the main subject or gave invalid responses (e.g., "-1"). These instances occurred across different presentation orders and languages: two outliers in the linear-order presentation (from ES), two in the reversed order (1 from ZH, 1 from KO), and three in the random-order presentation (1 from JP, 1 from ES, 1 from KO).

To examine potential differences in LLMs’ behavior across languages, we conducted a statistical analysis using mixed-effects logistic regression (see Appendix C). All model results are reported as an average over three runs.

4.3 Results and Discussion

4.3.1 Human vs. LLMs in English

First, we compare the performance of humans and LLMs on two aspects: (1) HA response rates in ambiguous conditions, and (2) matched answer rates for DP1- and DP2-biased conditions.

Table 2 presents the HA response rates in ambiguous conditions. We observe that both humans and LLMs demonstrate a strong preference for LA (> 0.70), with humans exhibiting a slightly higher rate of HA responses compared to the LLMs.

Table 3 presents the matched answer rates for DP1- and DP2-biased conditions, where the sentence bias promotes attachment to either DP1 or DP2. In these unambiguous cases, all models performed exceptionally well, particularly in the

Congruency	Human	Claude 3	GPT4o	Llama 3.1
ambiguous	0.268	0.154	0.197	0.157

Table 2: Comparison of human and model performance on HA answer rates in ambiguous conditions in EN.

Congruency	Human	Claude 3	GPT4o	Llama 3.1
DP1	0.641	0.726	0.743	0.575
DP2	0.885	1.000	0.978	0.954

Table 3: Comparison of human and model matched answer rates for DP1 and DP2 biased conditions in EN.

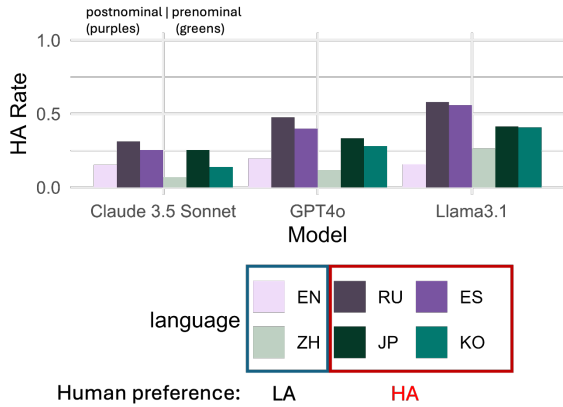


Figure 3: High attachment (HA) response rates in **ambiguous** conditions (Attachment Preference)

DP2-biased condition, where they demonstrated near-perfect accuracy and outperformed human participants. For the DP1-biased condition, GPT-4o (0.743) and Claude (0.726) still outperformed human participants (0.641), indicating a stronger alignment with the provided bias. Llama 3.1 exhibited the lowest matched answer rate (0.575) for DP1-biased sentences.⁴

While LLMs demonstrate a robust ability to integrate world knowledge and explicit biases, they often exhibit rigidity in their interpretations, frequently reinforcing existing social stereotypes. In contrast, human participants display a notable flexibility, adapting their interpretations to align with evolving social norms and contextual subtleties.

4.3.2 Multilingual Syntactic Attachment Preferences in Ambiguous Cases

Next, we examine the HA vs. LA preferences of LLMs under ambiguous conditions in six lan-

⁴We thank the anonymous reviewer for suggesting an additional analysis. To examine how much variance in human responses can be explained by LLM responses, we conducted further analysis. Please see Appendix B for details.

guages. The results are summarized in Figure 3. According to previous psycholinguistic studies, EN and ZH speakers demonstrated an LA preference, whereas JP, KO, RU, and ES speakers displayed an HA preference (Cuetos and Mitchell, 1988; Lee, 2018, 2021; Mitchell et al., 2012; Sekerina, 2003; Shen, 2006, a.o.). In contrast, in our study, LLMs exhibited an overall **LA preference across all languages**, regardless of the attachment tendencies reported in previous psycholinguistic studies (detailed results in Table 5 in Appendix C). In addition to not reflecting language-specific preferences, the models also do not exhibit a specific pattern based on syntactic structures, such as the difference between prenominal and postnominal RC languages.

This general tendency towards LA suggests that LLMs could be defaulting to simpler syntactic structures when resolving ambiguities rather than adapting to language-specific syntactic rules. One possible explanation is that the models may not have fully adapted to the linguistic structure of these high-attachment languages or it may be over-relying on its training data or on an innate bias learned from predominantly low-attachment languages, like EN. The one notable exception is Llama 3.1, which exhibited a slightly more HA preference in Russian, aligning with the psycholinguistic findings that often report an HA bias in human processing for this language.

4.3.3 Multilingual World Knowledge and Bias Alignment in Unambiguous Cases

Figures 4a and 4b show that under unambiguous conditions, LLMs across various languages exhibit high alignment with world knowledge and bias cues. This suggests that LLMs show strong sensitivity to world knowledge, effectively incorporating it when explicit biases are present, achieving high accuracy in interpreting syntactically unambiguous sentences (see Appendix D). However, similar to humans in EN, there was a more pronounced alignment with the intended biases in DP2-biased conditions compared to DP1-biased conditions (p close to 0), with Llama 3.1 showing exceptions in Russian and Spanish.

4.3.4 Influence of Presentation Order of Answer Choices

We analyzed the results separately based on presentation order. The models generally responded consistently regardless of presentation order of response choices, under both ambiguous (Table 4)

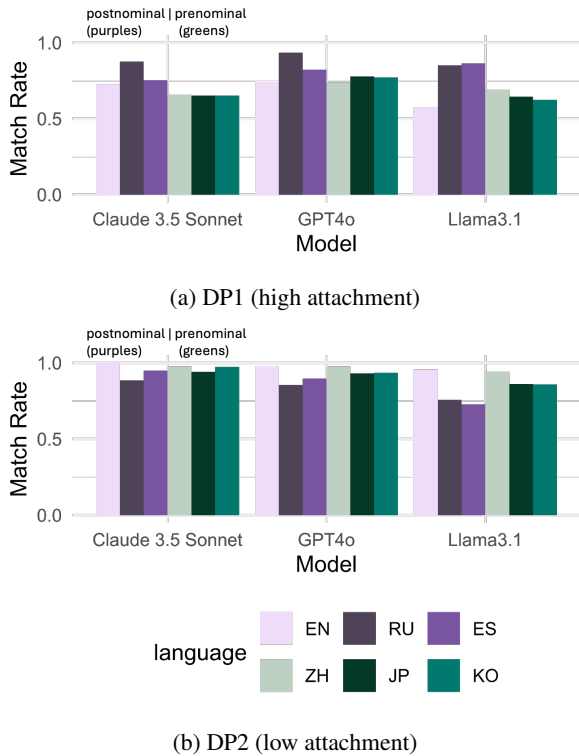


Figure 4: Average matched responses with the given world knowledge and bias toward DP1 and DP2 (high and low attachment) in **unambiguous** conditions in English

and unambiguous cases (Appendix E). In the ambiguous cases, the models consistently showed very strong preference towards LA response rather than the ordering of answer choices, suggesting that the model is prioritizing syntactic parsing strategies over surface-level presentation biases. The model’s bias toward low attachment may lead to poorer performance in languages with high attachment preferences, particularly in tasks involving syntactic parsing, translation, or question answering. Further research is needed to ensure that the models properly adapt to the syntactic preferences of different languages.

5 General Discussion

The findings in Sections 3.3 and 4.3 can be summarized as follows: While LLMs like GPT-4o and Claude 3 Sonnet exhibit high accuracy in unambiguous conditions by effectively leveraging world knowledge to resolve syntactic ambiguities, this success may reflect a reliance on stereotypes and explicit biases embedded in their training data (Bolukbasi et al., 2016; Sheng et al., 2019; Lucy and Bamman, 2021), leading to rigidity in interpretations and reinforcing existing biases. In contrast, humans

Language	Model	linear	reverse	random
EN	Claude 3	0.084	0.232	0.147
ZH	Claude 3	0.083	0.042	0.083
JP	Claude 3	0.375	0.094	0.292
KO	Claude 3	0.219	0.031	0.167
RU	Claude 3	0.240	0.396	0.302
ES	Claude 3	0.146	0.375	0.240
EN	GPT-4o	0.147	0.253	0.189
ZH	GPT-4o	0.115	0.083	0.156
JP	GPT-4o	0.281	0.396	0.323
KO	GPT-4o	0.302	0.281	0.260
RU	GPT-4o	0.385	0.552	0.490
ES	GPT-4o	0.323	0.479	0.396
EN	Llama-3.1	0.105	0.200	0.168
ZH	Llama-3.1	0.292	0.229	0.271
JP	Llama-3.1	0.458	0.354	0.432
KO	Llama-3.1	0.229	0.552	0.442
RU	Llama-3.1	0.573	0.594	0.573
ES	Llama-3.1	0.495	0.625	0.558

Table 4: Attachment preferences under varying response orders. The grayed rows (EN and ZH) indicate where humans typically show LA preference whereas all other languages show HA preferences. The cells in bold indicate where the human and LLM preferences align.

demonstrate greater flexibility by overriding syntactic biases when world knowledge conflicts with their default attachment preferences, considering rare but plausible interpretations through pragmatic reasoning, context, and social norms. Moreover, the LLMs’ overall preference for LA indicates insensitivity to syntactic differences between prenominal and postnominal structures and an inability to adapt to language-specific attachment preferences.

These findings offer important insights into how LLMs and humans process syntactic ambiguities and incorporate world knowledge, carrying substantial implications for both the advancement of LLM technology and the understanding of human language processing. While LLMs are becoming increasingly sophisticated, they do not yet fully capture the intricate interplay of syntactic, semantic, and pragmatic factors that characterize human language processing. Human comprehension involves dynamic cognitive strategies that adapt to evolving social norms, contextual cues, and world knowledge. Integrating insights from psycholinguistics and cognitive science is essential for developing LLMs that replicate not just outputs but also the underlying cognitive mechanisms of human language processing. Future models would benefit from this integration, enabling them to handle ambiguities and culturally specific knowledge in a manner that reflects genuine human adaptability and flexibility.

6 Related Work

Syntactic Ambiguity Early efforts used classifiers to reflect human preferences in pronoun resolution by leveraging coreference features (Seminck and Amsili, 2017). Other studies focused on RC attachment preferences, revealing mixed outcomes. For example, English RNNs learned low attachment preferences but over-generalized, while Spanish RNNs struggled to learn human-typical high attachment preferences unless biases in the training data were manipulated to be balanced (Davis and van Schijndel, 2020). In contrast, the MultiWho sentences are designed to always be semantically coherent, full sentences, because they must be realistic objects of human annotation. BERT-based parsers performed poorly on Dutch RC ambiguity unless bias-correcting priming was applied (Wijnholds and Moortgat, 2023). LLMs have been shown to diverge from human-like prepositional phrase attachment preferences (Cai et al., 2024).

Syntactic Ambiguity Datasets A tradeoff exists between small, high-quality datasets and large, synthetic ones. For instance, BLiMP (Warstadt et al., 2020), a set of syntactic minimal pairs generated automatically from templates, has inspired similar datasets in other languages like JBLiMP (Someya and Oseki, 2023) (a Japanese dataset of 331 minimal pairs based directly on example sentences taken from theoretical linguistics publications) and CLiMP (Xiang et al., 2021) (a Chinese version generated synthetically from BLiMP translations, though it faced quality issues (Song et al., 2022), illustrating the limitations of automatic, template-based generation and translation of linguistic test sets). SLING (Song et al., 2022) addressed these difficulties by utilizing a TreeBank to extract lexically diverse and ecologically valid sentences.

Prompt Design & Survey Bias LLMs are sensitive to minor prompt changes (Zhao et al., 2021), with survey question modifications shifting model responses away from known human biases (Tjuatja et al., 2024). Additionally, order bias (or position bias) can be introduced by the order of possible answers (Wang et al., 2023; Zheng et al., 2023b; Herr et al., 2024; Zheng et al., 2023a; Shi et al., 2024; Li et al., 2023; Liusie et al., 2024). To mitigate order bias, MultiWho experiments systematically examine all answer order cases to assess the impact of order on the results.

7 Conclusion

In summary, while LLMs have made considerable strides in handling syntactic ambiguities and incorporating world knowledge, they still exhibit limitations compared to human language processing. Addressing these gaps will require a concerted effort to train LLMs with more diverse and nuanced data, enhancing their adaptability and integrating insights from human cognition to create models capable of truly human-like language comprehension. Only by incorporating more context-sensitive, pragmatically rich, and culturally diverse training data can we develop LLMs that approach the depth and flexibility of human language understanding. By doing so, we will not only advance the capabilities of LLMs but also gain deeper insights into the complexities of human language processing, bridging the gap between artificial and human intelligence.

Limitations

Sinha et al. (2022) point out that language models' syntactic acceptability judgments improve in accuracy with longer context. Our experiments contain only a single sentence, and it is possible that longer context would improve the models' alignment with human judgment.

We sought to create an interdisciplinary team to ground our research well in both linguistics and natural language processing. Along the way, we discovered several hurdles involved in this type of work. We, as linguists and computer scientists, speak different languages⁵. We look at data differently, use different terms, and consider different outcomes significant in different ways. Each of these differences posed a barrier, and overcoming these barriers grew our capabilities as interdisciplinary researchers. We have learned much from each other, and would encourage others to take on similar challenges.

Acknowledgments

We are thankful to the anonymous reviewers for their helpful feedback.

References

Dale J Barr, Roger Levy, Christoph Scheepers, and Harry J Tily. 2013. Random effects structure for

⁵Though we also speak several different natural languages, that's not quite what we mean.

- confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language*, 68(3):255–278.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to home-maker? debiasing word embeddings. *Advances in neural information processing systems*, 29.
- Zhenguang G. Cai, Xufeng Duan, David A. Haslett, Shuqi Wang, and Martin J. Pickering. 2024. [Do large language models resemble humans in language use?](#) *Preprint*, arXiv:2303.08014.
- Charles Clifton Jr, Matthew J Traxler, Mohamed Taha Mohamed, Rihana S Williams, Robin K Morris, and Keith Rayner. 2003. The use of thematic role information in parsing: Syntactic processing autonomy revisited. *Journal of Memory and Language*, 49(3):317–334.
- Fernando Cuertos and Don C Mitchell. 1988. Cross-linguistic differences in parsing: Restrictions on the use of the late closure strategy in spanish. *Cognition*, 30(1):73–105.
- Forrest Davis and Marten van Schijndel. 2020. [Recurrent neural network language models always learn English-like relative clause attachment](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1979–1990, Online. Association for Computational Linguistics.
- Vittoria Dentella, Fritz Günther, and Evelina Leivada. 2023. [Systematic testing of three language models reveals low language accuracy, absence of response stability, and a yes-response bias](#). *Proceedings of the National Academy of Sciences*, 120(51):e2309583120.
- Fernanda Ferreira and Nikole D Patson. 2007. The ‘good enough’ approach to language comprehension. *Language and linguistics compass*, 1(1-2):71–83.
- Janet Dean Fodor. 1978. Parsing strategies and constraints on transformations. *Linguistic inquiry*, 9(3):427–473.
- Janet Dean Fodor. 1998. Learning to parse? *Journal of psycholinguistic research*, 27:285–319.
- Lyn Frazier and Charles Clifton. 1997. Construal: Overview, motivation, and some new evidence. *Journal of Psycholinguistic Research*, 26:277–295.
- Matthieu Futeral, Cordelia Schmid, Ivan Laptev, Benoît Sagot, and Rachel Bawden. 2023. [Tackling ambiguity with images: Improved multimodal machine translation and contrastive evaluation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5394–5413, Toronto, Canada. Association for Computational Linguistics.
- Elizabeth Gilboy, Josep-Maria Sopena, Charles Clifton Jr, and Lyn Frazier. 1995. Argument structure and association preferences in spanish and english complex nps. *Cognition*, 54(2):131–167.
- Nino Grillo and João Costa. 2014. A novel argument for the universality of parsing principles. *Cognition*, 133(1):156–187.
- Ali Hatami, Paul Buitelaar, and Mihael Arcan. 2022. [Analysing the correlation between lexical ambiguity and translation quality in a multimodal setting using WordNet](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Student Research Workshop*, pages 89–95, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.
- Yangfan He, Yuxuan Bai, and Tianyu Shi. 2024. Enhancing intent understanding for ambiguous prompt: A human-machine co-adaption strategy. In *ICML 2024 Workshop on Models of Human Feedback for AI Alignment*.
- B Hemforth, L Konieczny, and C Scheepers. 1996. Syntactic and anaphoric processes in modifier attachment. In *The 9th Annual CUNY Conference on Human Sentence Processing*, pages 21–23.
- Nathan Herr, Fernando Acero, Roberta Raileanu, María Pérez-Ortiz, and Zhibin Li. 2024. [Are large language models strategic decision makers? a study of performance and bias in two-player non-zero-sum games](#). *Preprint*, arXiv:2407.04467.
- Istvan Kecskes. 2023. The interplay of linguistic, conceptual, and encyclopedic knowledge in meaning construction and comprehension. In *The Cambridge Handbook of Language in Context*, pages 268–288. Cambridge University Press.
- So Young Lee. 2018. A minimalist parsing account of attachment ambiguity in English and Korean. *Journal of Cognitive Science*, 19(3):291–329.
- So Young Lee. 2021. The effect of honorific affix on pro-processing of an attachment ambiguity. *Japanese/Korean Linguistics*, 28:1–10.
- Junlong Li, Shichao Sun, Weizhe Yuan, Run-Ze Fan, Hai Zhao, and Pengfei Liu. 2023. [Generative judge for evaluating alignment](#). *Preprint*, arXiv:2310.05470.
- Adian Liusie, Potsawee Manakul, and Mark Gales. 2024. [LLM comparative assessment: Zero-shot NLG evaluation through pairwise comparisons using large language models](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 139–151, St. Julian’s, Malta. Association for Computational Linguistics.

- Lin Long, Rui Wang, Ruixuan Xiao, Junbo Zhao, Xiao Ding, Gang Chen, and Haobo Wang. 2024. [On llms-driven synthetic data generation, curation, and evaluation: A survey](#). *Preprint*, arXiv:2406.15126.
- Li Lucy and David Bamman. 2021. Gender and representation bias in gpt-3 generated stories. In *Proceedings of the third workshop on narrative understanding*, pages 48–55.
- Ninareh Mehrabi, Palash Goyal, Apurv Verma, Jwala Dhamala, Varun Kumar, Qian Hu, Kai-Wei Chang, Richard Zemel, Aram Galstyan, and Rahul Gupta. 2023. [Resolving ambiguities in text-to-image generative models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14367–14388, Toronto, Canada. Association for Computational Linguistics.
- Michael Meng and Markus Bader. 2000. Ungrammaticality detection and garden path strength: Evidence for serial parsing. *Language and Cognitive processes*, 15(6):615–666.
- Don C Mitchell, Fernando Cuetos, and Daniel Zagar. 2012. Reading in different languages: Is there a universal mechanism for parsing sentences? In *Comprehension processes in reading*, pages 285–302. Routledge.
- Abhijnan Nath, Huma Jamil, Shafiuddin Rehan Ahmed, George Arthur Baker, Rahul Ghosh, James H. Martin, Nathaniel Blanchard, and Nikhil Krishnaswamy. 2024. [Multimodal cross-document event coreference resolution using linear semantic transfer and mixed-modality ensembles](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 11901–11916, Torino, Italia. ELRA and ICCL.
- Ayana Niwa and Hayate Iso. 2024. [Ambignlg: Addressing task ambiguity in instruction for nlg](#). *Preprint*, arXiv:2402.17717.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kopic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qim-

- ing Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Asad Sayeed and Alessandra Zarcone. Explicit world-knowledge and distributional semantic representations.
- Irina Sekerina. 2003. The late closure principle in processing of ambiguous russian sentences. In *The Proceedings of the Second European Conference on Formal Description of Slavic Languages*. Universität Potsdam, Germany, pages 1–17.
- Olga Seminck and Pascal Amsili. 2017. [A computational model of human preferences for pronoun resolution](#). In *Proceedings of the Student Research Workshop at the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 53–63, Valencia, Spain. Association for Computational Linguistics.
- Xingjia Shen. 2006. Late assignment of syntax theory: evidence from chinese and english.
- Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. *arXiv preprint arXiv:1909.01326*.
- Lin Shi, Weicheng Ma, and Soroush Vosoughi. 2024. Judging the judges: A systematic investigation of position bias in pairwise comparative assessments by llms. *arXiv preprint arXiv:2406.07791*.
- Koustuv Sinha, Jon Gauthier, Aaron Mueller, Kanishka Misra, Keren Fuentes, Roger Levy, and Adina Williams. 2022. [Language model acceptability judgments are not always robust to context](#). *Preprint*, arXiv:2212.08979.
- Edward E Smith and Stephen Michael Kosslyn. 2007. Cognitive psychology: Mind and brain. (*No Title*).
- Taiga Someya and Yohei Oseki. 2023. [JBLiMP: Japanese benchmark of linguistic minimal pairs](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1581–1594, Dubrovnik, Croatia. Association for Computational Linguistics.
- Yixiao Song, Kalpesh Krishna, Rajesh Bhatt, and Mohit Iyyer. 2022. [SLING: Sino linguistic evaluation of large language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4606–4634, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Lindia Tjuatja, Valerie Chen, Tongshuang Wu, Ameet Talwalkar, and Graham Neubig. 2024. Do llms exhibit human-like response biases? a case study in survey design. *Transactions of the Association for Computational Linguistics*, 12:1011–1026.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *Preprint*, arXiv:2302.13971.
- Hélène Tran, Issam Fali, Xavier Goblet, and Engelbert Mephu Nguifo. 2022. [Do multimodal emotion recognition models tackle ambiguity?](#) In *Proceedings of the 2nd Workshop on People in Vision, Language, and the Mind*, pages 6–11, Marseille, France. European Language Resources Association.
- Matthew J Traxler, Martin J Pickering, and Charles Clifton Jr. 1998. Adjunct attachment is not a form of lexical ambiguity resolution. *Journal of memory and language*, 39(4):558–592.
- Amos Tversky and Daniel Kahneman. 1981. The framing of decisions and the psychology of choice. *science*, 211(4481):453–458.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023. [Large language models are not fair evaluators](#). *Preprint*, arXiv:2305.17926.
- Xinru Wang, Hannah Kim, Sajjadur Rahman, Kushan Mitra, and Zhengjie Miao. 2024. [Human-llm collaborative annotation through effective verification of llm labels](#). In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI ’24, New York, NY, USA. Association for Computing Machinery.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Gijs Wijnholds and Michael Moortgat. 2023. [Structural ambiguity and its disambiguation in language model based parsers: the case of dutch clause relativization](#). *Preprint*, arXiv:2305.14917.
- Beilei Xiang, Changbing Yang, Yu Li, Alex Warstadt, and Katharina Kann. 2021. [CLiMP: A benchmark for Chinese language model evaluation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2784–2790, Online. Association for Computational Linguistics.
- Yuewei Yuan, Chaitanya Malaviya, and Mark Yatskar. 2023. [AmbiCoref: Evaluating human and model sensitivity to ambiguous coreference](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1023–1030, Dubrovnik, Croatia. Association for Computational Linguistics.
- Jeremy Zehr and Florian Schwarz. 2018. Penncontroller for internet based experiments (ibex).

Tony Z. Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. [Calibrate before use: Improving few-shot performance of language models](#). Preprint, arXiv:2102.09690.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023a. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023b. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). Preprint, arXiv:2306.05685.

A Cross-Linguistic Adaptations and Challenges

Creating a multilingual dataset for syntactic ambiguity resolution posed numerous challenges due to the diverse linguistic and cultural features of the target languages. This section highlights key adaptations and considerations encountered while adapting our dataset from EN to ES, JP, KO, RU, and ZH.

Grammatical Considerations

In Spanish and Russian, both languages with grammatical gender, we ensured that relative clauses remained syntactically gender-neutral to maintain the required ambiguity. For example, in Spanish, "pregnant" was rephrased in one set as "in a state of pregnancy" (*quedó en estado de embarazo*, #94-96⁶) to avoid gender-specific adjectives. We carefully constructed sentences to avoid adjectives that modify the RC subject, which would normally provide a syntactic clue about the attachment. In Russian, we avoided past-tense verbs when one DP is feminine and the other is masculine, as past-tense verbs are gender-marked in Russian and would provide a syntactic disambiguation cue. For professions typically associated with a specific gender, we used gender-neutral alternatives. For instance, "wet nurse" in "the daughter of the man who became a wet nurse" (#415-417) was translated to lactation professional (*profesional de lactancia*, a gender-neutral term in Spanish; #127-129).

Cultural Adaptations

Localization in the form of cultural adaptations were necessary to ensure the stimuli were appro-

priate and meaningful in each language context. "Choirmaster" sounds anomalous when translated into Chinese, so was replaced with "道" *daozhang* (Daoist Temple Master, #88-90), aligning with local cultural contexts. In Spanish, "quinceañera" was used to create an age-differentiating relative clause, leveraging a culture-specific celebration (#1324-1326). The English "tooth fairy" (#442-444) was adapted to "ratoncito" in Spanish (#1306-1308) and to "包" *hongbao* (red envelope for gifting money to children at festivals) in Chinese (#151-153), reflecting culture-specific concepts. Names were also adapted to be culturally appropriate, changing "John" to "Juan" in Spanish (#1162, etc.) or "Taro" in Japanese (#640, etc.), for example.

Language-specific Phraseology

Each language presented unique challenges in expressing certain concepts while maintaining ambiguity. Succinct idiomatic expressions like "to father a child" (sets #406-408) required creative adaptations, such as "*que embarazó a una mujer*" (who got a woman pregnant) in Spanish (#1270-1272). Register contrasts, exemplified by "The student met the preschooler of the boss who was learning to use the potty" (set #452), posed translation difficulties. Fixed expressions like the Japanese "妻せずに過ごした" (#757-759), meaning "was a bachelor" but literally "spent time without entering the state of having a wife," needed careful handling to preserve semantic richness and structural ambiguity. These challenges highlight the importance of understanding each language's preferred phraseology.

The above adaptations highlight the complexity of creating equivalent stimuli across languages, and make clear that adaptations taking careful account of grammatical, cultural, and pragmatic equivalence are more appropriate than literal translations.

B Exploring Variance in Human Responses via LLM Predictions

We conducted additional analyses to explore how much variance in human responses can be explained by LLM responses.

To assess the relationship between LLM predictions and human behavior, we analyzed items (N=18) where all three LLMs provided incongruent responses. These items were then categorized based on human response patterns as follows:

Highly incongruent: 10–18 participants agreed on the incongruent answer (11/18 items). Moder-

⁶Numbers in this section refer to the index numbers of the data items in the MultiWho dataset

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	-2.04959	0.15231	-13.456	< 2e-16
languageEN	0.16083	0.14084	1.142	0.253
languageJP	1.23326	0.12964	9.513	< 2e-16
languageKO	0.89965	0.13161	6.836	8.16e-12
languageRU	1.86667	0.12861	14.514	< 2e-16
languageSP	1.60480	0.12869	12.470	< 2e-16
ModelLlama3.1	0.54073	0.08088	6.685	2.30e-11
ModelSonnet	-0.68742	0.08857	-7.761	8.41e-15

Table 5: Statistical Results for Ambiguous Conditions

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	1.41169	0.14735	9.580	< 2e-16
languageEN	-0.10699	0.11431	-0.936	0.349
languageJP	-0.02991	0.11442	-0.261	0.794
languageKO	-0.08231	0.11404	-0.722	0.470
languageRU	1.42448	0.13989	10.183	< 2e-16
languageSP	0.75876	0.12423	6.107	1.01e-09
ModelLlama3.1	-0.59088	0.08815	-6.703	2.04e-11
ModelSonnet	-0.52534	0.08844	-5.940	2.84e-09

Table 6: Statistical Results for DP1-biased Conditions

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	4.3175	0.2586	16.693	< 2e-16
languageEN	0.5305	0.2985	1.778	0.075473
languageJP	-1.0867	0.2255	-4.819	1.44e-06
languageKO	-0.9333	0.2291	-4.074	4.63e-05
languageRU	-1.9654	0.2133	-9.216	< 2e-16
languageSP	-1.7224	0.2156	-7.991	1.34e-15
ModelLlama3.1	-1.0036	0.1244	-8.067	7.18e-16
ModelSonnet	0.5292	0.1539	3.440	0.000582

Table 7: Statistical Results for DP2-biased Conditions

ately incongruent: 1–9 participants agreed on the incongruent answer (6/18 items). No incongruence: No participants provided the incongruent answer (1/18 item). The alignment between LLM and humans regarding incongruency suggests LLMs can help identify ambiguous or challenging stimuli likely to result in variable human responses. This can make LLMs a useful tool for psycholinguistic research, particularly for selecting stimuli for further testing.

C Statistical Results

We conducted a statistical analysis using mixed-effects logistic regression. The primary model included LLMs and languages as fixed effects, while random intercepts were assigned to items. Following best practices, we initially employed the maximal random effects structure and progressively simplified it until model convergence was achieved (Barr et al., 2013). The analysis produced coefficients, standard errors, Z-scores, and p-values for

each fixed effect and interaction, with statistical significance determined at a threshold of 0.05.

Tables 5, 6, and 7 present the statistical results for ambiguous, DP1-biased, and DP2-biased conditions, respectively.

D More Details in Human Results

The summary of the human results indicates that, among 96 sets, human participants provided incongruent answers in 91 sets. The number of participants per item ranged from 1 to 18. For example, in Set 2, the sentence was:

"Mr. Johnson visited the baby of the mother who was in a stroller."

Question: Who was in a stroller?

Expected Answer: The baby.

In Set 2, for instance, the expected answer "the baby" aligns with world knowledge and typical bias, as it is more plausible for a baby to be in a stroller than for a mother. This frequency-driven plausibility makes "the baby" the favored interpretation for the question "Who was in a stroller?" However, human participants frequently chose the incongruent response "the mother," likely due to a preference for local attachment, prioritizing the noun phrase closer to the relative clause. These incongruent human responses reflect a flexibility in interpretation that overrides world knowledge and bias in certain contexts.

E Presentation Order Results

Table 8 presents the detailed results obtained using different presentation of responses - linear, reversed, and random.

language	Model	congruency	linear	reverse	random
EN	GPT4o	DP1	0.663	0.831	0.736
ZN	GPT4o	DP1	0.802	0.677	0.739
JP	GPT4o	DP1	0.791	0.770	0.770
KO	GPT4o	DP1	0.781	0.739	0.791
RU	GPT4o	DP1	0.937	0.937	0.927
ES	GPT4o	DP1	0.760	0.906	0.802
EN	Llama3.1	DP1	0.526	0.631	0.568
ZN	Llama3.1	DP1	0.718	0.642	0.718
JP	Llama3.1	DP1	0.708	0.614	0.614
KO	Llama3.1	DP1	0.479	0.770	0.625
RU	Llama3.1	DP1	0.843	0.864	0.843
ES	Llama3.1	DP1	0.810	0.916	0.864
EN	Sonnet	DP1	0.694	0.757	0.726
ZN	Sonnet	DP1	0.760	0.562	0.656
JP	Sonnet	DP1	0.760	0.520	0.677
KO	Sonnet	DP1	0.739	0.562	0.656
RU	Sonnet	DP1	0.843	0.906	0.875
ES	Sonnet	DP1	0.697	0.812	0.750
CH	GPT4o	DP2	0.989	0.968	0.968
EN	GPT4o	DP2	0.989	0.968	0.978
JP	GPT4o	DP2	0.947	0.916	0.927
KO	GPT4o	DP2	0.927	0.937	0.937
RU	GPT4o	DP2	0.864	0.833	0.864
SP	GPT4o	DP2	0.947	0.843	0.895
CH	Llama3.1	DP2	0.916	0.947	0.958
EN	Llama3.1	DP2	0.968	0.936	0.957
JP	Llama3.1	DP2	0.885	0.854	0.843
KO	Llama3.1	DP2	0.916	0.757	0.895
RU	Llama3.1	DP2	0.802	0.708	0.760
SP	Llama3.1	DP2	0.812	0.677	0.697
CH	Sonnet	DP2	0.979	0.979	0.968
EN	Sonnet	DP2	1.000	1.000	1.000
JP	Sonnet	DP2	0.916	0.968	0.937
KO	Sonnet	DP2	0.968	0.979	0.968
RU	Sonnet	DP2	0.906	0.864	0.885
SP	Sonnet	DP2	0.958	0.937	0.947

Table 8: Matched answer rates for DP1- and DP2-biased conditions obtained using different response orders